Working Paper on

Comparison of Performance over IPv6 versus IPv4

By

Arthur Berger

Akamai Technologies

Abstract

As IPv4 address space gets tighter, there is increasing pressure to deploy IPv6. The Internet Assigned Number Authority (IANA) allocated the last of the available /8's of the v4 address space to the Regional Internet Registries (RIR's) on February 3, 2011. Currently, the RIR's are restricting allocations to cover only about 3 months of growth. A market for legacy v4 address has begun: In March, 2011, as part of Nortel's bankruptcy, Microsoft bought 667,000 legacy v4 addresses for $11/address.

Since the transition to IPv6 will be slow, there will be a long period where many end-points will be dual stack. Thus, the ability to pick the better performing path over v4 versus v6 will be a valuable feature. We have done a performance comparison of v4 versus v6 latency and loss, with results by continent, and by tunneled versus native v6 addresses. Although overall performance is better over v4, it is not always so; for example 10% of the time the latency between the U.S. and Europe is shorter over v6 by at least 10 ms, and to Asia is shorter by at least 38 ms. Latency and loss over v6 is in general higher to tunneled v6 destinations, as compared with native. Somewhat surprisingly, the latency and loss over _v4_ is also higher to nameservers whose v6 interface is tunneled, as compared with nameservers whose v6 interface is native. We conjecture that nameservers with a tunneled v6 interface are more likely to be in smaller networks, lower down in the hierarchy. Thus, the common observation that v6 latency is higher over tunnels is not due exclusively to the poorer v6 architecture of tunnels, but also is partially due to other factors, such as the topological location.

Table of Contents

Performance Comparison of v6 versus v4

1 Dataset

Pings were sent to 6,864 globally distributed dual-stack nameservers from three locations in the U.S.: Dallas, TX; San Jose, CA; and Reston, VA. The present report considers measurements for the period from April 12, 2010 through December 19, 2010. For this period, we have 44 million measurements on 9,223 distinct time epochs.

There are time periods where no measurements were collected, most notably April 24 to 25, June 25 to July 1, August 11 to 24, and October 27 to November 8.

For 2,085 of the 6,864 nameservers, the IPv6 interface is a 6to4 tunnel (address 2002::/16) and 33 are a Teredo tunnel (address 2001:0::/32), where these are the two most popular tunneling methodologies currently in use. We have partitioned some of the results below into "tunneled" and "native" based on the IPv6 address of the nameservers. Caution: it is possible that a path to a "native" nameserver does contain a tunnel.

2 Summary Stats by Geo Region

In terms of the summary statistics, Table 1 shows the summary statistics of median, mean, and ninety-fifth percentile of latency over v4 and over v6, conditioned on the geographic region of the nameserver and whether the v6 interface of the nameserver is native, tunneled, or either.

A first observation is that in terms of these summary statistics, the latency is less over v4 than v6. For example, for destinations in the North America, the mean latency is 55 ms over v4 but substantially higher, 101 ms, over v6.

A second observation is that, except for South America, the latency is higher to destinations where the v6 interface is tunneled, as opposed to native, and this pertains for both the v6 and v4 path. For example, for destinations in Asia where the v6 interface is native, the mean latency over v6 was 212 ms. For v6 interfaces that are tunneled, the mean latency over v6 was significantly higher at 317 ms. And also over v4, the mean latency is again higher to destinations where the v6 interface is tunneled, 205 ms versus 245 ms.

That the latency over v6 is higher to tunneled destinations is consistent with common expectations; however, it is somewhat a surprise that the latency over v4 is also higher. How could the v6 interface affect the latency over the v4 path? Admittedly the set of nameservers are distinct; that is, we are comparing v4 latency to one set of destinations, those whose v6 interface is tunneled, with the v4 latency to another set, albeit in the same geographic region. However, this alone does not imply any intrinsic bias and thus does not explain why the latency (in terms of the summary statistics) is consistently higher to one of the two sets. Also, the number of destinations in each set is reasonably large, 2,085 and 4,779, and the cause is not due to a few outliers, as the affect also pertains for the median and 95 percentile. Although I do not have the definitive explanation, a plausible explanation is that the nameservers with a tunneled v6 interface are more likely to be in smaller networks, less well-connected networks, lower down in the hierarchy. Supporting this explanation we have estimates of v4 load from the nameservers, as seen by Akamai, and the load from the nameservers with tunneled v6 interfaces overall is indeed lower.

Geo-Region	Set of Nameservers based on v6 interface	Latency [ms]					
		Median		Mean		95th percentile	
		v4	v6	v4	v6	v4	v6
North America	native	47	86	52	95	101	172
	all	49	92	55	101	108	192
	tunneled	53	101	61	114	119	216
Europe	native	151	162	154	163	218	231
	all	154	166	158	168	224	240
	tunneled	167	182	172	188	252	273
Asia	native	184	198	205	212	359	331
	all	196	215	216	240	367	388
	tunneled	229	313	245	317	378	469
South America	native	183	198	188	208	272	345
	all	176	217	186	235	306	392
	tunneled	172	233	186	246	330	404
Africa	native	344	357	337	350	438	454
	all	348	368	356	379	481	529
	tunneled	355	393	377	415	557	697
Australia	native	208	216	211	232	293	317
	all	210	227	216	244	298	384
	tunneled	225	275	235	288	329	401

Table 1

Regardless of the correct explanation, Table 1 shows that the common observation that v6 latency is higher over tunnels, as compared with native, is not due exclusively to the poorer v6 architecture of tunnels, but also is partially due to other factors, in particular as suggested above, the topological location.

In subsequent results where we split out the v6 measurements based on tunneled or native, we do likewise for v4, which yields a perspective on "the higher v6 latency over tunnels" that is due to the tunnels.

A third observation is that the extent v4 is better than v6 (in the sense of lower latency) is more substantial for destinations with v6 tunnels. For example, for destinations in Asia, the reduction in the median latency is 84 ms (229 minus 313) given tunneled destinations, and is only 14 ms (184 minus 198) given native destinations. This affect is more noticeable for Asia, South America, Africa, and Australia, than for North America and Europe, though it is still present.

The following Table 2 provides the summary statistics on packet loss. The median values are omitted from the table as they were all 0, except for v6 paths to tunneled interfaces in Africa, where the median was 0.3%. Observations:

Geo-Region	Set of Nameservers based on v6 interface	Percent Packet Loss [0, 100]			
		Mean		95th percentile	
		v4	v6	v4	v6
North America	native	0.4	2.1	0.5	6.6
	all	0.6	3.2	1.2	23.2
	tunneled	1.0	5.2	3.2	32.4
Europe	native	0.5	0.7	1.3	2.0
	all	0.7	1.8	2.2	10.3
	tunneled	1.4	6.0	6.5	31.8
Asia	native	3.0	1.2	20.8	6.3
	all	2.8	2.7	19.0	18.6
	tunneled	2.2	6.9	11.4	34.8
South America	native	0.8	1.0	4.8	4.7
	all	1.7	4.4	8.6	27.9
	tunneled	2.1	5.7	9.8	31.6
Africa	native	2.0	4.8	9.5	32.8
	all	2.8	6.7	12.4	40.0
	tunneld	3.9	9.0	19.1	43.4
Australia	native	0.4	1.0	1.6	5.0
	all	0.6	2.0	2.9	11.5
	tunneled	1.4	5.6	7.8	31.3

Table 2

The packet loss over v4 is less than over v6, except for destinations in Asia, where interestingly the loss is higher over v4 to those nameservers with a native v6 interface.

The 95th percentile of packet loss is quite high at the v6 tunneled interfaces.

The packet loss to the v4 interface is lower to the set of nameservers whose v6 interface is native, as opposed to the set of nameservers whose v6 interface is tunneled, except again for Asia. This is consistent with the heuristic explanation above that nameservers with tunneled interfaces tend to be in networks that are smaller, further down in the hierarchy.

Appendix A contains plots of the complementary distribution functions of packet loss. These plots emphasize the points above.

3 Comparison of V6 and V4 Latency: Distributions

The following two plots show the distribution of the difference: IPv6 latency minus IPv4 latency, partitioned by geo region. Positive values on the x-axis correspond to v6 latency being greater than v4. The range on the x-axis was chosen so as to highlight most of the action, although it cuts-off each end of the distributions. If the full range were shown, the distributions would go from 0 to 1.

For example, consider the first plot and the distribution given destinations in Europe. At the zero point on the x-axis, the distribution is 0.29. Thus, for 29% of the measurements the latency over v6 was less than or equal to that over v4. Likewise for 71% of the measurements, the latency over v6 was greater. Looking at the -10 ms on the x-axis, we find that for 11% of the measurements, the latency of v6 was smaller (better) by at least 10 ms. Considering the ticks at -10 ms and 10 ms, for 48% (59% minus 11%) of the measurements, the latency over v6 was within 10 ms of that over v4. If one didn't care about latency differences within 10 ms, then for about half of the time, one would be indifferent (considering only this factor) between the two protocols.

As first-order summary, the plots show that more often the latency is greater over v6. However, there will be given clients for which this is not the case. In the context of optimizing performance, one would ideally like to be able to distinguish which would be better.

