

Internet Architecture Board (IAB)
Request for Comments: 7094
Category: Informational
ISSN: 2070-1721

D. McPherson
Verisign, Inc.
D. Oran
Cisco Systems
D. Thaler
Microsoft Corporation
E. Osterweil
Verisign, Inc.
January 2014

Architectural Considerations of IP Anycast

Abstract

This memo discusses architectural implications of IP anycast and provides some historical analysis of anycast use by various IETF protocols.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Architecture Board (IAB) and represents information that the IAB has deemed valuable to provide for permanent record. It represents the consensus of the Internet Architecture Board (IAB). Documents approved for publication by the IAB are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7094>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Overview	2
2. Background	3
2.1. Anycast History	3
2.2. Anycast in IPv6	6
2.3. DNS Anycast	6
2.4. BCP 126 on Operation of Anycast Services	8
3. Principles	8
3.1. Layering and Resiliency	8
3.2. Anycast Addresses as Destinations	9
3.3. Anycast Addresses as Sources	10
3.4. Service Discovery	10
4. Analysis	11
4.1. Regarding Widespread Anycast Use	11
4.2. Transport Implications	11
4.3. Stateful Firewalls, Middleboxes, and Anycast	12
4.4. Security Considerations	12
4.5. Deployment Considerations	15
5. Conclusions	16
6. Acknowledgements	16
7. Informative References	16
Appendix A. IAB Members at the Time of Approval	21

1. Overview

IP anycast is a technique with a long legacy and interesting engineering challenges. However, at its core, it is a relatively simple concept. As described in BCP 126 [RFC4786], the general form of IP anycast is the practice of making a particular Service Address available in multiple, discrete, autonomous locations, such that datagrams sent are routed to one of several available locations.

IP anycast is used for at least one critical Internet service: that of the Domain Name System [RFC1035] root servers. By late 2007, at least 10 of the 13 root name servers were already using IP anycast [RSSAC29]. Use of IP anycast is growing for other applications as well. It has been deployed for over a decade for DNS resolution services and is currently used by several DNS Top Level Domain (TLD) operators. IP anycast is also used for other services in operational environments, including Network Time Protocol (NTP) [RFC5905] services.

Anycast addresses are syntactically indistinguishable from unicast addresses. Anycast addressing is equivalent to that of unicast in multiple locations. Destination-based routing does best-effort delivery of a packet to one interface among the set of interfaces asserting reachability for the address. The expectation of delivery

is to the "closest" instance as determined by unicast routing topology metric(s), and there is also a possibility that various load-balancing techniques (e.g., per-packet, per-microflow) may be used among multiple equal-cost routes to distribute load for an anycasted prefix.

Unlike IP unicast, it is not considered an error to assert the same anycast address on multiple interfaces within the same or multiple systems.

When IP anycast is employed, many pitfalls and subtleties exist with applications and transports as well as for routing configuration and operation. In this document, we aim to capture many of the architectural implications of IP anycast.

BCP 126 [RFC4786] discusses several different deployment models with IP anycast. Two additional distinctions beyond that document involve "off-link anycast" and "on-link anycast". "Off-link anycast" takes advantage of routing protocol preferences and the IP hop-by-hop destination-based forwarding paradigm in order to direct packets to the "closest" destination. This is the traditional method of anycast largely considered in BCP 126 [RFC4786] and can be used for IPv4 and IPv6. "On-link anycast" is the formal support of anycast in the address resolution (duplicate address detection) protocol and is only standardized for IPv6, with the introduction of designated anycast addresses on the anycasted hosts, and the Override flag in Neighbor Discovery (ND) Neighbor Advertisements (NAs) [RFC4861]. There is no standardized mechanism for this in IPv4.

2. Background

As of this writing, the term "anycast" appears in 176 RFCs and 144 active Internet-Drafts. The following sections capture some of the key appearances and discussion of anycasting within the IETF over the years.

2.1. Anycast History

The first formal specification of anycast was provided in "Host Anycasting Service" [RFC1546]. The authors of this document did a good job of capturing most of the issues that exist with IP anycast today.

