

IPv6 Survey: October 2011

Geoffrey Sisson

© 2011 The Measurement Factory

November 11, 2011

Abstract

In this study, commissioned by Infoblox, we undertake to measure some basic statistics on IPv6 deployment. We approach the subject from a DNS-centric perspective, taking a 1% sample of domain names in the .com, .net, and .org zones and investigating to what extent IPv6 has been deployed for key services associated with each domain.

1 Introduction

With the impending exhaustion of IPv4 addresses, the deployment of IPv6 has begun to gain momentum. Earlier this year the Internet Society hosted World IPv6 Day, during which a number of websites and Internet service providers conducted a test deployment of IPv6 services.[1] More than in any previous year, IPv6 is in the spotlight, receiving more attention than other significant Internet infrastructure initiatives such as DNSSEC and Internationalized Domain Names.

In this study we examine various aspects of IPv6 deployment on the Internet.

- We survey the number of name servers with the AAAA resource records that are required for them to communicate via IPv6.
- We also survey the availability of web and mail services via IPv6.
- We test IPv6 connectivity to these services from servers in three different autonomous systems.
- Finally, we take a generalized look at IPv6 adoption by country based on our results.

This is not intended to be a fully comprehensive survey of IPv6 use. For example, we do not measure levels of traffic on public networks or arriving at a particular destination. Nonetheless, we hope it offers useful insights into the current state of IPv6 deployment.

2 Data set used

As the basis for our survey, we used a 1% sampling of the domains in the .com, .net, and .org zones. We obtained zone files from VeriSign and Public Interest Registry (PIR) for .com, .net, and .org containing approximately 97.6 million, 14.2 million, and 9.5 million domains respectively, and randomly selected 1% (1,212,227) to probe. Names were chosen from each zone using the Mersenne Twister MT19937 algorithm as implemented by the GNU Scientific Library.

3 Zones with name servers with AAAA RRs

3.1 Methodology

We recorded the number of zones with one or more name servers with a globally routable IPv6 address. The name servers were obtained from authoritative answers for the child zone only. If none of the authoritative name servers in the child zone had an associated AAAA RR, we determined that the zone had no IPv6-based name service, even if there was an NS RR with an associated AAAA RR in the parent zone.

As using a CNAME in the NSDNAME field of an NS RR is prohibited by RFC 2181, we did not follow any CNAMEs encountered in NS RRs. (0.7% of the zones in the sample had CNAMEs for name servers.) Another common misconfiguration that is tolerated by popular resolvers is the use of IP addresses in the NSDNAME field of NS RRs. We did not check for IPv6 addresses in the NSDNAME.

3.2 Results

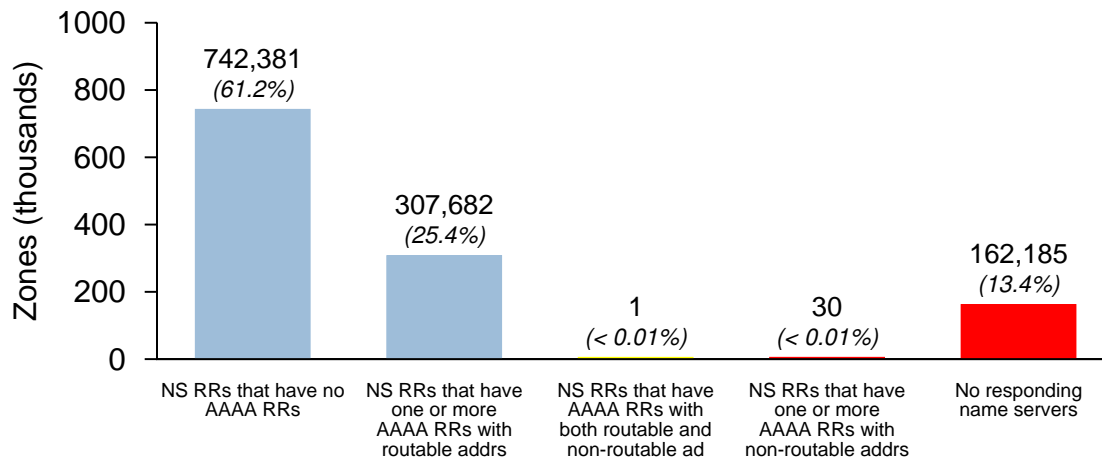


Figure 1: Number of zones with name servers that have AAAA RRs.

Figure 1 shows the number of zones with name servers that have AAAA RRs. Note that a very small number of zones have name servers with non-routable addresses. These were mostly loopback, link-local, IPv4-mapped IPv6 addresses, and TEREDO addresses.

We have tracked the number of zones with name servers with AAAA RRs in our annual DNS Survey since 2006.[2] Figure 2 shows the history of this value.

The increase since last year's survey is striking. In October 2010, only 14,347 (1.27%) of the zones in the sample had name servers with IPv6 addresses. This year it shot up to 307,682 (25.4%), a twenty-fold increase. At first glance this seems implausibly high – IPv6 adoption has been steady but hasn't shown signs of exuberance. However, there is a simple explanation: nearly all of the increase is attributable to one registrar, GoDaddy.com, that added IPv6 addresses to the name servers they use for hosting.