Figure 1. Distribution of difference in v6 versus v4 latencies

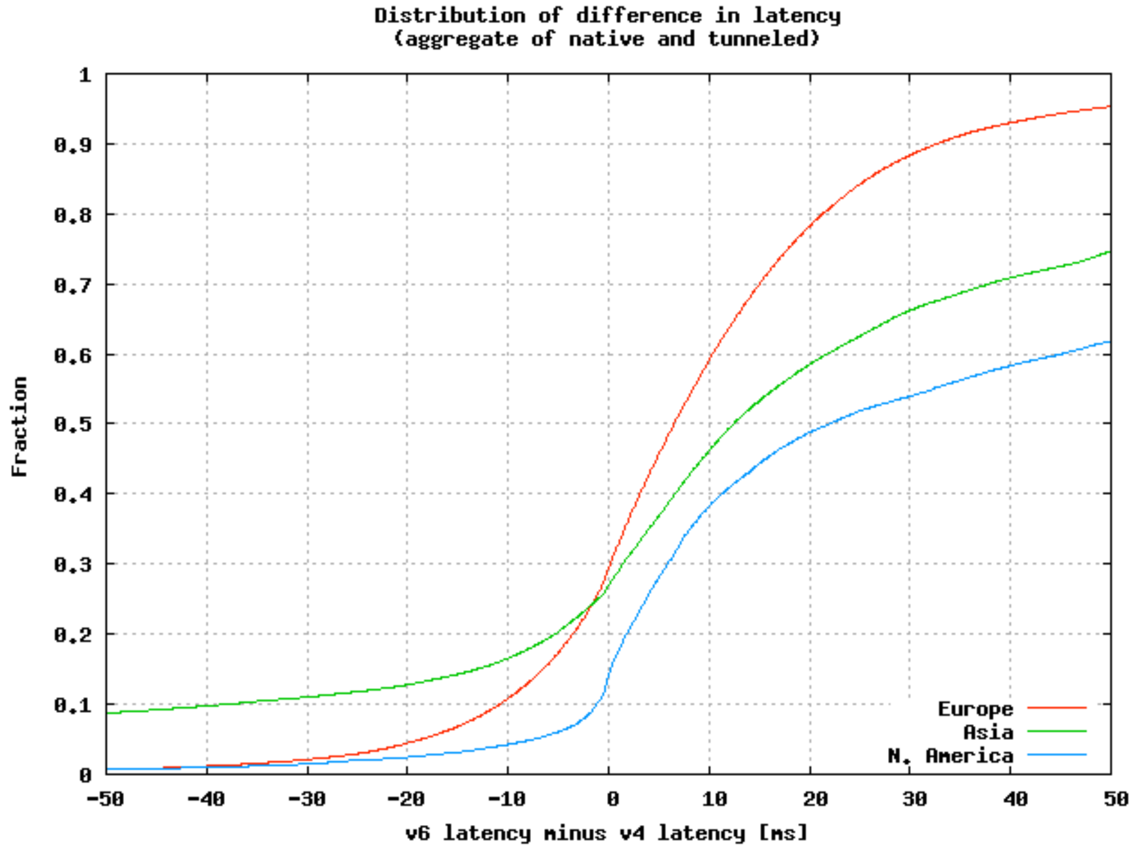
Distribution of difference in latency
(aggregate of native and tunneled)
Fraction
v6 latency minus v4 latency [ms]
Europe
Asia
N. America
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

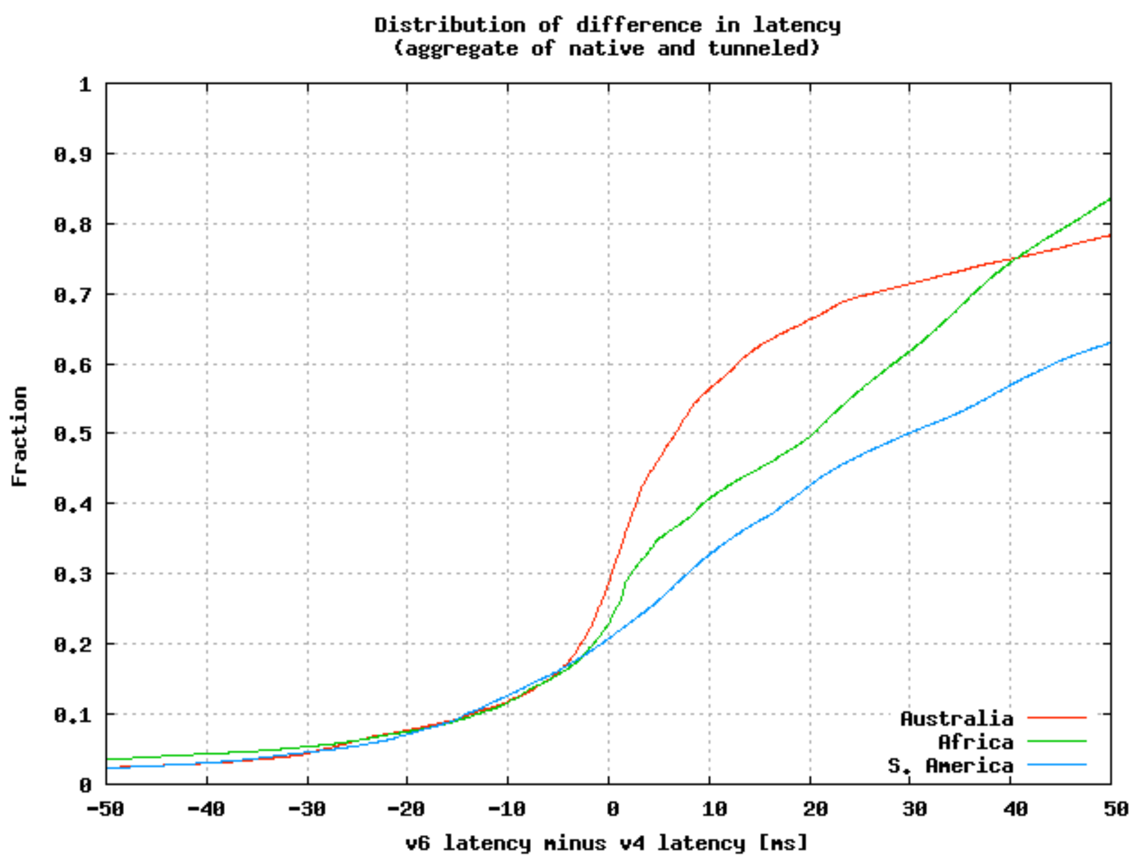
Distribution of difference in latency
(aggregate of native and tunneled)
Fraction
v6 latency minus v4 latency [ms]
Australia
Africa
S. America
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

The following two plots provide another viewpoint on the same data by considering the ratio of v6 latency divided by v4 latency and where the x-axis is on a log scale.

Figure 2. Distribution of ratio of latencies

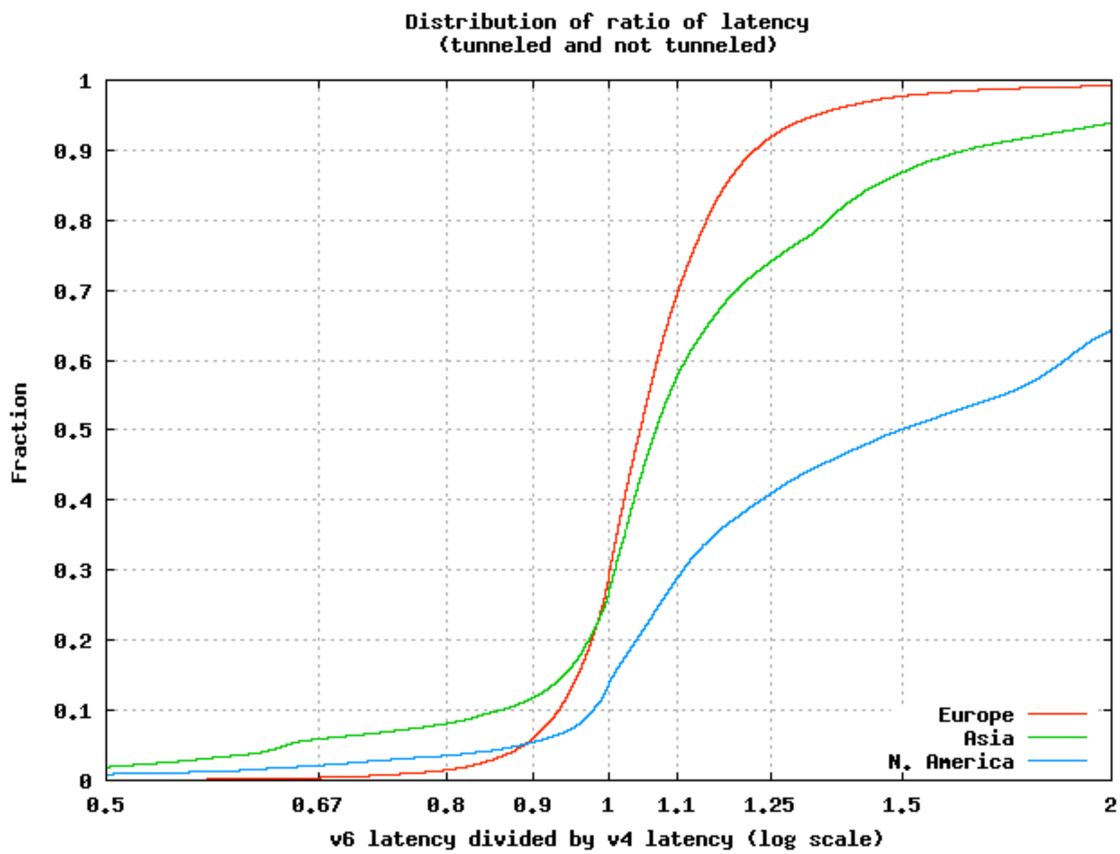

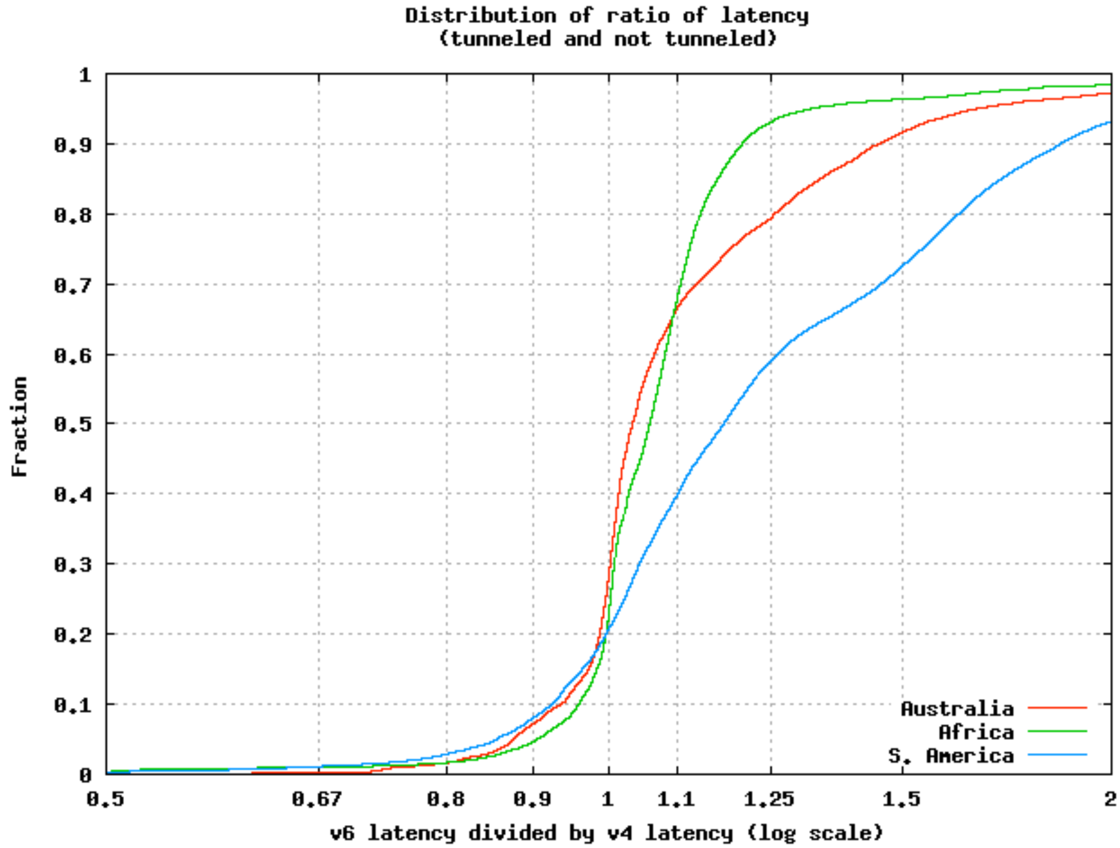

Appendix B contains analogous plots, except partitioned by tunneled versus native.