One of the first documented uses of anycast was in 1994 for a "Video Registry" experiment [IMR9401]. In the experiment, a UDP query was transmitted to an anycasted address to locate the topologically closest "supposedly equivalent network resource":

A video resource (for example, a catalog server that lists available video clips) sends an anycast UDP datagram to locate the nearest video registry. At most one registry responds with a unicast UDP datagram containing the registry's IP address. Said resource then opens a TCP connection to that [the received registry address] address and sends a request to register itself. Every 5 minutes or so, each registry multicasts to all other registries all of the resources it knows from local registration requests. It also immediately announces newly registered resources. Remotely registered resources not heard about for 20 minutes are dropped.

There is also discussion that ISPs began using anycast for DNS resolution services around the same time, although no public references to support this are available.

In 1997, the IAB clarified that IPv4 anycast addresses were pure "locators" and could never serve as "identifiers" of hosts or interfaces [RFC2101].

In 1998, the IAB conducted a routing workshop [RFC2902]. Of the conclusions and output action items from the report, an Anycast section is contained in Section 2.10.3. Specifically called out is the need to describe the advantages and disadvantages of anycast and the belief that local-scoped well-known anycast addresses will be useful to some applications. In the subsequent section, an action item was outlined that suggested a BOF should be held to plan work on anycast, and if a working group forms, a paper on the advantages and the disadvantages of anycast should be included as part of the charter.

As a result of the recommendation in [RFC2902], an Anycast BOF [ANYCASTBOF] was held at IETF 46 in November of 1999. A number of uses for anycast were discussed. No firm conclusion was reached regarding use of TCP with anycasted services. However, it was observed that anycasting was useful for DNS, although it did introduce some new complexities. The use of global anycast was not expected to scale (see Section 4.1 below for more discussion) and, hence, was expected to be limited to a small number of key uses.

In 2001, the Multicast and Anycast Group Membership [MAGMA] WG was chartered to address host-to-router signaling, including initial authentication and access control issues for multicast and anycast group membership, but other aspects of anycast, including architecture and routing, were outside the group's scope.

Simple Network Time Protocol (SNTP) Version 4 [RFC2030] defined how to use SNTP anycast for server discovery. This was extended in [RFC4330] as an NTP-specific "manycast" service, in which anycast was used for the discovery part.

IPv6 defined some reserved subnet anycast addresses [RFC2526] and assigned one to "Mobile IPv6 Home-Agents" [RFC3775] (obsoleted by [RFC6275]).

The original IPv6 transition mechanism [RFC2893] made use of IPv4 anycast addresses as tunnel endpoints for IPv6 encapsulated in IPv4, but this was later removed [RFC4213]. The 6to4 tunneling protocol [RFC3056] was augmented by a 6to4 relay anycast prefix [RFC3068] in a move aimed at simplifying the configuration of 6to4 routers. Incidentally, 6to4 deployment has shown a fair number of operational and security issues [RFC3964] that result from using anycast as a discovery mechanism. Specifically, one inference is that operational consideration is needed to ensure that anycast addresses get advertised and/or filtered in a way that produces the intended scope (e.g., only advertise a route for your 6to4 relay to Autonomous Systems (ASes) that conform to your own acceptable usage policy), an attribute that can easily become quite operationally expensive.

In 2002, DNS' use of anycast was first specified in "Distributing Authoritative Name Servers via Shared Unicast Addresses" [RFC3258]. It is notable that it used the term "shared unicast address" rather than "anycast address" for the service. This distinction was made due to the IPv6 differentiation in the on-link model. "Shared unicast" addresses are unicast (not multicast) in the IPv6 model and, therefore, support the off-link anycast model (described earlier) but not the on-link anycast model. At the same time, site-local-scoped well-known addresses began being used for recursive resolvers [DNS-DISC], but this use was never standardized (see below in Section 3.4 for more discussion).

Anycast was used for routing to rendezvous points (RPs) for PIM [RFC4610].

"Operation of Anycast Services" BCP 126 [RFC4786] deals with how the routing system interacts with anycast services and the operation of anycast services.