Figure 2: Percentage of zones with name servers with one or more AAAA RRs (historical).

GoDaddy.com – or, more accurately, their subsidiary, Special Domain Services, Inc. – appears to have 92 name servers with domain names ending with “domaincontrol.com” in the 2607:F208::/32 network that serve approximately 22% of all domains in the .com, .net, and .org TLDs. As we will see later, GoDaddy.com’s deployment – with a handful of exceptions – is limited to DNS only. Mail servers and web servers for the hosted zones are almost entirely IPv4 only.

It’s interesting to look at what growth might have been like had GoDaddy.com not deployed IPv6 on their name servers. Figure 3 shows the historical growth when GoDaddy.com’s name servers are omitted from the sample. The number of zones with name servers with AAAA records is now 37,815, 3.12% of the zones in the sample. This is a more than two-fold increase from last year. It’s evident that not only is IPv6 deployment increasing, but that rate of deployment is increasing, too.

The GoDaddy.com deployment is significant, but it has the effect of overshadowing and potentially obscuring interesting results. In this survey we occasionally omit the 269,867 zones with GoDaddy.com name servers (which we will refer to as the “Go Daddy zones”) from consideration in an attempt to discern more general results. Note that there were 221 zones that had both GoDaddy.com and non-GoDaddy.com name servers. These are *not* included in the set of Go Daddy zones.

Figure 4 shows the number of zones with one, two, etc., name servers with associated AAAA RRs. No zone in the sample had more than nine name servers with AAAA RRs. The overwhelming majority had two name servers with AAAA RRs. This does not take into account the Go Daddy zones, which had almost exclusively two name servers with AAAA RRs.

Figure 5 shows the percentage of zones per TLD that had at least one name server with an AAAA RR. Again, this does not include the Go Daddy zones. VeriSign began accepting AAAA glue records for publication in the .com and .net zones in 2002, while PIR began accepting them for




Figure 3: Historical values with GoDaddy.com-hosted zones omitted.




Figure 4: Numbers of name servers per zone that had associated AAAA RRs.

the .org zone in 2008. In spite of this, .org now has the highest proportion of zones with IPv6 name servers.

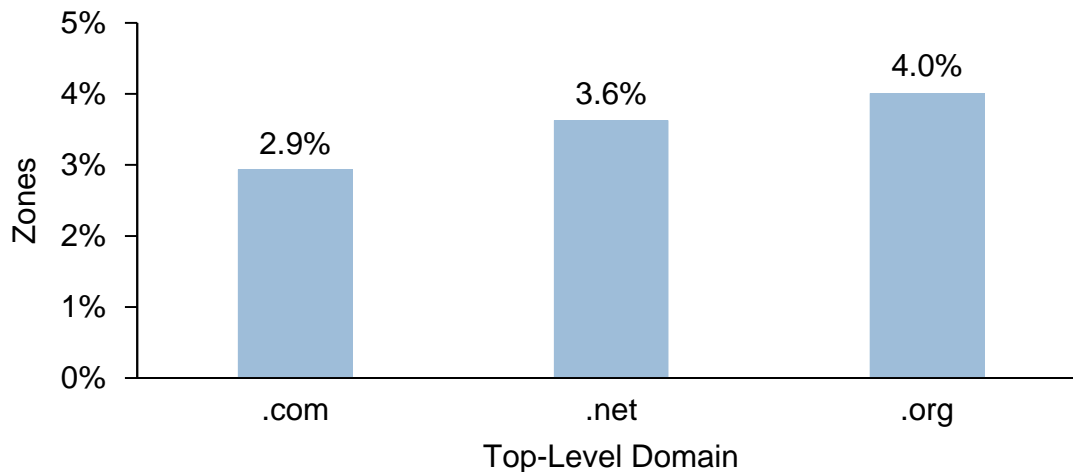


Figure 5: Percentage of zones per TLD that had at least one name server with an AAAA RR.

4 Zones with mail servers with AAAA RRs

Besides DNS, the two services that most zones have is SMTP (mail) and HTTP (web). We surveyed the extent to which the zones in our sample had enabled IPv6 for these services.

4.1 Methodology

We sent queries to an authoritative name server for each zone in the sample to see whether MX RRs existed for the zone. If any of the MX RRs with the lowest preference value (in other words, the terminal mail exchanger) had AAAA RRs with a routable IPv6 address for the zone, we considered the zone to have IPv6 mail (SMTP) service.

As CNAMEs in the EXCHANGE field of an MX RR are prohibited by RFC 2181, we did not follow any CNAMEs encountered. However, only 0.6% of the zones in the sample had CNAMEs for their lowest-preference-value exchangers.

Note that just because a zone lacks MX RRs does not mean it's incapable of receiving mail. In the absence of MX RRs, an A RR for the origin of the zone is treated as a mail server by mail transfer agents. However, we did not count hosts without MX RRs as mail servers.

4.2 Results

Figure 6 shows the number of zones with MX RRs that have hosts with AAAA RRs. As with the name servers, a very small number of zones (177) had mail servers with non-routable IPv6 addresses.