4 Comparison Across Time

This section compares v4 versus v6 latency and loss, viewed across time.

4.1 April to December, 2010

In the following Figure 3, each plotted point is the mean over a 24-hour period. (Recall that there are gaps in the measurement data as is evident in the plots.) There are six plots of latency, one for each geographic area, followed by six plots of packet loss. Note that the range on the y-axis varies from plot to plot.

The most obvious feature of the latency plots is the clear variation across time. As each point is the average over 24 hours, the variation in latency is on longer time scales than daily variation (which is considered in the next section). Sometimes higher latency can persist for months, as for example over v6 to tunneled destinations (the green line) in North America, Asia and South America. Also, all of the geographic regions have spikes in latency. Note that sometimes the variation is strongly correlated across the four cases, as for Europe, and sometimes not, as for North America. These observations suggest that a more detailed study could make inferences as to *where* congestion was occurring.   This could be possible future work for the IPv6 project.

Consider now the difference in latency across the four cases within each geographic region. In each geographic region, the highest latency is on the v6 path to tunneled destinations, which is consistent with the summary statistics in Section 2. Likewise, the lowest latency is on the v4 path to destinations whose v6 interface is native, except for South America. Consider the difference between the green and blue lines, i.e. the latency over v6 to tunneled versus native destinations. Some of the difference is due to the well-known poorer architecture of v6 tunnels and, some, as suggested in Section 2, is likely due the tunneled destinations being in networks that are smaller, further down in the hierarchy (call this second factor "inferior topology").   As a rough estimate of this second factor, we can use the difference between the *v4* latency to destinations whose v6 interfaces are tunneled versus native, i.e. the difference between the yellow and red lines in the plots. For example, in North America, the space between the yellow and red lines is fairly small and so is the space between the blue and yellow lines until roughly Oct. 1. The subsequent jump in the green line is probably due to something regarding the tunnels, as opposed to the inferior topology, as the yellow line remains relatively flat.

It is worth emphasizing that the difference between the yellow and red lines is just a rough estimate of the "inferior topology" factor: for example, in the North-America plot there are points where the difference between the yellow and red lines is greater than that between the green and blue, and thus makes no sense to be viewed as a piece of the latter.

In the Europe plot, the space between the yellow and red lines is about equal to that between the green and blue thus suggesting the poorer performance with v6 tunnels, versus native, is due to the inferior topology. Notice also that the four lines are quite in sync with each other, which suggests that the variation in latency is due to congestion on facilities that are shared by all four cases, such as trans-oceanic optical cables.

The remaining geographic regions can be viewed with the above comments in mind. In addition, note that in the Asia plot, the red line is significantly more variable than the yellow (though is still lower). I haven't thought of a likely explanation though a further examination of the data might be illuminating. Notice also that in Asia the blue line is sometimes below the red - the latency over v6, to native destinations, is sometimes faster than over v4. This is examined more thoroughly in Appendix B.

In South America, the red and yellow lines roughly overlap, and, as reported in Table 1 of Section 2, the v4 latency to destinations whose v6 interface is native is actually a bit higher than to destinations whose v6 interface is tunneled.   Thus, this data does not support the supposition of "inferior topology" for tunneled interfaces for South America.