"Requirements for a Mechanism Identifying a Name Server Instance" [RFC4892] cites the use of anycast with DNS as a motivation to identify individual name server instances, and the Name Server ID (NSID) option was defined for this purpose [RFC5001]. One could view

the addition of NSID as an incarnation of locator and identifier separation (where the anycast address is a locator and the NSID is an identifier).

The IAB's "Reflections on Internet Transparency" [RFC4924] briefly mentions how violating transparency can also damage global services that use anycast.

2.2. Anycast in IPv6

Originally, the IPv6 addressing architecture [RFC1884] [RFC2373] [RFC3513] severely restricted the use of anycast addresses. In particular, the architecture provided that anycast addresses must not be used as source addresses and must not be assigned to IPv6 hosts (i.e., only routers). These restrictions were later lifted in 2006 [RFC4291].

In fact, the more recent "IPv6 Transition/Co-existence Security Considerations" [RFC4942] overview now recommends:

To avoid exposing knowledge about the internal structure of the network, it is recommended that anycast servers now take advantage of the ability to return responses with the anycast address as the source address if possible.

As discussed in the Overview, "on-link anycast" is employed expressly in IPv6 via ND NAs; see Section 7.2.7 of [RFC4861] for additional information.

2.3. DNS Anycast

"Distributed Authoritative Name Servers via Shared Unicast Addresses" [RFC3258] described how to reach authoritative name servers using multiple unicast addresses, each one configured on a different set of servers. It stated in Section 2.3:

This document presumes that the usual DNS failover methods are the only ones used to ensure reachability of the data for clients. It does not advise that the routes be withdrawn in the case of failure; it advises instead that the DNS process shutdown so that servers on other addresses are queried. This recommendation reflects a choice between performance and operational complexity. While it would be possible to have some process withdraw the route for a specific server instance when it is not available, there is considerable operational complexity involved in ensuring that this occurs reliably. Given the existing DNS failover methods, the marginal improvement in performance will not be sufficient to justify the additional complexity for most uses.

In anycast more generally, most anycast benefits cannot be realized without route withdrawals, since traffic will continue to be directed to the link with the failed server. When multiple unicast addresses are used with different sets of servers, a client can still fail over to using a different server address and, hence, a different set of servers. There can still be reliability problems, however, when each set contains a failed server. If all servers in the same set are on the same subnet, such problems could be minimized where address resolution within the subnet will cause traffic to go to an available server.

Other assertions included:

- o It asserted (as an advantage) that no routing changes were needed.
- o It recommended stopping DNS processes rather than withdrawing routes to deal with failures, data synchronization issues, and failover, as provided in the quoted text above. The spirit of this advice was that DNS resolvers may (indeed) reach out and query unavailable DNS name servers, but as their queries time out, they will elect to pin themselves to other server addresses and, hence, different servers.
- o It argued that failure modes involving state were not serious, because:
 - * the vast majority of DNS queries are UDP
 - * large routing metric disparity among authoritative server instances would localize queries to a single instance for most clients
 - * when the resolver tries TCP and it breaks, the resolver will try to move to a different server address. In order to ensure that this is possible, it is important that the DNS zone be configured with multiple server addresses for different sets of name servers. The advice given in Section 3.3 of [DNS-DISC] describes, in more detail, why using multiple addresses is important.

"Unique Per-Node Origin ASNs for Globally Anycasted Services" [RFC6382] makes recommendations regarding the use of per-node unique origin Autonomous System Numbers (ASNs) for globally anycasted critical infrastructure services in order to provide routing system discriminators for a given anycasted prefix. The object was to allow network management and monitoring techniques, or other operational

mechanisms to employ this new origin AS as a discriminator in whatever manner fits their operating environment, either for detection or policy associated with a given anycasted node.

2.4. BCP 126 on Operation of Anycast Services

"Operation of Anycast Services" BCP 126 [RFC4786] was a product of the IETF's GROW working group. The primary design constraint considered was that routing "be stable" for significantly longer than a "transaction time", where "transaction time" is loosely defined as "a single interaction between a single client and a single server". It takes no position on what applications are suitable candidates for anycast usage.