5 ZONES WITH WEB SERVERS WITH AAAA RRS




Figure 6: Number of zones with MX RRs that have associated AAAA RRs.

5 Zones with web servers with AAAA RRs

5.1 Methodology

If either `<domain>` or `www.<domain>` had an AAAA RR, we surmised that the zone *probably* had IPv6 web service. (In subsequent testing, 99.6% of the zones with A RRs and 85.7% with AAAA RRs were found to have web servers.) When resolving AAAA RRs, we followed CNAMEs and DNAMEs. The longest CNAME chains encountered were five CNAMEs deep.

5.2 Results

Figure 7 shows the number of zones that have AAAA RRs. 910 zones had AAAA records with non-routable addresses. While this number is small (0.07%) compared to the sample size, it represents a surprising 7.6% of all zones with AAAA RRs. It turns out that most (704) of these are attributable to the use of IPv4-mapped IPv6 addresses (addresses like “`::ffff:192.0.2.126`”) by an ICANN registrar who appears to summarily add them for all of the domains they host.

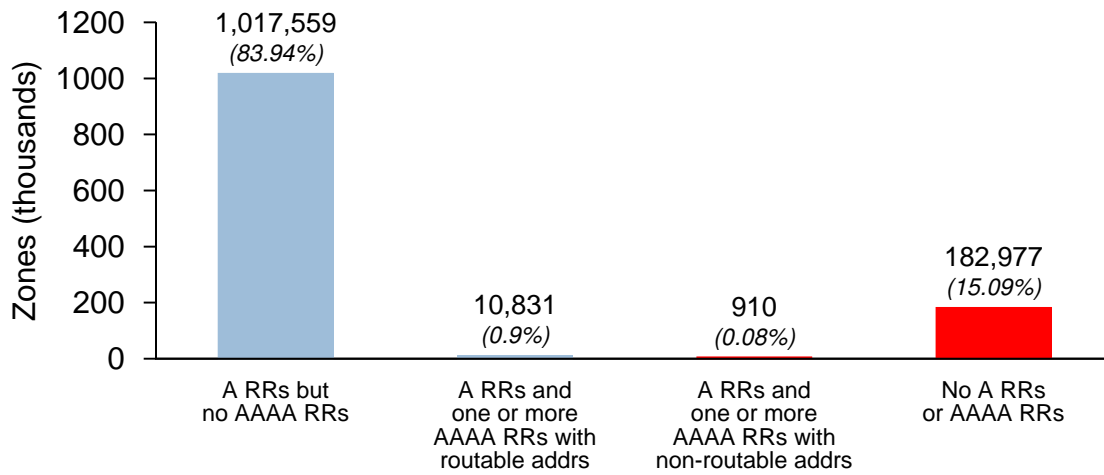


Figure 7: Number of zones with AAAA RRs for either $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$.

6 Relative allocation of services

We were interested in understanding how IPv6-enabled services were distributed relative to one another. For example, did zones with IPv6 mail servers tend also to have IPv6 name servers? Or was there no correlation?

We first looked at how these services were allocated in the sample for their IPv4 counterparts as a basis for comparison.

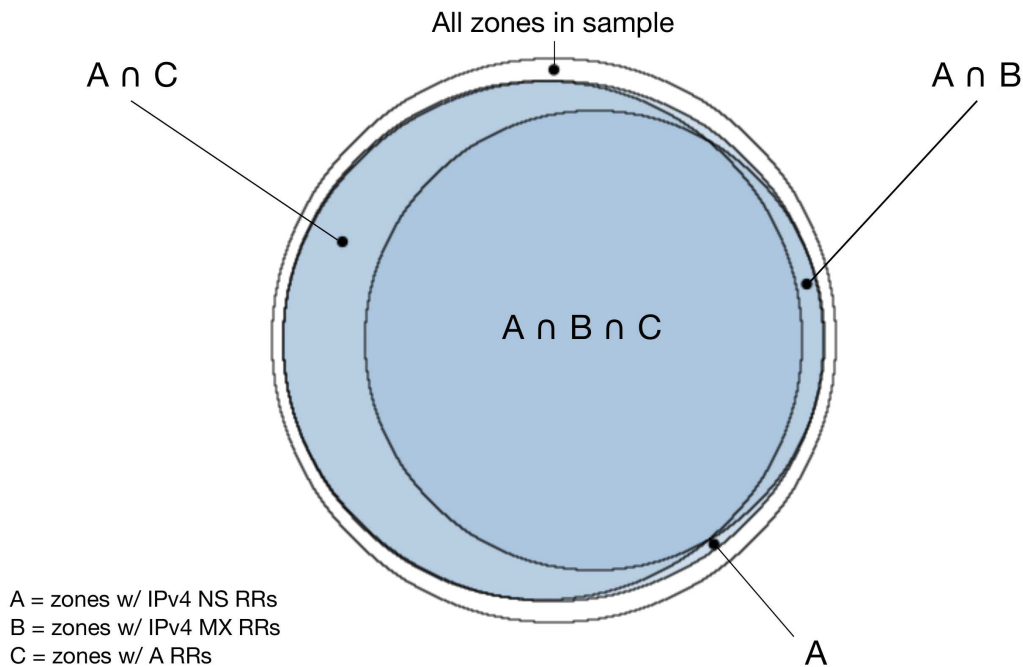


Figure 8: Relative allocation of IPv4 services by zone.