Figure 3. Time History of Latency, April – Dec. , 2010

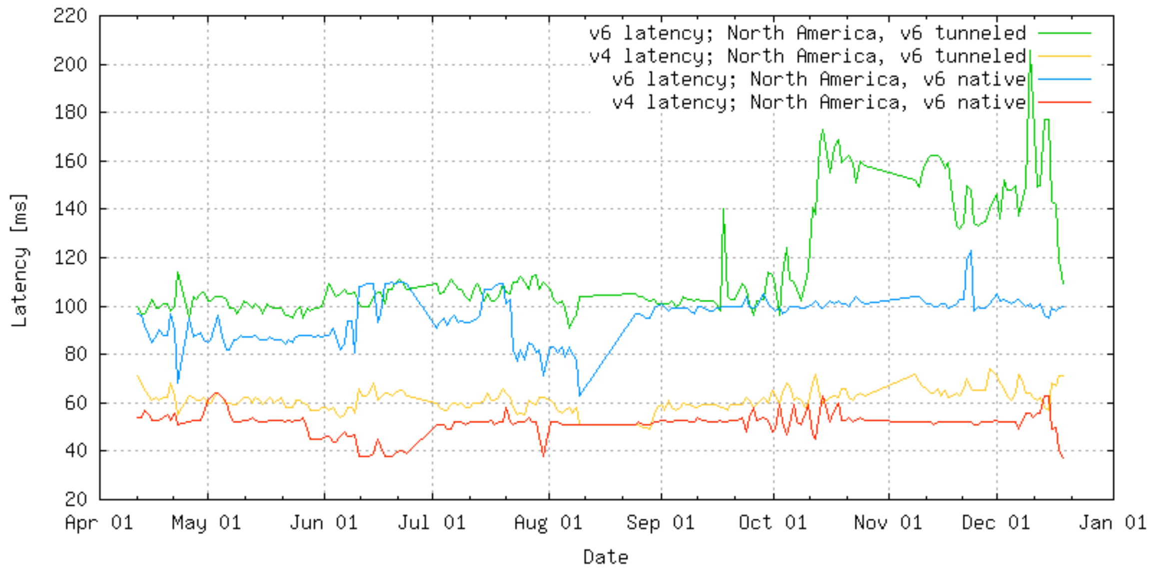

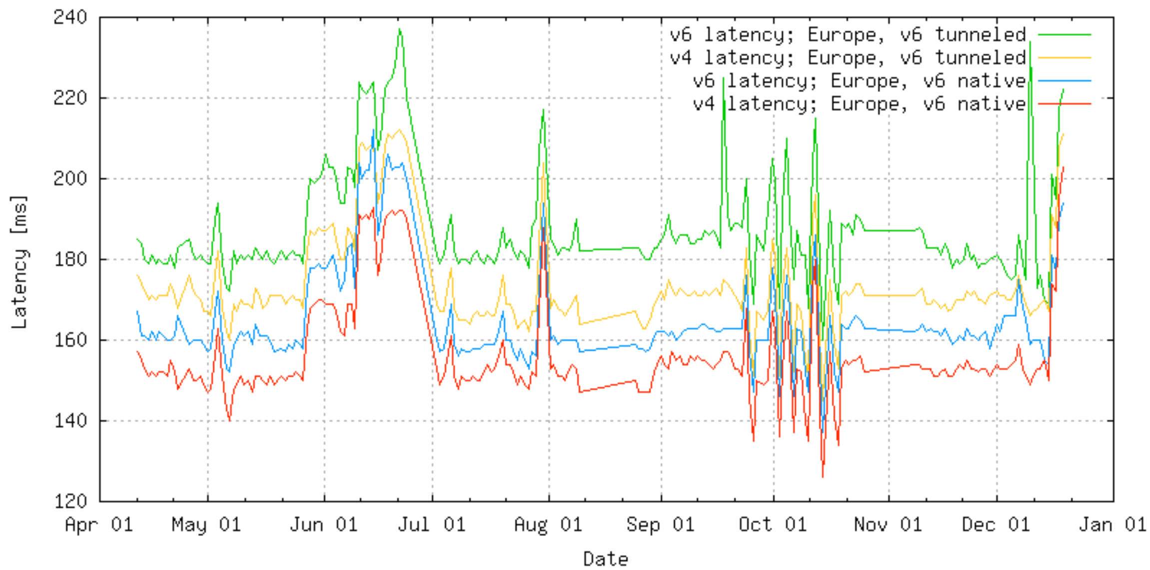

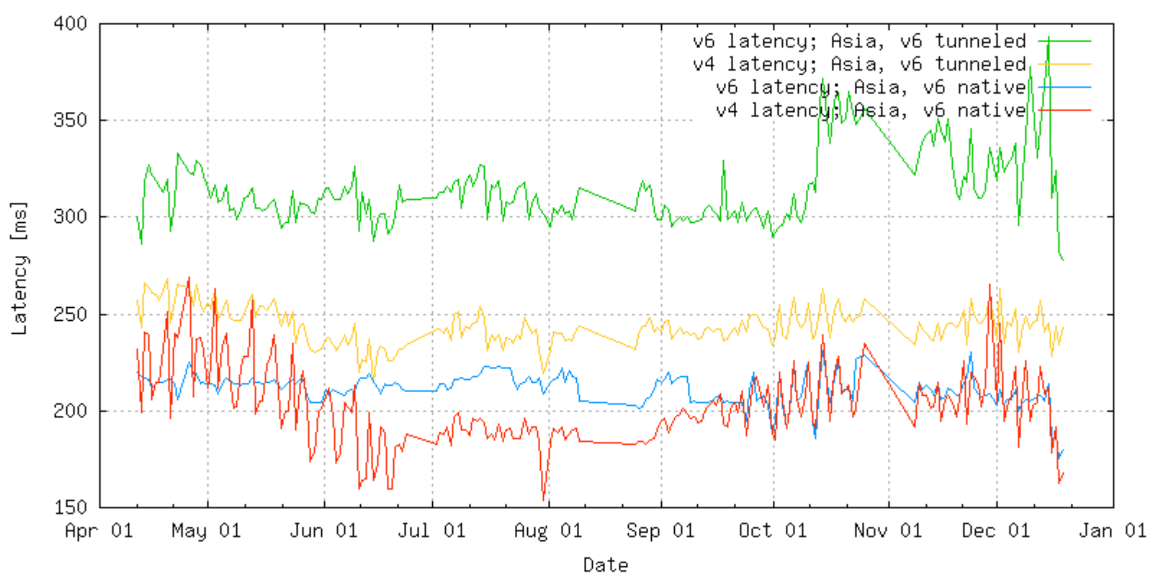

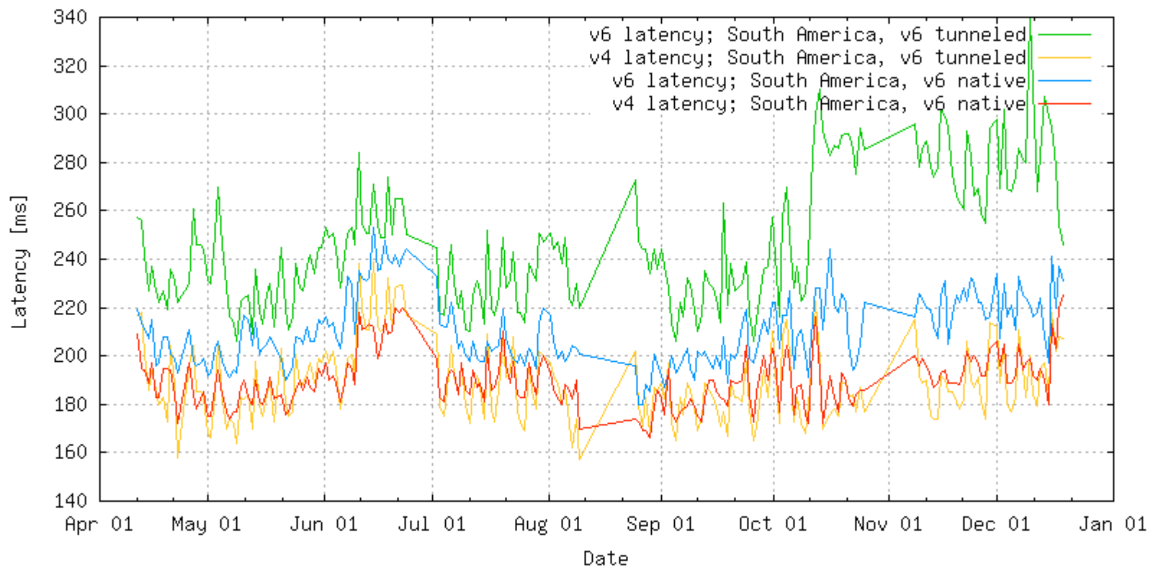
v6 latency; South America, v6 tunneled
v4 latency; South America, v6 tunneled
v6 latency; South America, v6 native
v4 latency; South America, v6 native
Latency [ms]
340
320
300
280
260
240
220
200
180
160
140
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01
Date

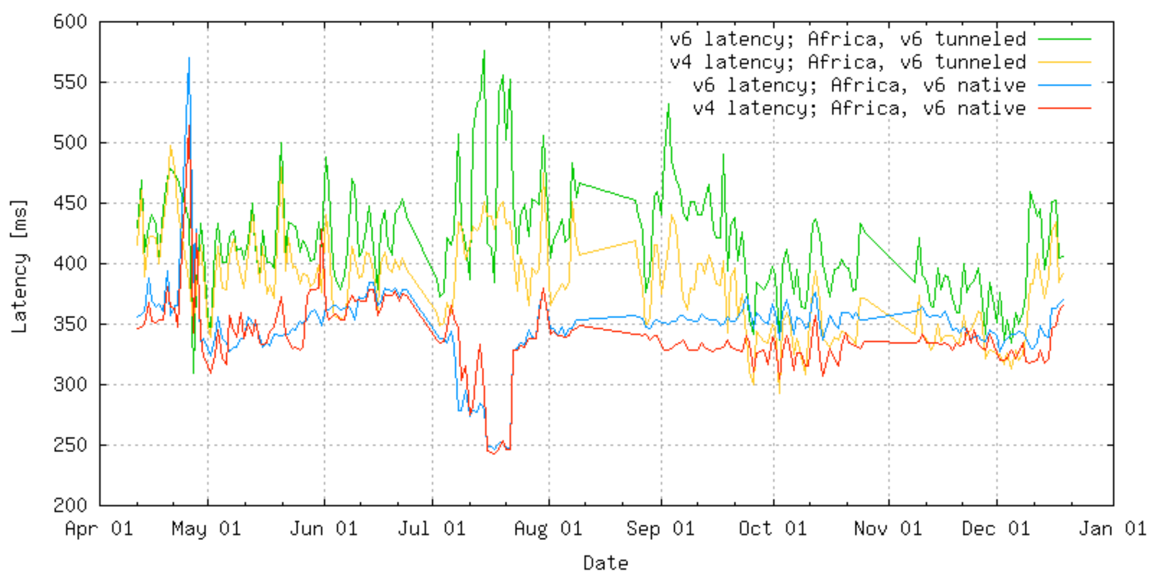
v6 latency; Africa, v6 tunneled
v4 latency; Africa, v6 tunneled
v6 latency; Africa, v6 native
v4 latency; Africa, v6 native
Latency [ms]
600
550
500
450
400
350
300
250
200
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01
Date

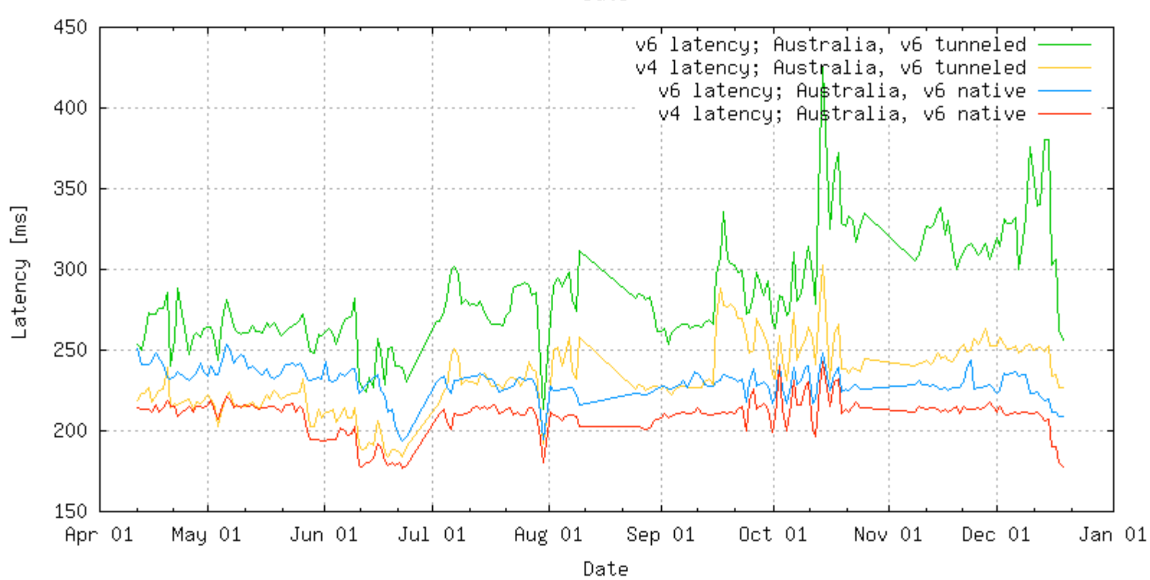
v6 latency; Australia, v6 tunneled
v4 latency; Australia, v6 tunneled
v6 latency; Australia, v6 native
v4 latency; Australia, v6 native
Latency [ms]
450
400
350
300
250
200
150
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01
Date

The following are the analogous six plots for packet loss. Note that the range on the y-axis varies.

The most striking feature is that there are periods of high loss. From September through December, there is high v6 loss to tunneled destinations in all geographic regions. All three origin regions had high loss. Thus the cause was broad enough to affect multiple origins. A possible explanation is that for all three origins, the route to the anycast 6to4 address, 2002::/16, led to relays that were overloaded, and that this condition persisted. Anycast routing is based on BGP, which does not adapt to congestion. As such, a network operator would need to intervene and change policy.

Also note that during June and July, there was high v6 loss to native destinations in one of the geographic regions: North America. In this case, just one of three origin regions had the high loss, thus the cause was localized.

Note that the loss to over v4 and over v6 to destinations with v6 tunneled interfaces (yellow and green lines) is higher than to destinations with v6 native interfaces, except for Asia where the reverse pertains for periods for v4. (This is consistent with Table 1 in Section 2.) I don't have an explanation for this exception in Asia.