Furthermore, it views anycast service disruptions as an operational problem: "Operators should be aware that, especially for long running flows, there are potential failure modes using anycast that are more complex than a simple 'destination unreachable' failure using unicast".

The document primarily deals with global Internet-wide services provided by anycast. Where internal topology issues are discussed, they're mostly regarding routing implications rather than application design implications. BCP 126 also views networks employing per-packet load balancing on equal cost paths as "pathological". This was also discussed in [RFC2991].

3. Principles

3.1. Layering and Resiliency

Preserving the integrity of a modular layered design for IP protocols on the Internet is critical to its continued success and flexibility. One such consideration is that of whether an application should have to adapt to changes in the routing system.

Applications should make minimal assumptions about routing stability, just as they should make minimal assumptions about congestion and packet loss. When designing applications, it would perhaps be safe to assume that the routing system may deliver each anycast packet to a different service instance, in any pattern, with temporal reordering being a not-so-rare phenomenon.

Most stateful transport protocols (e.g., TCP), without modification, do not understand the properties of anycast; hence, they will fail probabilistically, but possibly catastrophically, when using anycast addresses in the presence of "normal" routing dynamics. Specifically, if datagrams associated with a given active transaction

are routed to a new anycasted end system and that end system lacks state data associated with the active transaction, the session will be reset; hence, it will need to be reinitiated. As another example, different networks have different routing properties and therefore will experience problems under different conditions. This can lead to a protocol working fine in, say, a test lab but not in the global Internet.

3.2. Anycast Addresses as Destinations

When an anycast address is used as a destination address, different packets with the same destination IP address may reach different destination hosts, even if the packets are generated by the same source host. Anycast addresses are thus "safe" to use as destination addresses for an application if the following design points are all met:

- o A request message or "one shot" message is self-contained in a single transport packet.
- o A stateless transport (e.g., UDP) is used for the above.
- o Replies are always sent to a unicast address; these can be multipacket since the unicast destination is presumed to be associated with a single "stable" end system and not an anycasted source address. Note that this constrains the use of anycast as source addresses in request messages, since reply messages sent back to that address may reach a device that was not the source that initially triggered it.
- o The server side of the application keeps no hard state across requests.
- o Retries are idempotent; in addition to not assuming server state, they do not encode any assumptions about loss of requests versus loss of replies.

It is noteworthy, though, that even under the above circumstances ICMP messages against packets with anycast source addresses may be routed to servers other than those expected. In addition, Path Maximum Transmission Unit Discovery (PMTUD) can encounter complications when employed against anycast addresses, since iterations in the PMTU discovery process may have packets routed to different anycast service instances.

3.3. Anycast Addresses as Sources

When an anycast address is used as a source address, the source address does not uniquely identify the source host; hence, replies might be sent to a different host. As noted earlier, this concept is sometimes referred to (e.g., in [RFC3258]) as a "shared unicast address". Anycast addresses are "safe" to use as source addresses for an application if all of the following design points are met:

- o No response message is generated by the receiver with the anycast source used as a destination unless the application has some private state synchronization that allows for the response message arriving at a different instance.
- o The source anycast address is reachable via the interface address if unicast reverse path forwarding (RPF) [RFC4778] checking is on, or the service address is explicitly provisioned to bypass RPF checks. In addition to the application defined in [RFC4778], Section 4.4.5 of BCP 126 [RFC4786] gives explicit consideration to RPF checks in anycasting operations.

3.4. Service Discovery

Applications able to tolerate an extra round-trip time (RTT) to learn a unicast destination address for multipacket exchanges might safely use anycast destination addresses for service instance discovery. For example, "instance discovery" messages are sent to an anycast destination address, and a reply is subsequently sent from the unique unicast source address of the interface that received the discovery message, or a reply is sent from the anycast source address of the interface that received the message, containing the unicast address to be used to invoke the service. Only the latter of these will avoid potential NAT binding and stateful firewall issues.