Figure 8 shows the allocation of IPv4 services. The diagram is an area-proportional Venn diagram, so the relative sizes are meaningful.¹ The outer circle represents the 1,212,277 zones in the sample. The inner circles represent the number of name servers, mail servers, and web servers with IPv4 addresses:

- A = zones with IPv4 NS RRs
- B = zones with IPv4 MX RRs
- C = zones with A RRs for $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$

It's difficult to see, but there are three inner circles: two intersecting ones contained (barely) by a third almost-circular ellipse. The ellipse represents the number of zones with IPv4 name servers. The larger circle represents the number of zones with A RRs for either $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$, and the smaller one the number of zones with MX RRs.

¹Note that the areas of intersections are not exact – they represent the best approximation that can be made with circles.

We can see that nearly all of the zones (86.6%) have at least one service. (The others simply do not resolve, as they have no name servers.) The set of zones with MX RRs and the set of zones with A RRs are completely contained in the set of zones with name servers, which is unsurprising – the only way a zone could have an A RR without an IPv4 name server would be if it were served by name servers with IPv6 only. The overlapping area in the middle represents the 66.3% of zones that have all three services: name servers, mail servers, and web servers.

	Zones	Percent
A	1,050,006	86.6%
B	811,920	67.0%
C	1,028,168	84.8%
$A \cap B$	811,920	67.0%
$A \cap C$	1,028,168	84.8%
$B \cap C$	804,232	66.3%
$A \cap B \cap C$	804,232	66.3%

Table 1: Relative allocation of IPv4 services by zone.

Table 1 contains the actual values for the allocation of IPv4 services.




Figure 9: Relative allocation of IPv6 services by zone.

Figure 9 shows the allocation of IPv6 services. In this diagram:

A = zones with IPv6 NS RRs

6 RELATIVE ALLOCATION OF SERVICES

B = zones with IPv6 MX RRs
 C = zones with AAAA RRs for $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$

It's immediately apparent that all zones with IPv6 services combined are a fraction of those for IPv4. The large circle represents zones with name servers with AAAA RRs; the bulk of these are the Go Daddy zones. The next-smaller circle represents the number of zones with MX RRs with associated AAAA RRs, and the smallest circle represents the number of zones with AAAA RRs. Perhaps surprisingly, the number of zones with all three services is tiny – just 0.1% of all zones. Note that the areas of intersections are not exact – they represent the best approximation that can be made with circles – so in this case the actual intersection $A \cap B \cap C$ is even smaller than depicted.

	Zones	Percent
A	307,682	25.4%
B	24,895	2.1%
C	10,831	0.9%
$A \cap B$	5,979	0.5%
$A \cap C$	1,888	0.2%
$B \cap C$	6,749	0.6%
$A \cap B \cap C$	1,139	0.1%

Table 2: Relative allocation of IPv6 services by zone.

Table 2 contains the actual values for the allocation of IPv6 services.

Finally, Figure 10 shows the allocation of IPv6 services with the Go Daddy zones omitted. In this diagram:

A = zones with IPv6 NS RRs (*Go Daddy zones omitted*)
 B = zones with IPv6 MX RRs
 C = zones with AAAA RRs for $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$

	Zones	Percent
A	37,815	3.1%
B	24,895	2.1%
C	10,831	0.9%
$A \cap B$	5,971	0.5%
$A \cap C$	6,749	0.6%
$B \cap C$	1,870	0.2%
$A \cap B \cap C$	1,138	0.1%

Table 3: Relative allocation of IPv6 services by zone (with Go Daddy zones omitted).

Table 3 contains the actual values for the allocation of IPv6 services modulo the Go Daddy zones.