Figure 4. Time History of Packet Loss, April – Dec., 2010

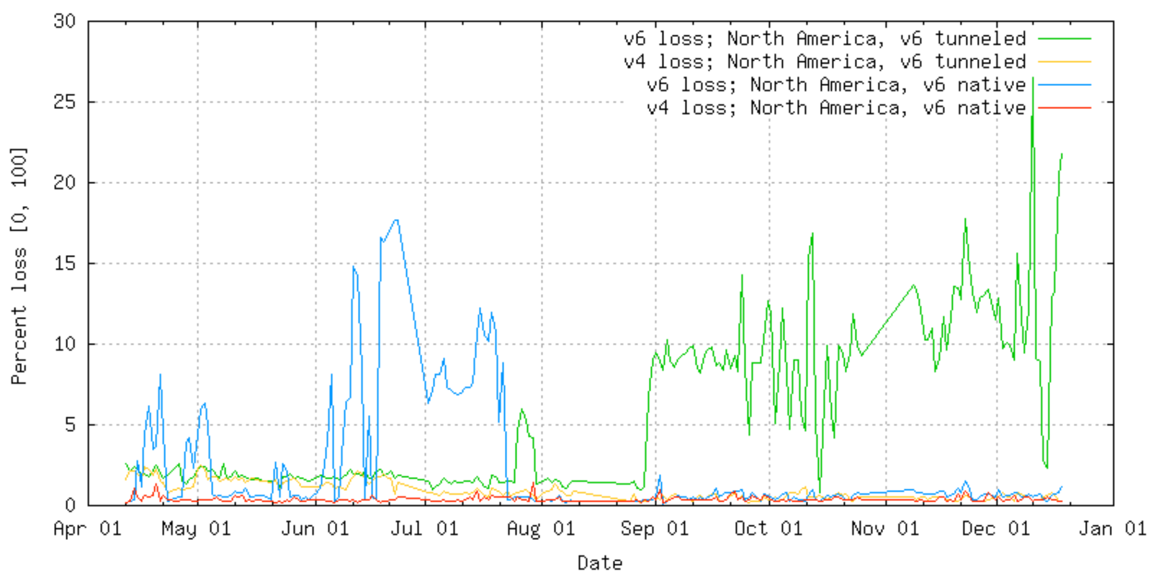

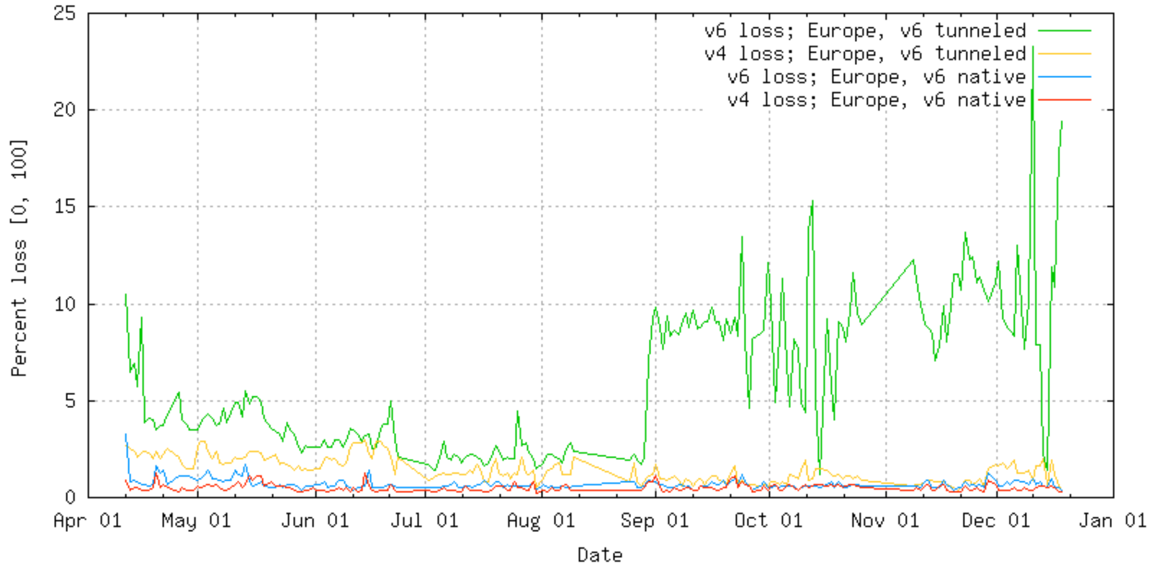
v6 loss; Europe, v6 tunneled
v4 loss; Europe, v6 tunneled
v6 loss; Europe, v6 native
v4 loss; Europe, v6 native
Percent loss [0, 100]
Date
0
5
10
15
20
25
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01

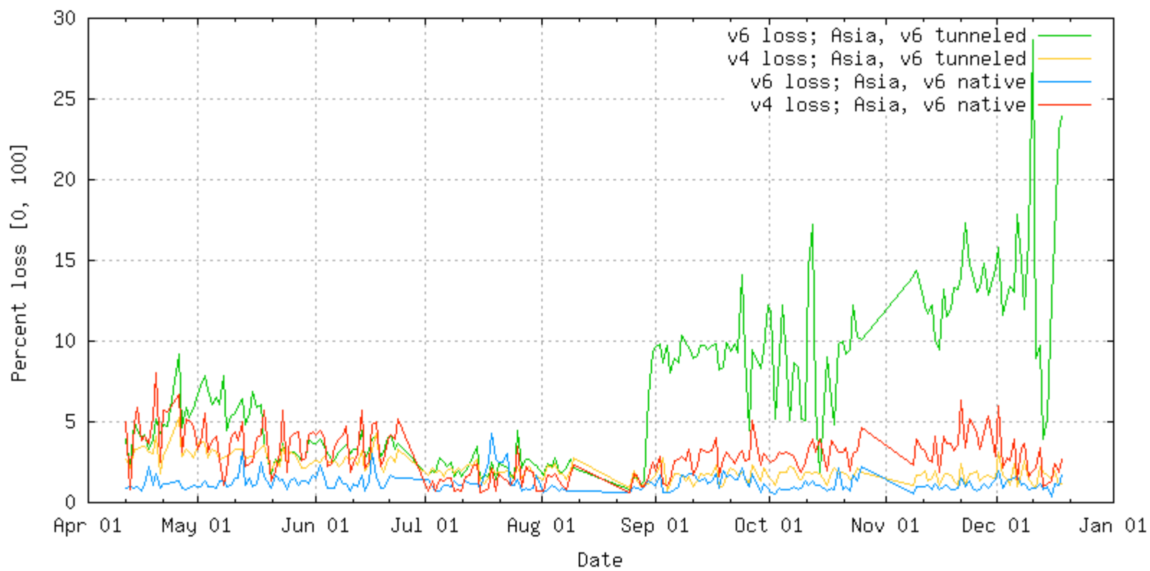
v6 loss; Asia, v6 tunneled
v4 loss; Asia, v6 tunneled
v6 loss; Asia, v6 native
v4 loss; Asia, v6 native
Percent loss [0, 100]
Date
0
5
10
15
20
25
30
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01

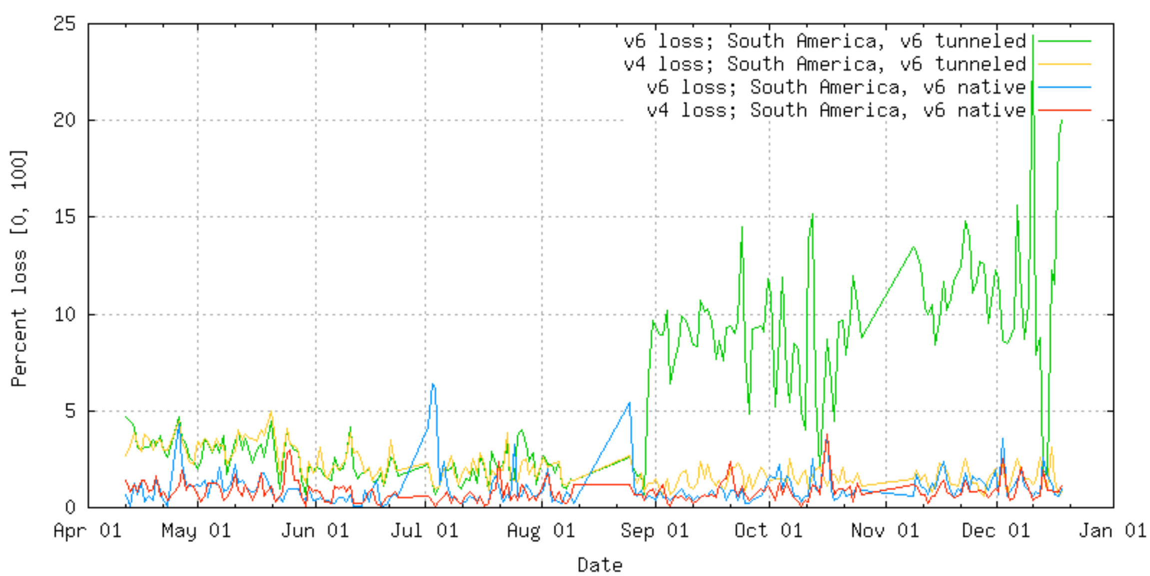
v6 loss; South America, v6 tunneled
v4 loss; South America, v6 tunneled
v6 loss; South America, v6 native
v4 loss; South America, v6 native
Percent loss [0, 100]
Date
0
5
10
15
20
25
Apr 01
May 01
Jun 01
Jul 01
Aug 01
Sep 01
Oct 01
Nov 01
Dec 01
Jan 01

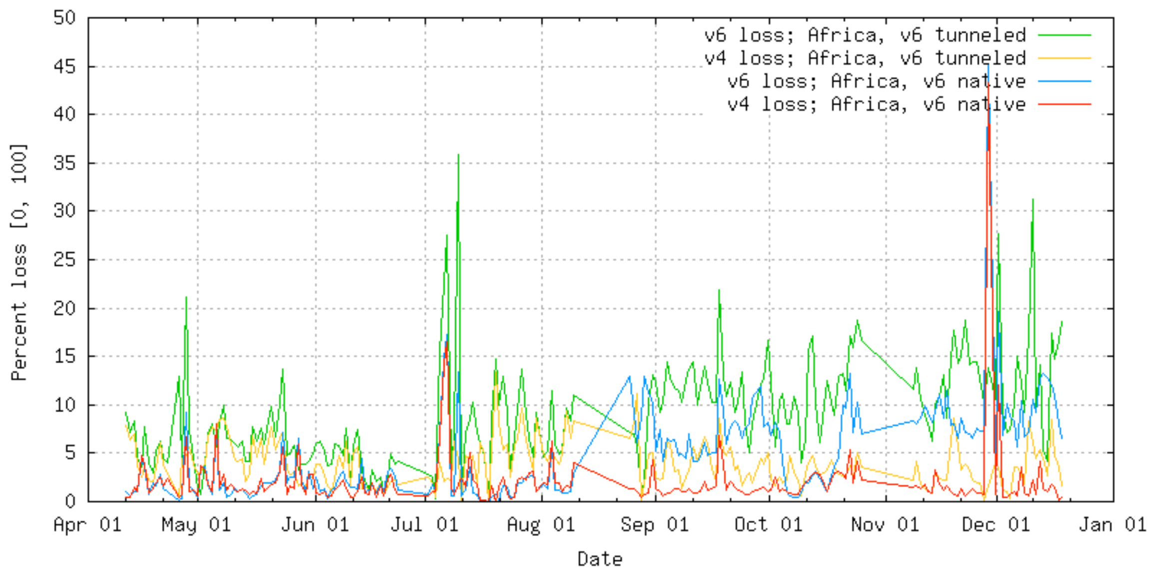

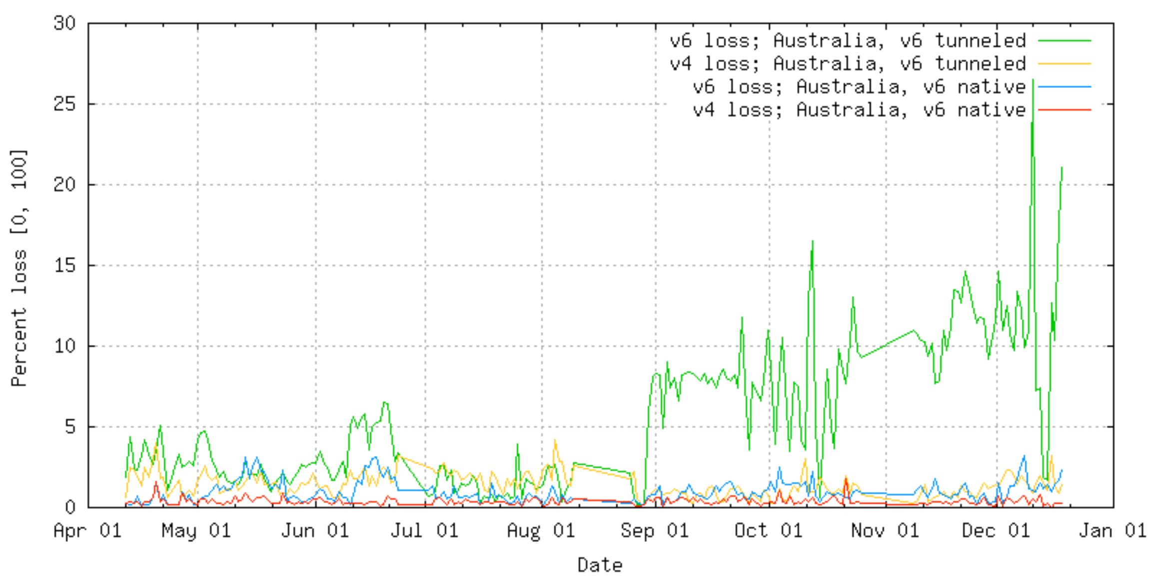

4.2 July 12 through 14, 2010

To get a sense of an hour-of-day pattern, if any, we plot a three-day period in July. Each plotted point is the mean over a one-hour period. Again there are six plots of latency, followed by six plots of packet loss. The range on the y-axis varies from plot to plot.