[DNS-DISC] discussed several options to address the need to configure DNS servers, including the use of a "Well-known Anycast Address" for recursive DNS service configuration in clients to ease configuration and allow those systems to ship with these well-known addresses configured "from the beginning, as, say, factory default". The proposal was later dropped, but the analysis was used in publishing [RFC4339].

After the final round of revisions to [DNS-DISC] was made, [RFC4339] was published with a very similar focus and overlapping content. The difference was that the writing in [RFC4339] focused on analysis, while [DNS-DISC] covered both the analysis and a specific proposal. The proposal details were removed in what became [RFC4339] although Section 3.3 of that RFC still discusses the approach of using a

well-known anycast address in this scenario. During publication, the IESG requested that the following "IESG Note" be contained in the document:

This document describes three different approaches for the configuration of DNS name resolution server information in IPv6 hosts.

There is not an IETF consensus on which approach is preferred. The analysis in this document was developed by the proponents for each approach and does not represent an IETF consensus.

The 'RA option' and 'Well-known anycast' approaches described in this document are not standardized. Consequently the analysis for these approaches might not be completely applicable to any specific proposal that might be proposed in the future.

4. Analysis

4.1. Regarding Widespread Anycast Use

Widespread use of anycast for global Internet-wide services or inter-domain services has some scaling challenges. Similar in ways to multicast, each service generates at least one unique route in the global BGP routing system. As a result, additional anycast instances result in additional paths for a given prefix, which scales super-linearly as a function of denseness of inter-domain interconnection within the routing system (i.e., more paths result in more resources, more network interconnections result in more paths).

This is why the Anycast BOF concluded that "the use of global anycast addresses was not expected to scale and hence was expected to be limited to a small number of key uses".

However, one interesting note is that multiple anycast services can share a route if they are all located in a single announced prefix and if all the servers of all the services are always collocated. If the announced prefix is aggregated differently in different locations though, longest-match routing might result in some anycast locations being unreachable. Hence, extra precaution must be taken when aggregating prefixes used by anycast services.

4.2. Transport Implications

UDP is the "lingua franca" for anycast today. Stateful transports could be enhanced to be more anycast friendly. This was anticipated in Host Anycasting Services [RFC1546], specifically:

The solution to this problem is to only permit anycast addresses as the remote address of a TCP SYN segment (without the ACK bit set). A TCP can then initiate a connection to an anycast address. When the SYN-ACK is sent back by the host that received the anycast segment, the initiating TCP should replace the anycast address of its peer, with the address of the host returning the SYN-ACK. (The initiating TCP can recognize the connection for which the SYN-ACK is destined by treating the anycast address as a wildcard address, which matches any incoming SYN-ACK segment with the correct destination port and address and source port, provided the SYN-ACK's full address, including source address, does not match another connection and the sequence numbers in the SYN-ACK are correct.) This approach ensures that a TCP, after receiving the SYN-ACK is always communicating with only one host.

The reason for such considerations can be illustrated through an example: one operationally observed shortcoming of using the Transmission Control Protocol (TCP) [RFC0793] and anycast nodes in DNS is that even during the TCP connection establishment, IP control packets from a DNS client may initially be routed to one anycast instance, but subsequent IP packets may be delivered to a different anycast instance if (for example) a route has changed. In such a case, the TCP connection will likely elicit a connection reset but will certainly result in the disruption of the connection.

Multi-address transports (e.g., SCTP) might be more amenable to such extensions than TCP.

The features needed for address discovery when doing multihoming in the transport layer are similar to those needed to support anycast.

4.3. Stateful Firewalls, Middleboxes, and Anycast

Middleboxes (e.g., NATs) and stateful firewalls cause problems when used in conjunction with some ways to use anycast. In particular, a server-side transition from an anycast source IP address to a unique unicast address may require new or additional session state, and this may not exist in the middlebox, as discussed previously in Section 3.4.