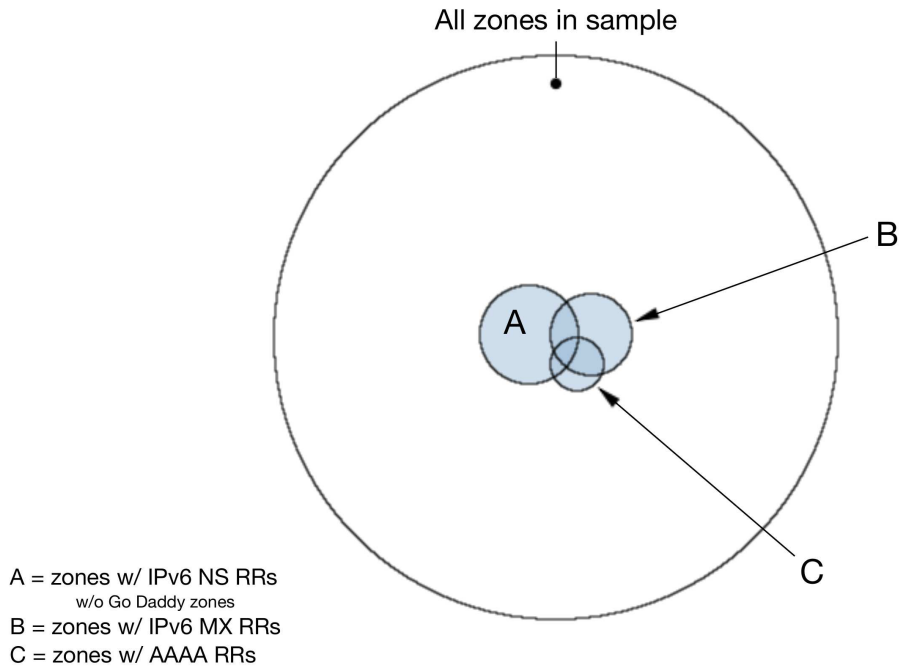


Figure 10: Relative allocation of IPv6 services by zone (Go Daddy zones omitted).

6.1 Comparison: Alexa 500

As a comparison, we also looked at the IPv6 records in use by the Alexa Top 500 Global Sites.[3]

Figure 11 shows the allocation of IPv6 services in the Alexa Top 500. In this diagram:

- A = Alexa Top 500 zones with IPv6 NS RRs
- B = Alexa Top 500 zones with IPv6 MX RRs
- C = Alexa Top 500 zones with AAAA RRs for $\langle \text{domain} \rangle$ or $\text{www}.\langle \text{domain} \rangle$

The distribution of services mirrors the ones we've seen for the zones in the sample: the bulk of the IPv6 deployment is NS RRs. There's one difference: in contrast to the zones in the sample, there are more AAAA RRs than there are IPv6 MX RRs. It's interesting to note that only one of the zones has all three IPv6 services.

Table 4 contains the actual values for the allocation of IPv6 in the Alexa Top 500.

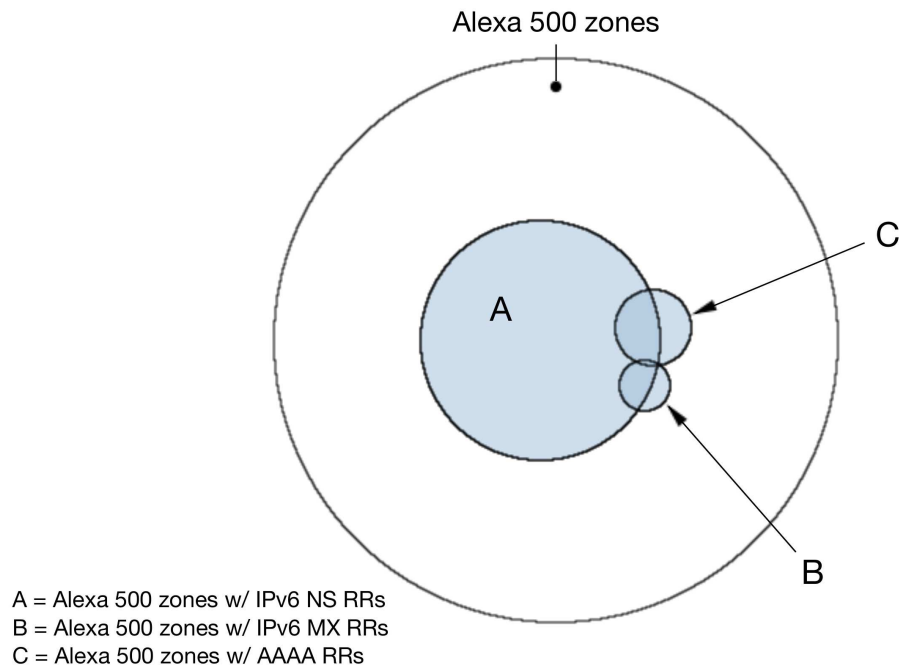


Figure 11: Relative allocation of IPv6 services, Alexa Top 500 Global Sites.

	Zones	Percent
A	90	18.0%
B	4	0.8%
C	9	1.8%
$A \cap B$	2	0.4%
$A \cap C$	3	0.6%
$B \cap C$	1	0.2%
$A \cap B \cap C$	1	0.2%

Table 4: Relative allocation of IPv6 services: Alexa Top 500 Global Sites.

7 IPv6 connectivity

Publishing an AAAA RR in the DNS is not synonymous with IPv6 deployment. The referenced host must also be reachable via IPv6. We sought to learn how many hosts with AAAA RRs gathered from our sample could be contacted via IPv6.

7.1 IPv6 DNS connectivity

7.1.1 Methodology

For each zone in the data set with at least one name server with at least one A RR containing a routable IPv4 address and at least one AAAA RR containing a routable IPv6 address, we sent UDP queries for the SOA RR of the zone to all IPv4 and IPv6 addresses associated with all name servers for the zone and recorded the results. We used a Perl script utilizing the Net::DNS CPAN module, which we ran on test boxes in three distinct autonomous systems. If we received a correct reply for the zone at any of the test boxes, we recorded the address as “good”. If we received no reply at all, we recorded the address as “unreachable”. Retries were set to 3, and the timeout was set to 15 seconds. To maximize the probability of a response, the test boxes sent queries at different times separated by at least four hours. The Go Daddy zones were not included.