Daily variation in latency and loss is typically due higher loads and thus increased congestion during the busy period of the day. Relatively constant latency and loss over the course of the day typically is indicative of either constant (light or heavy) load, or variable load that however remains light.

The latency plot for South America shows strong daily variation for all four cases, indicating congestion on the paths between there and the US. For destinations in

Europe, the daily variation in latency is also evident, though less extreme. For North America the daily variation is a slight, if any. For Australia, none is apparent.

The latency plot for Asia is more curious: there is strong daily variation for three of the four cases, but the latency is rather constant for traffic over v6 to native destinations. Possibly the routes from the three origins to the native v6 addresses in Asia are on an uncongested trans-oceanic channel though I would be surprised if a fiber channel were reserved for v6 traffic, but maybe it is. (The routes to the tunneled v6 addresses most likely go via anycast 6to4 relays in the U.S., in which case the trans-oceanic hop in all likelihood is over v4.)

The latency plot for Africa in one sense is the reverse of Asia's – three of the cases have little to no daily variation, and one does: v6 to tunneled destinations. A possibility is that at the far-end of the tunnel, the v6 networks are resource constrained. (Although latency over v4 to destinations with native v6 interfaces is also variable, the red line, the variability is not so much in a daily pattern.)

If there is interest, future work could investigate the above conjectures and in general investigate the causes for various behavior displayed in these plots.

As a final note: when I examined other three-day intervals, I would sometimes see the same patterns as here and sometimes see different ones. What pertained for one three-day interval need not pertain months later.

Figure 5. Time History of Latency, July 12 -14 , 2010

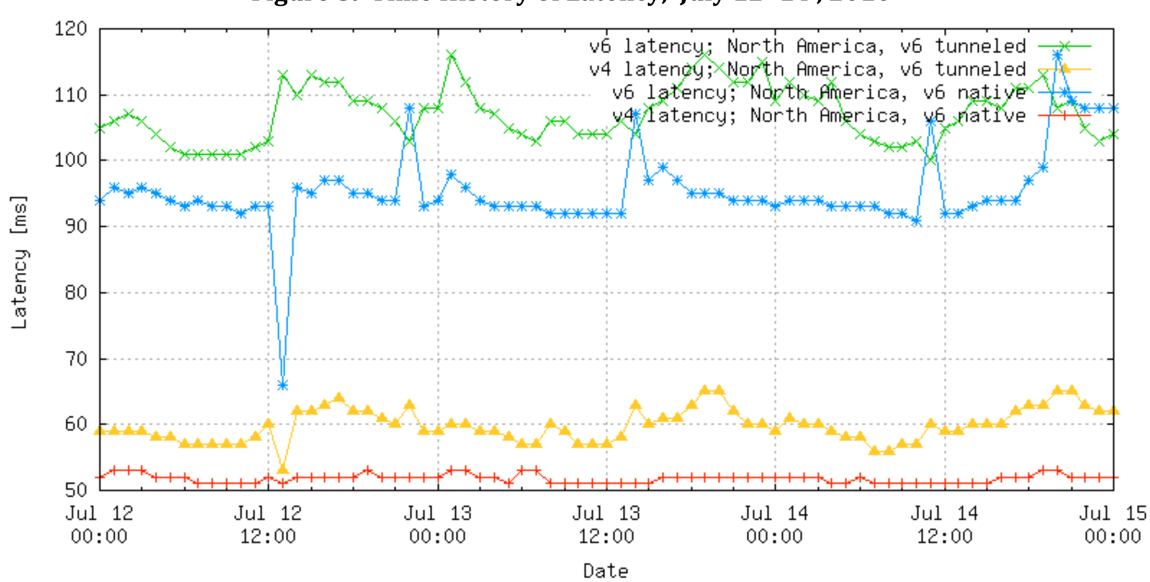

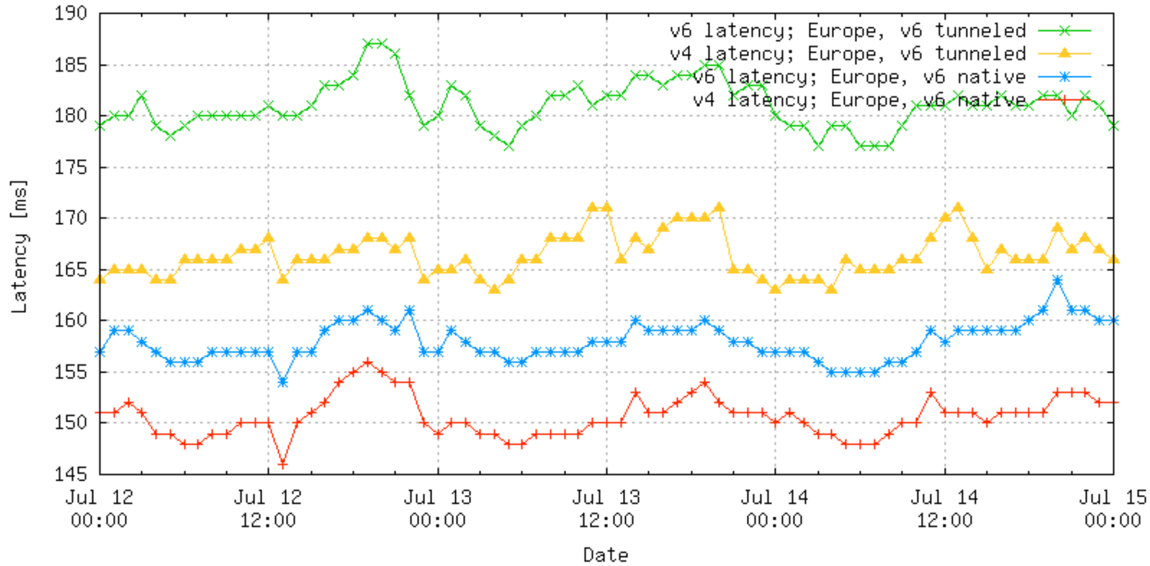
v6 latency; Europe, v6 tunneled
v4 latency; Europe, v6 tunneled
v6 latency; Europe, v6 native
v4 latency; Europe, v6 native
Latency [ms]
Date
190
185
180
175
170
165
160
155
150
145
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

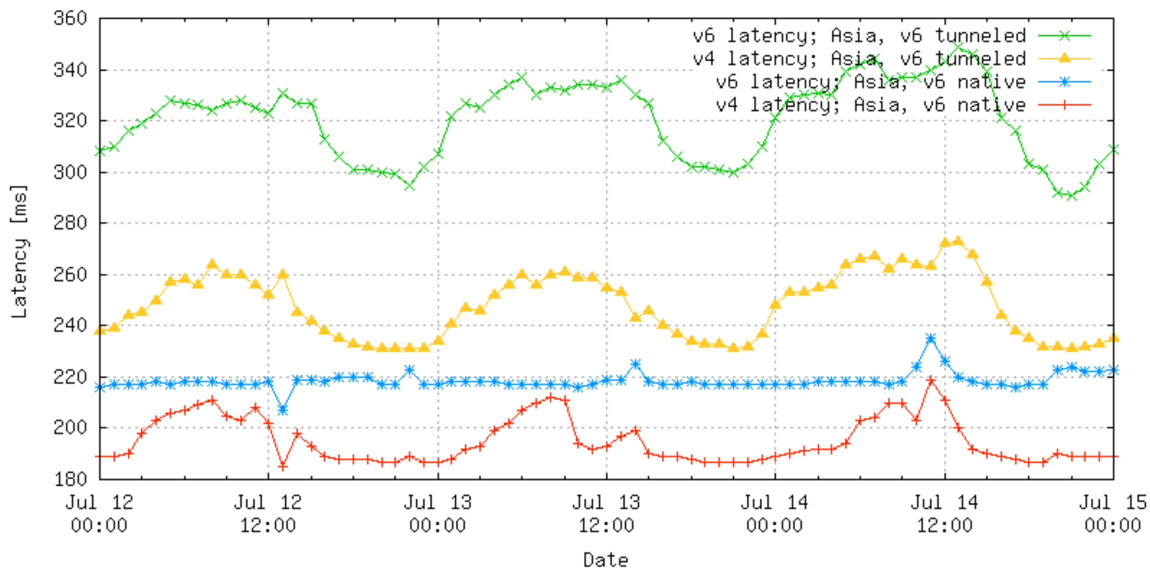
v6 latency; Asia, v6 tunneled
v4 latency; Asia, v6 tunneled
v6 latency; Asia, v6 native
v4 latency; Asia, v6 native
Latency [ms]
Date
360
340
320
300
280
260
240
220
200
180
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

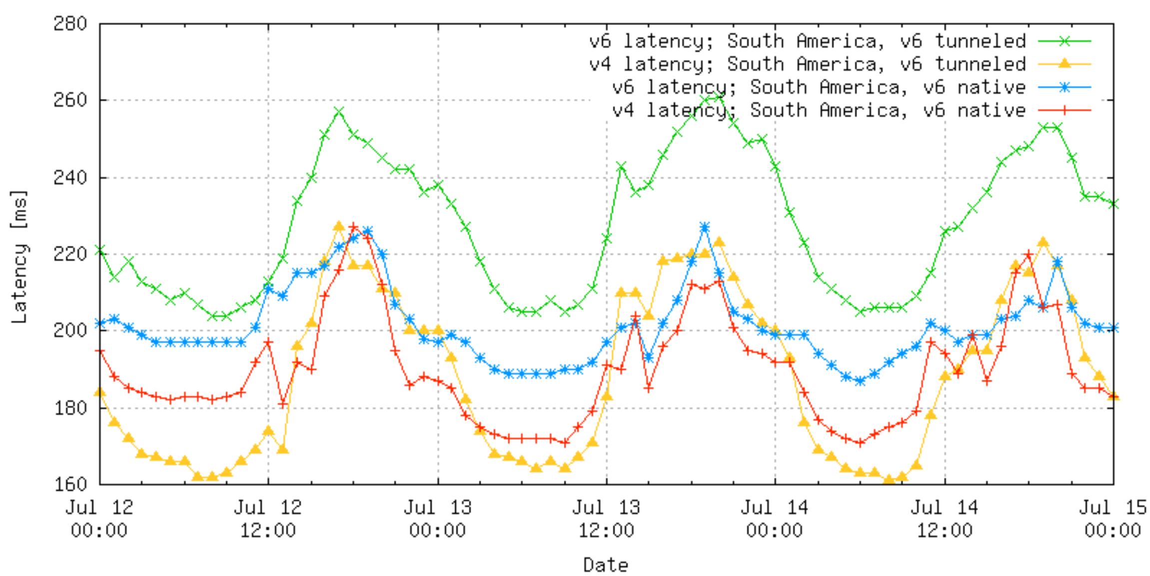
v6 latency; South America, v6 tunneled
v4 latency; South America, v6 tunneled
v6 latency; South America, v6 native
v4 latency; South America, v6 native
Latency [ms]
Date
280
260
240
220
200
180
160
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