4.4. Security Considerations

Anycast is often deployed to mitigate or at least localize the effects of distributed denial-of-service (DDoS) attacks. For example, with the Netgear NTP fiasco [RFC4085] anycast was used in a distributed sinkhole model [RFC3882] to mitigate the effects of embedded globally routed Internet addresses in network elements.

"Internet Denial-of-Service Considerations" [RFC4732] notes that: "A number of the root nameservers have since been replicated using anycast to further improve their resistance to DoS".

"Operation of Anycast Services" BCP 126 [RFC4786] cites DoS mitigation, constraining DoS to localized regions, and identifying attack sources using spoofed addresses as some motivations to deploy services using anycast. Multiple anycast service instances such as those used by the root name servers also add resiliency when network partitioning occurs (e.g., as the result of transoceanic fiber cuts or natural disasters).

When using anycast, care must be taken not to simply withdraw an anycast route in the presence of a sustained DoS attack, since the result would simply move the attack to another service instance, potentially causing a cascaded failure. Anycast adds resiliency when such an attack is instead constrained to a single service instance.

It should be noted that there is a significant man-in-the-middle (MITM) exposure in either variant of anycast discovery (see Section 3.4) that, in many applications, may necessitate the need for end-to-end security models (e.g., using IPsec [RFC6071] or even DNSSEC [RFC4033]) that enable end systems to authenticate one another, or the data itself.

However, when considering the above suggestion of enabling end systems to authenticate each other, a potential complication can arise. If the service nodes of an anycast deployment are administered by separate authorities, any server-side authentication credentials that are used must (necessarily) be shared across the administrative boundaries in the anycast deployment. This would likely also be the case with Secure Neighbor Discovery, described in [RFC5909].

Furthermore, as discussed earlier in this document, operational consideration needs to be given to ensure that anycast addresses get advertised and/or filtered in a way that produces intended scope (for example, only advertise a route to your 6to4 relay to ASes that conform to your own Acceptable Use Policy (AUP)). This seems to be operationally expensive, and is often vulnerable to errors outside of the local routing domain, in particular when anycasted services are deployed with the intent to scope associated announcements within some local or regional boundary.

As previously discussed, [RFC6382] makes recommendations regarding the use of per-node unique origin ASNs for globally anycasted critical infrastructure services in order to provide routing system discriminators for a given anycasted prefix. Network management and

monitoring techniques, or other operational mechanisms, may then employ this new discriminator in whatever manner fits their operating environment, for either detection or policy associated with a given anycasted node.

Moreover, the use of per-node unique origin ASNs has the additional benefit of overcoming complications that might arise with the potential deployment of the Resource Public Key Infrastructure (RPKI) [RFC6480]. Without per-node unique origin ASNs, the cryptographic certificates needed to attest to the Route Origin Authorizations (ROAs) of a multi-administrative deployment of anycast would need to be shared. However, if each service instance has a separate ASN, then those ASNs can be managed separately in the RPKI.

Unlike multicast (but like unicast), anycast allows traffic stealing. That is, with multicast, joining a multicast group doesn't prevent anyone else who was receiving the traffic from continuing to receive the traffic. With anycast, adding an anycasted node to the routing system can prevent a previous recipient from continuing to receive traffic because it may now be delivered to the new node instead. As such, if an unauthorized anycast node can inject a route into the network, or be resolved using ARP/Neighbor Discovery on a link with an authorized anycast node, traffic can be diverted thereby triggering DoS or other attacks. Section 6.3 of BCP 126 [RFC4786] provides expanded discussion on "Service Hijacking" and "traffic stealing", and [FanInfocom13] discusses measured instances of anycast nodes and "benign masquerading or hostile hijacking of anycast services", by unauthorized nodes.

Unlike unicast (but like multicast), the desire is to allow applications to cause route injection. In multicast, one often allows arbitrary applications on hosts to join multicast groups, resulting in multicast routing state. Trying to apply that same model to anycast would present new security concerns, which is why [MAGMA] only got so far. The security concerns include:

1. Allowing route injection can cause DOS to a legitimate address owner.
2. Allowing route injection consumes routing resources and can hence cause DOS to the routing system and impact legitimate communications as a result.