7.1.2 Results




Figure 12: DNS query results by zone – IPv4 vs IPv6.

Figure 12 shows the range of responses by zone for the 37,814 zones in the sample with both IPv4 and IPv6 name servers. “Good” means at least one name server replied with an authoritative SOA RR response. “Unreachable” means no response was received from any name server to queries sent from all three test boxes. “Other” means that at least one response was received, but that none of the responses were valid, usually because the server wasn’t authoritative for the zone.

Note that “good” doesn’t mean all replies for the zone were good, just that at least one reply was good, which is sufficient to resolve the domain. However, “Unreachable” replies were not received

from any name server. We did not test whether unresponsive name servers were reachable via other protocols, e.g. ICMP.

For IPv4, we see that no responses were received for four zones. These appear to be outliers. In at least two cases, the domains appear to have been deregistered in the 24-hour period between when the sample was initially probed for the existence of IPv6 name servers, and the connectivity tests were performed.

We also see that 459 (1.2%) zones received only non-authoritative responses via IPv4. This appears to be attributable to the practice by some name service providers of maintaining a shadow TLD or root zone rather than multiple SLD zones, typically to avoid managing large numbers of zone files. In this case, no SOA RR is returned, or the provider returns a CNAME for the SOA RR. In either case, the zone will not be marked good by our methodology.

The IPv6 results are more mixed. We received no responses via IPv6 for 818 (2.2%) zones. We received no “good” responses for another 1,098 (2.9%), though some of these will be attributable to the practice described above. It will come as little surprise that there is a higher incidence of DNS misconfiguration associated with IPv6. Because of the low number of IPv6 users and also the redundancy provided by the availability of the IPv4 name servers, a zone with a broken IPv6 configuration is more likely to escape detection.

7.2 IPv6 SMTP connectivity

7.2.1 Methodology

For each zone in the data set whose lowest-preference-value MX RRs specified at least one mail exchanger with at least one A RR containing a routable IPv4 address and at least one mail exchanger with at least one AAAA RR containing a routable IPv6 address, we attempted to connect via SMTP to all IPv4 and IPv6 addresses associated with all lowest-preference-value mail exchangers for the zone and recorded the results. We used a Perl script using Perl’s built-in low-level socket routines as opposed to higher-level modules such as `Net::SMTP` or `IO::Socket` to avoid any application-layer interactions that might obscure the result. If we received a response code of “220” we recorded the zone as having “good” SMTP. We also counted response codes of “421” (connection rate limit exceeded) as good, as there were cases where multiple zones in the sample shared the same mail exchanger, and then repeated connections caused the mail exchanger to throttle.

If the connection attempt timed out or we received an error from the socket layer indicative of any ICMP unreachable message for all mail exchangers, all we recorded the zone as having “unreachable” mail exchangers. The connection timeout was set to 30 seconds. If we received a RST message or any other error, we categorized the connection attempt as “other”. To maximize the probability of a connection, the test made attempts at different times separated by at least four hours. We included the Go Daddy zones; however it should be noted that only 21 such zones had AAAA RRs. To maximize the probability of a connection, connection attempts on different test boxes were made at different times, separated by at least four hours. We included the GoDaddy zones; however it should be noted that only eight such zones had IPv6 mail exchangers.

Note that we did not attempt to send messages to the mail exchanger, so it’s possible that a mail exchanger marked as “good” might not actually accept mail for the zone.

Note also that MX RRs are not required for SMTP service; in the absence of MX RRs, any A or AAAA RR is treated as an implicit MX RR. However, we did not attempt SMTP connections if MX RRs were not present, as it’s difficult to know whether the operator intends for SMTP to be

available.²

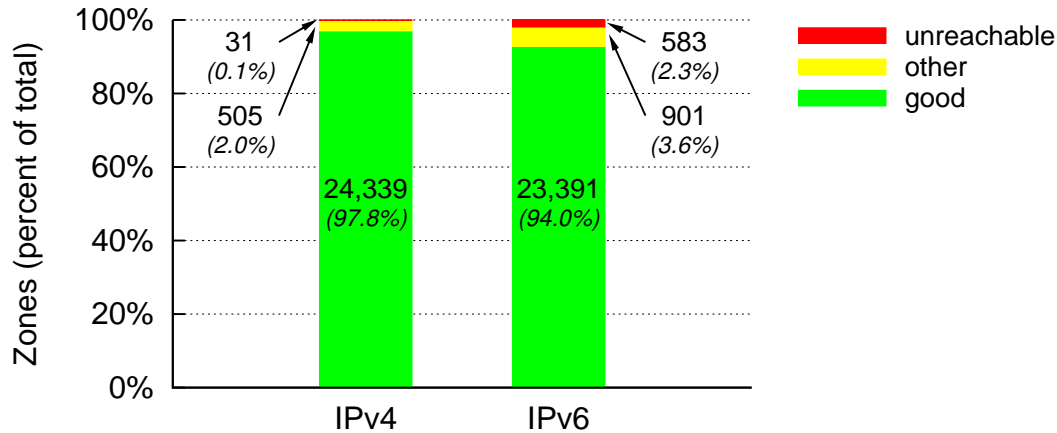


Figure 13: SMTP connectivity by zone – IPv4 vs IPv6.