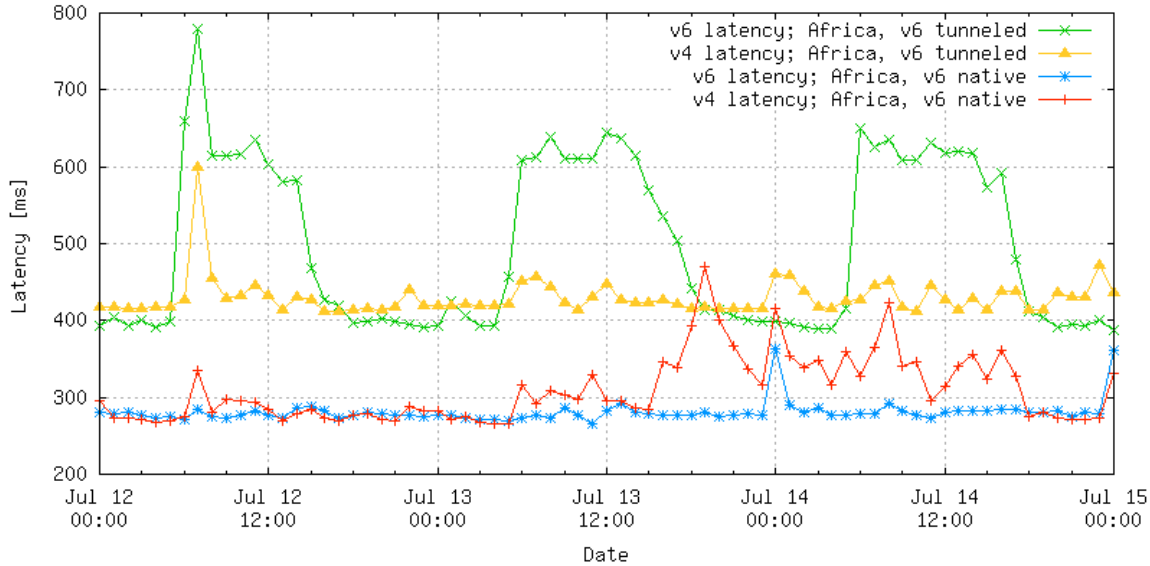

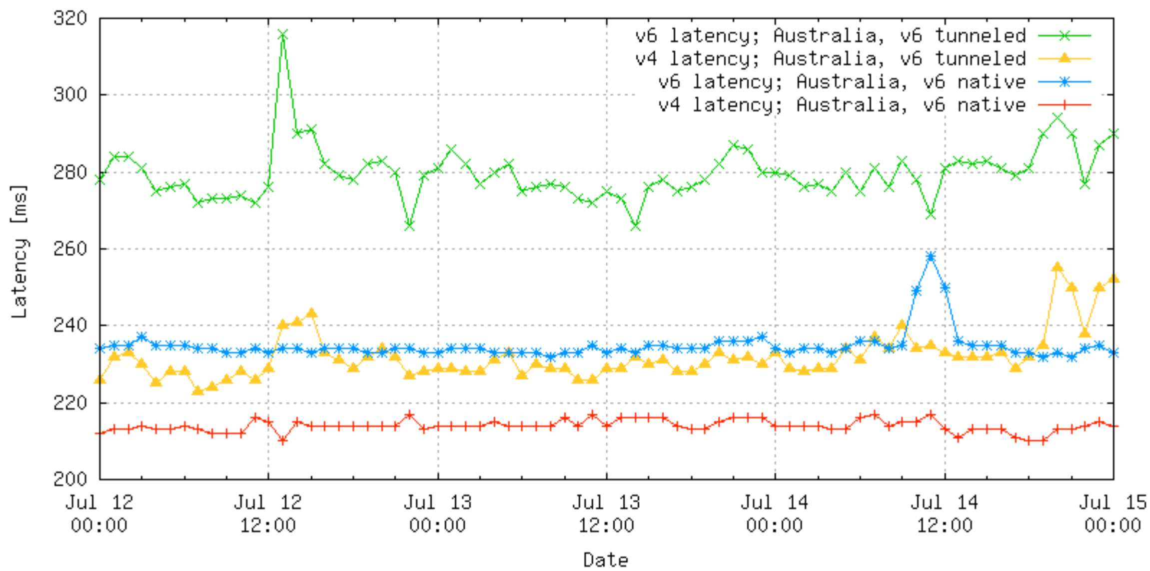

The following are the corresponding plots for packet loss. Note that the range on the y-axis varies from plot to plot

As with latency, South America has very evident daily variation in loss. Again as with latency, a daily variation in loss is also evident for Europe, though more modest. North America has no evident daily variation; also, as discussed above in Section 4.1, the relatively high v6 loss to native destinations in North America is due to just one of the three origins.

Asia has daily variation in loss for all four cases, including, though more modest, v6 probes to native v6, whose latency had been flat.

Africa is notable for high spikes in loss.

Australia, which had no daily variation in latency, has noticeable variation in loss, that somewhat follows a daily pattern.

As a rough summary, the presence or absence of daily variation in loss mirrors that of latency.