These are two of the core issues that were part of the discussion during [RFC1884], the [ANYCASTBOF], and the MAGMA [MAGMA] chartering.

Additional security considerations are scattered throughout the list of references provided herein.

4.5. Deployment Considerations

BCP 126 [RFC4786] provides some very solid guidance related to operations of anycasted services and, in particular, the operations of DNS.

This document covers issues associated with the architectural implications of anycast. This document does not address, in any depth, the fact that there are deployed services with TCP transport using anycast today. Evidence exists to suggest that such practice is not "safe" in the traditional and architectural sense (as described in Section 4.2). These sorts of issues are indeed relative, and we recognize sometimes unpredictability in the routing system beyond the local administrative domain can be manageable. That is, despite the inherent architectural problems in the use of anycast with stateful transport and connection-oriented protocols, there is expanding deployment (e.g., for content distribution networks) and situations exist where it may make sense (e.g., such as with service discovery, short-lived transactions, or in cases where dynamically directing traffic to topologically optimal service instances is required). In general, operators should consider the content and references provided herein and evaluate the benefits and implications of anycast in their specific environments and applications.

In addition, (as noted in Section 2.3) the issue of whether to withdraw anycast routes when there is a service failure is only briefly broached in [RFC3258]. The advice given is that routes should not be withdrawn, in order to reduce operational complexity. However, the issue of route advertisements and service outages deserves greater attention.

There is an inherent trade-off that exists between the operational complexity of matching service outages with anycast route withdrawals, and allowing anycast routes to persist for services that are no longer available. [RFC3258] maintains that DNS' inherent failure recovery mechanism is sufficient to overcome failed nodes, but even this advice enshrines the notion that these decisions are both application-specific and subject to the operational needs of each deployment. For example, the routing system plays a larger role in DNS when services are anycast. Therefore, operational consideration must be given to the fact that relying on anycast for DNS deployment optimizations means that there are operational trade-offs related to keeping route advertisements (and withdrawals) symmetric with service availability. For example, in order to ensure that the DNS resolvers in a failed anycast instance's catchment [RFC4786] are able to fail over and reach a non-failed catchment, a route withdrawal is almost certainly required. On the other hand,

instability of a DNS process that triggers frequent route advertisement and withdrawal might result in suppression of legitimate paths to available nodes, e.g., as a result of route flap damping [RFC2439].

Rather than prescribing advice that attempts to benefit all situations, it should simply be recognized that when using anycast with network services that provide redundancy or resilience capabilities at other layers of the protocol stack, operators should carefully consider the optimal layer(s) at which to provide said functions.

As noted in Section 2.3, use of anycast within a subnet does not necessarily suffer from the potential issues with route withdrawals. As such, use of anycast to reach servers that reside in the same subnet can be made more reliable than use of anycast to reach topologically disparate server instances. Within a subnet, however, care must be taken as stated in Section 5.4 of [RFC4862], "Duplicate Address Detection MUST NOT be performed on anycast addresses"; hence, the servers must be configured appropriately.

5. Conclusions

In summary, operators and application vendors alike should consider the benefits and implications of anycast in their specific environments and applications and also give forward consideration to how new network protocols and application functions may take advantage of anycast or how they may be negatively impacted if anycasting is employed.

6. Acknowledgements

Many thanks to Kurtis Lindqvist for his early review and feedback on this document. Thanks to Brian Carpenter, Alfred Hoenes, and Joe Abley for their usual careful review and feedback, as well as Mark Smith, Lixia Zhang, Stephane Bortzmeyer, Masataka Ohta, and S. Moonesamy for their detailed reviews. Helpful feedback was also received from others including Edward Lewis, Jean-Michel Combes, Wolfgang Nagele, Mark Townsley, and Abdussalam Baryun.

7. Informative References

[ANYCASTBOF]

Deering, S., "IAB Anycast BOF Announcement", October 1999, <<http://www.ietf.org/mail-archive/web/ietf/current/msg11182.html>>.