7.2.2 Results

Figure 13 shows the range of responses by zone for the 24,875 zones in the sample with both IPv4 and IPv6 mail exchangers. The results are not dissimilar to the DNS results in the previous section. We were initially surprised by the high proportion of mail exchangers that were available even via IPv4, as we had anticipated a higher incidence of server (and service) misconfiguration. However, we suspect there’s a correlation between operators with IPv6 deployments and better-than-average network management practices.

As with the DNS results, we see what appears to be a higher incidence of misconfigurations of IPv6 SMTP than we see for its IPv4-based counterpart. Nearly all of the connection results categorized as “other” were TCP RST messages.

7.3 IPv6 HTTP connectivity

7.3.1 Methodology

For each zone in the data set with one or more A RRs for either `<domain>` or `www.<domain>` and one or more AAAA RRs for either `<domain>` or `www.<domain>`, we attempted to connect via HTTP to all IPv4 and IPv6 addresses specified by the A and AAAA RRs and recorded the results. As with the SMTP connectivity tests, we used a Perl script using Perl’s built-in low-level socket routines as opposed to higher-level modules such as LWP, Net::HTTP, or IO::Socket to avoid any application-layer interactions that might obscure the result. We sent an HTTP request message with the following format:

```
GET /robots.txt HTTP/1.0
```

²Technically we could check for SRV RRs for `_smtp._tcp`, but the use of SRV RRs for SMTP is nearly nonexistent.

```
User-Agent: test_http
Host: <hostname>
```

where `<hostname>` was the domain name of the target zone.

We used `robots.txt` in an attempt to minimize any concern that the connection attempts might cause to website operators; `GET /robots.txt` requests are typical of responsibly managed web crawlers. As we were not interested in the content of the response, it was inconsequential whether the `robots.txt` file actually existed or not (it existed only around 30% of the time).

If we received any valid HTTP response (e.g., a response beginning with “HTTP/1.0” or “HTTP/1.1”), we marked the zone as having “good” HTTP connectivity. If the connection attempt timed out or we received an error from the socket layer indicative of any ICMP unreachable message, we recorded the zone as having “unreachable” web service. (The timeout interval was 30 seconds.) If we received a TCP RST message or any other error, we categorized the connection attempt as “other”. To maximize the probability of a connection, the test boxes made connection attempts at different times separated by at least four hours.

For this test we included the Go Daddy zones; however there were only 21 Go Daddy zones that had AAAA RRs.

7.3.2 Results




Figure 14: HTTP connectivity by zone – IPv4 vs IPv6.

Figure 14 shows the range of responses by zone for the 10,806 zones in the sample with both A and AAAA RRs. The large number of “unreachables” is primarily due to the use of a single unreachable address (2001:4860:800b::79) by 1,285 zones, as configured by a single hosting provider. Without this anomaly, there would be only 118 (1.1%) of zones with AAAA RRs that were unreachable via IPv6.

8 IPv6 adoption by country

We were interested to gauge the level of IPv6 adoption by country. To this end, we attempted to determine the location by country of the zones in the data set. We used a heuristic that inferred the zone’s location based on the IP addresses for the NS, MX, A, and AAAA RRs for the zone. We used a recent copy of the MaxMind GeoIP Country geolocation database to determine the locations of the IP addresses.[4] The heuristic gave precedence to the location of the addresses associated with the lowest-preference-value MX RRs, then the A and AAAA RRs for the zone, and then the NS RRs for the zones. The heuristic is described in Appendix A. We note that this approach is imperfect, but alternative methods for achieving better geolocation were impractical. These limitations should be considered when evaluating the results in this section.




Figure 15: IPv6 adoption by country.

Figure 15 shows the distribution of the top 40 countries (by zone count in the sample) and the percentage zones from each country that had one or more associated AAAA RR with a routable IP address. France figures prominently at over 57%. This is principally due to France’s two largest registrars – Gandi and OVH – deploying IPv6 on their name servers. The United States is second, largely on the strength of the GoDaddy.com deployment. The Czech Republic is third, largely due the popularity of Active 24, a pan-European hosting company that has deployed IPv6 on its name servers. This pattern among the top IPv6 adopters by country is consistent: one or two registrars popular in the country has deployed IPv6 on their name servers.

Figure 16 shows the distribution of the top 40 countries with the Go Daddy zones omitted. Unsurprisingly, the US moves from second to sixteenth. A few other countries swap positions, as GoDaddy.com hosts some zones identified as non-US zone (though, based on our heuristic, it’s less than 1%).

It’s apparent that the results in Figure 15 and Figure 16 are dominated by the configuration of a small number of large registrars and web hosting companies who have deployed IPv6 on their

8 IPV6 ADOPTION BY COUNTRY




Figure 16: IPv6 adoption by country (Go Daddy zones omitted).

name servers. Another way of attempting to gain an appreciation of IPv6 deployment is to look at zones which have “full IPv6”, that is, AAAA RRs for all three services (DNS, mail and web). Organizations with these zones can be characterized as having fully embraced IPv6.