- [DNS-DISC] Durand, A., Hagino, J., and D. Thaler, "Well known site local unicast addresses for DNS resolver", Work in Progress, September 2002.
- [FanInfocom13]
Fan, X., Heidemann, J., and R. Govindan, "Evaluating Anycast in the Domain Name System", Proceedings of the IEEE Infocom 2013, April 2013.
- [IMR9401] RFC Editor, "INTERNET MONTHLY REPORT", January 1994, <<ftp://ftp.rfc-editor.org/in-notes/museum/imr/imr9401.txt>>.
- [MAGMA] MAGMA (concluded), "Multicast and Anycast Group Membership (MAGMA)", April 2006, <<http://www.ietf.org/wg/concluded/magma>>.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC1546] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.
- [RFC1884] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 1884, December 1995.
- [RFC2030] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 2030, October 1996.
- [RFC2101] Carpenter, B., Crowcroft, J., and Y. Rekhter, "IPv4 Address Behaviour Today", RFC 2101, February 1997.
- [RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, July 1998.
- [RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", RFC 2439, November 1998.
- [RFC2526] Johnson, D. and S. Deering, "Reserved IPv6 Subnet Anycast Addresses", RFC 2526, March 1999.
- [RFC2893] Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 2893, August 2000.

- [RFC2902] Deering, S., Hares, S., Perkins, C., and R. Perlman, "Overview of the 1998 IAB Routing Workshop", RFC 2902, August 2000.
- [RFC2991] Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, November 2000.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.
- [RFC3258] Hardie, T., "Distributing Authoritative Name Servers via Shared Unicast Addresses", RFC 3258, April 2002.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3882] Turk, D., "Configuring BGP to Block Denial-of-Service Attacks", RFC 3882, September 2004.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, December 2004.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [RFC4085] Plonka, D., "Embedding Globally-Routable Internet Addresses Considered Harmful", BCP 105, RFC 4085, June 2005.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4330] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 4330, January 2006.
- [RFC4339] Jeong, J., "IPv6 Host Configuration of DNS Server Information Approaches", RFC 4339, February 2006.

- [RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, August 2006.
- [RFC4732] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.
- [RFC4778] Kaeo, M., "Operational Security Current Practices in Internet Service Provider Environments", RFC 4778, January 2007.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, December 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4892] Woolf, S. and D. Conrad, "Requirements for a Mechanism Identifying a Name Server Instance", RFC 4892, June 2007.
- [RFC4924] Aboba, B. and E. Davies, "Reflections on Internet Transparency", RFC 4924, July 2007.
- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", RFC 4942, September 2007.
- [RFC5001] Austein, R., "DNS Name Server Identifier (NSID) Option", RFC 5001, August 2007.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.
- [RFC5909] Combes, J-M., Krishnan, S., and G. Daley, "Securing Neighbor Discovery Proxy: Problem Statement", RFC 5909, July 2010.
- [RFC6071] Frankel, S. and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", RFC 6071, February 2011.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

- [RFC6382] McPherson, D., Donnelly, R., and F. Scalzo, "Unique Origin Autonomous System Numbers (ASNs) per Node for Globally Anycasted Services", BCP 169, RFC 6382, October 2011.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.
- [RSSAC29] "RSSAC 29 Meeting Minutes", December 2007, <<http://www.icann.org/en/groups/rssac/meetings/rssac-29-en.pdf>>.

Appendix A. IAB Members at the Time of Approval

Bernard Aboba
Jari Arkko
Marc Blanchet
Ross Callon
Alissa Cooper
Joel Halpern
Russ Housley
Eliot Lear
Xing Li
Erik Nordmark
Andrew Sullivan
Dave Thaler
Hannes Tschofenig

Authors' Addresses

Danny McPherson
Verisign, Inc.
12061 Bluemont Way
Reston, VA
USA

EEmail: dmcpherson@verisign.com

Dave Oran
Cisco Systems
USA

EEmail: oran@cisco.com

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA
USA

EEmail: dthaler@microsoft.com

Eric Osterweil
Verisign, Inc.
12061 Bluemont Way
Reston, VA
USA

EEmail: eosterweil@verisign.com