Figure 17 shows the distribution of the countries with five or more zones in the sample with all three IPv6 services as a proportion of all zones from that country in the sample. Slovakia tops the list, with 15 of its 119 zones having all IPv6 services. This is due to two Slovakian hosting companies (Yegon s.r.o. and ELBIA s.r.o.). Too much should not be read into this: between the imprecise nature of the geolocation heuristic and the small numbers involved, the prominence of Slovakia may be artifactual.

The result for the Netherlands is more credible. 664 zones had all IP services out of 13,896 active zones in the sample. 390 of these are attributable to one Dutch hosting company (TransIP B.V.) and 184 to another (Tiscomhosting B.V.), but the remaining 90 are spread across a number of different providers and organizations.

The US ranks 10th, in part because of the large number of zones that the heuristic identified as belonging to the US (696,622, which is around 66% of the zones with responding name servers in the sample). It is possible that this is an inflated number, as many foreign entities use US hosting services and registrars. Also, many of the most popular domain parking sites are in the US. However, even if the count is inflated by an order of magnitude, the US would rank only eighth – just behind Germany – if the count were to be adjusted.

Note that Asian countries known to have significant IPv6 deployments, such as China and Japan, may be underrepresented here because they have relatively fewer registrations in the `.com`, `.net`, and `.org` zones than the US and (to a lesser extent) European countries.




Figure 17: IPv6 adoption by country – Zones with “full IPv6”.

A Geolocation heuristic

The following heuristic was used to infer the location of a zone. The heuristic assumes that any CNAMEs for either `<domain>` or `www.<domain>` have already been resolved.

1. If MX RRs exist for the zone, use the IP addresses of any A or AAAA RRs for the lowest-preference mail exchangers. If there is more than one lowest-preference value mail exchanger, use the IP address(es) associated with any in-bailiwick mail exchangers. If there are multiple IP addresses with conflicting locations, ignore any that are for “EU” (rationale: other country codes likely represent a more specific location). If there are still conflicting locations, pick one at random and take that to be the location of the zone.
2. If there are no MX RRs but there are A RRs or AAAA RRs for either `<domain>` or `www.<domain>`, use the location of these addresses as the location of the zone. If there are no A RRs or AAAA RRs but there is one or more CNAMEs, resolve the CNAMEs and continue. If there are multiple IP addresses with conflicting locations, ignore any that are for “EU”. If there are still conflicting locations, pick one at random and take that to be the location of the zone.
3. Use the location of the A and AAAA RRs associated with the authoritative NS RRs for the zone. If there are both in-bailiwick and out-of-bailiwick name servers, prefer the location of in-bailiwick name servers. If there are multiple IP addresses with conflicting locations, ignore any that are for “EU”. If there are still conflicting locations, pick one at random and take that to be the location of the zone.

Glossary

Go Daddy zones

The set of 269,867 zones in the sample that had name servers exclusively operated by Special Domain Services, Inc., a subsidiary of GoDaddy.com, Inc., as evidenced by name servers with names in the `domaincontrol.com` domain. This set is excluded from some tests and results in the survey, as it potentially obscures more general patterns. 88% of the zones in the sample with IPv6 name servers were GoDaddy.com zones, comprising 22% of all zones in the sample. 221 zones in the sample had both GoDaddy.com and non-GoDaddy.com name servers; these are not included in the set of Go Daddy zones.

routable address

An IP address that is globally routed. For IPv4 addresses, this included addresses with prefixes marked as ALLOCATED or LEGACY in the IANA IPv4 Address Space Registry with the exception of any prefixes documented in RFC 5735 or in Team Cymru “fullbogons” list.[5][6][7] For IPv6 addresses, this included addresses with a prefix in the IANA IPv6 Global Unicast Address Assignment registry marked as ALLOCATED with the exception of any addresses with a prefix in the IANA IPv6 Special Purpose Address Registry or the Team Cymru “fullbogons” list[7][8][9].

References

- [1] World IPv6 Day. <http://www.worldipv6day.org/>
- [2] Measurement Factory DNS Surveys. <http://dns.measurement-factory.com/surveys/>
- [3] Alexa Top 500 Global Sites. <http://www.alexa.com/topsites>
- [4] MaxMind GeoIP Country Database. <http://www.maxmind.com/app/country>
- [5] IANA IPv4 Address Space Registry.
<http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml>
- [6] Cotton, M., and L. Vegoda, *Special Use IPv4 Addresses*, RFC 5735, BCP 153, January 2010.
<http://tools.ietf.org/html/rfc5735>
- [7] Team Cymru “fullbogons” list. <http://www.team-cymru.org/Services/Bogons/>
- [8] IANA IPv6 Global Unicast Address Assignments.
<http://www.iana.org/assignments/ipv6-unicast-address-assignments/>
- [9] IANA IPv6 Special Purpose Address Registry.
<http://www.iana.org/assignments/iana-ipv6-special-registry/>