Figure 6. Time History of Packet Loss, July 12 -14 , 2010

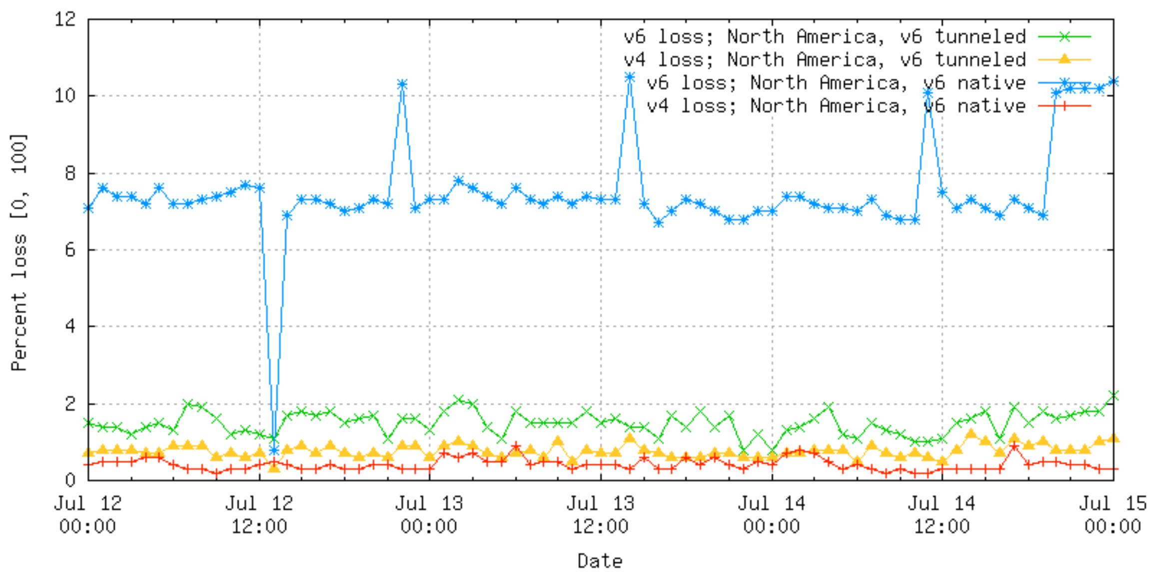

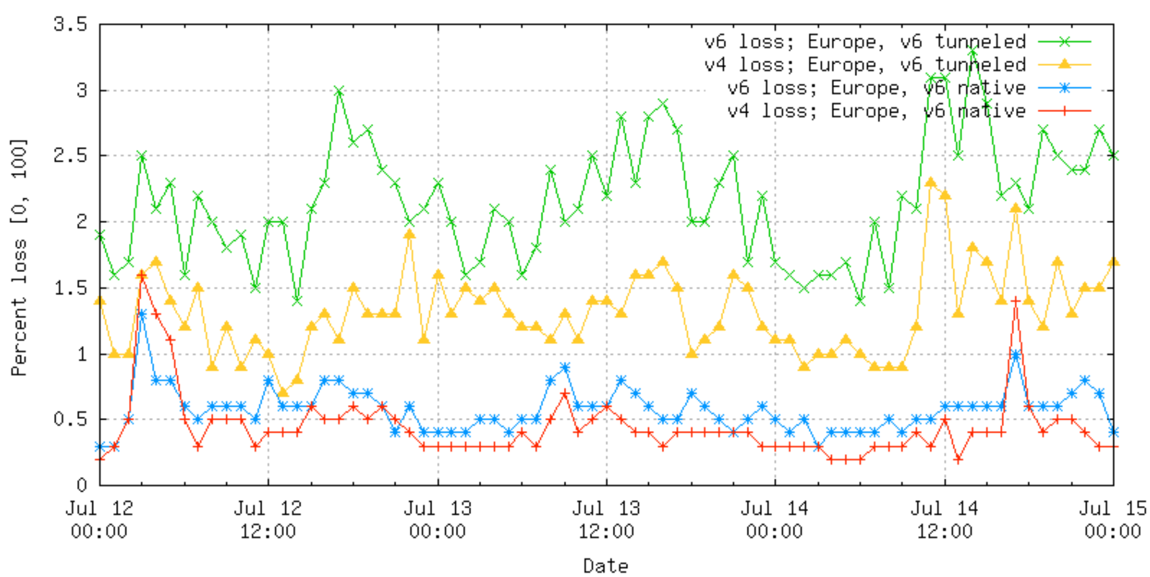

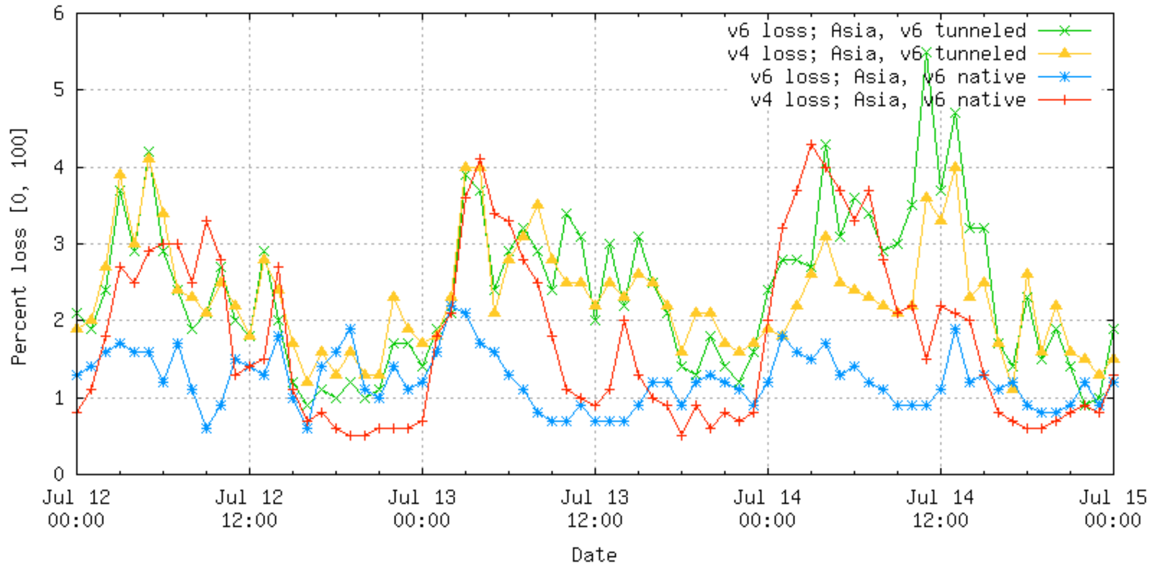
v6 loss; Asia, v6 tunneled
v4 loss; Asia, v6 tunneled
v6 loss; Asia, v6 native
v4 loss; Asia, v6 native
Percent loss [0, 100]
Date
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

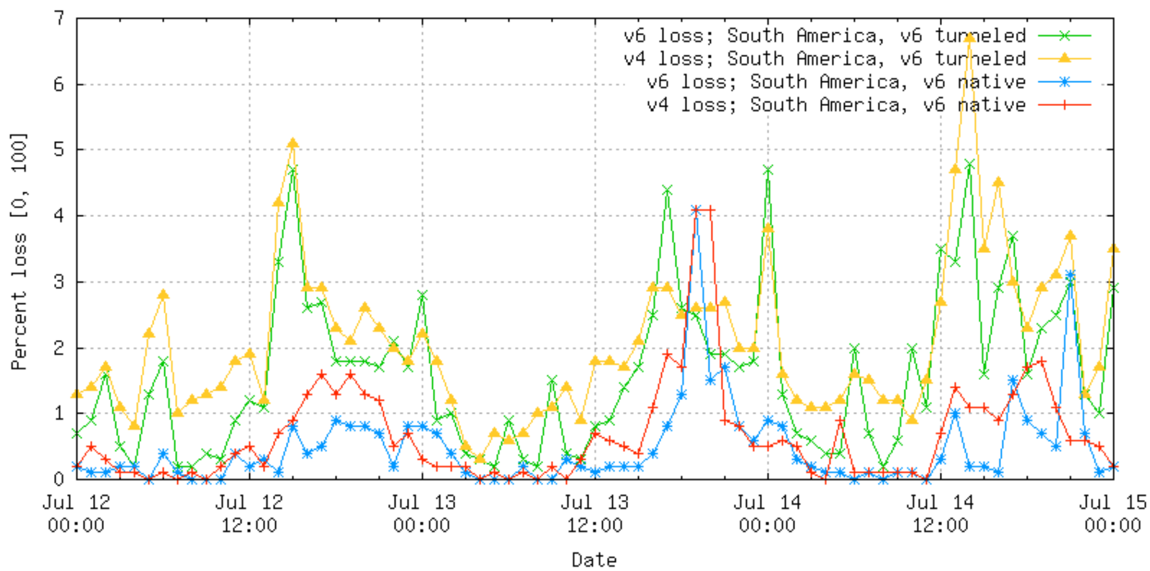
v6 loss; South America, v6 tunneled
v4 loss; South America, v6 tunneled
v6 loss; South America, v6 native
v4 loss; South America, v6 native
Percent loss [0, 100]
Date
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

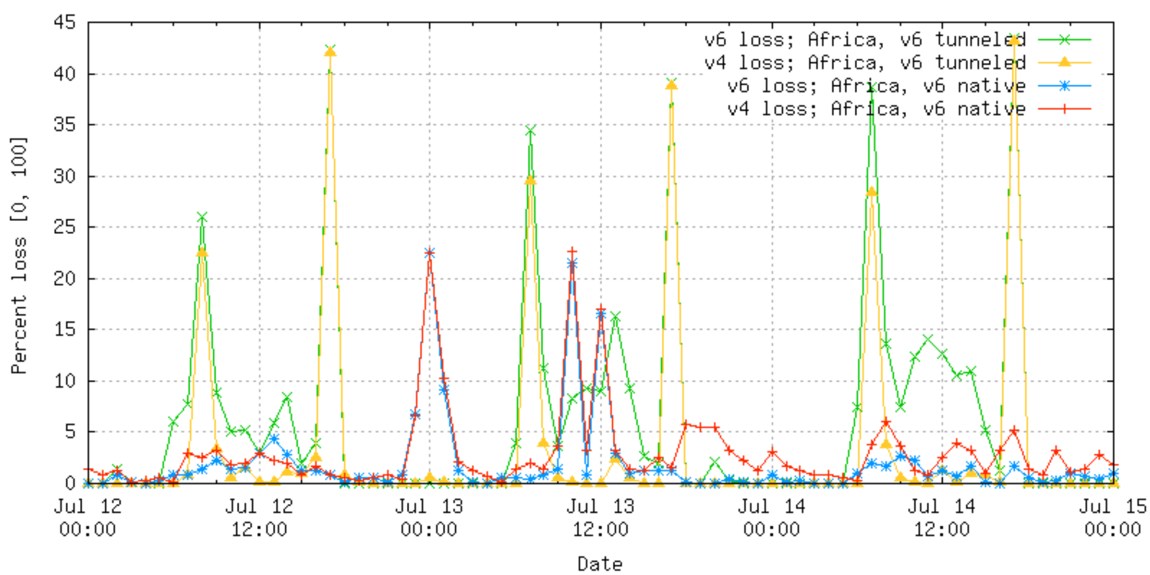
v6 loss; Africa, v6 tunneled
v4 loss; Africa, v6 tunneled
v6 loss; Africa, v6 native
v4 loss; Africa, v6 native
Percent loss [0, 100]
Date
Jul 12 00:00
Jul 12 12:00
Jul 13 00:00
Jul 13 12:00
Jul 14 00:00
Jul 14 12:00
Jul 15 00:00

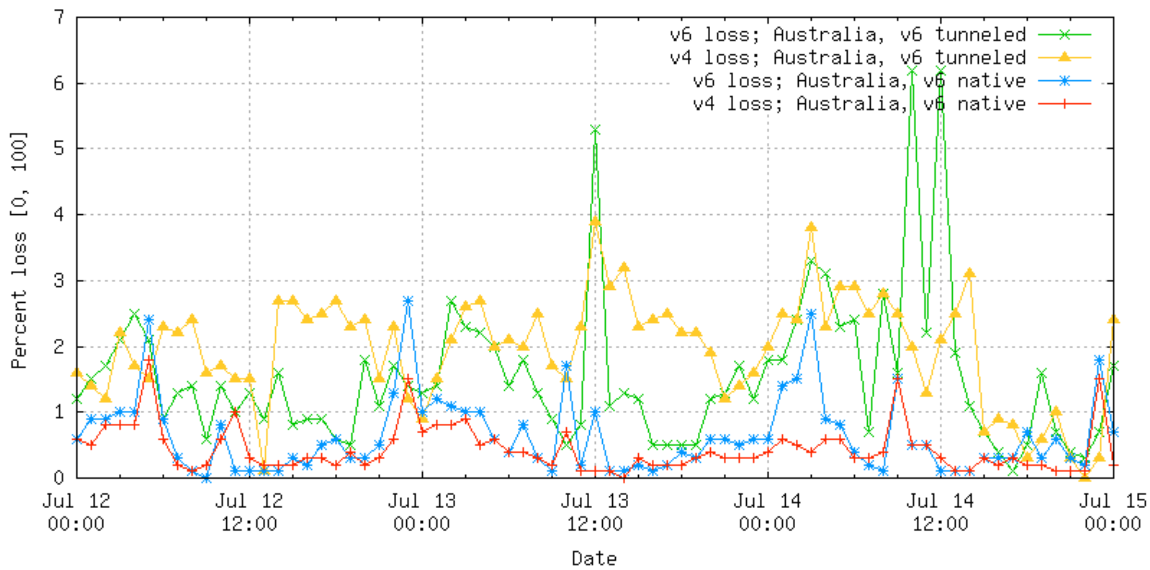

5 Related Work

Here is a sampling, in reverse chronological order, of studies that compare v4 and v6 performance.

Narayan et al. [1] on a testbed, compare v4 and v6 on Windows Vista and Ubuntu. They find that v4 had slightly higher throughput. v6 had significantly higher latency on Ubuntu, as compared with Vista. Law et al. [2] probe from a location in Hong Kong to 2,000 dual-stack, global hosts. v6 had lower hop-counts and higher RTT and higher throughput. RTT was 40% higher on tunneled v6 versus native v6. Zhuo et al. [3] used 26 test boxes of RIPE, globally distributed, though concentrated in Europe, and 600 end-to-end paths. They report that IPv6 has higher loss and latency, mainly due to tunneling. Siau et al. [4] in a large-scale network environment find a minor degradation in throughput of TCP, a slightly higher throughput of UDP, a lower packet loss rate and a slightly longer round trip time over v6 as compared with v4. Zhou et al. [5] report that v6 paths had larger delay variation, and longer delay. From about 1,000 dual-stack web servers in 44 countries, Wang et al. [6] found that v6 connections tend to have smaller RTTs, but suffer higher packet loss. The authors also find that tunneled paths do not show a notable degraded performance compared with native. Cho et al. [7] introduce techniques for identifying IPv6 network problems at dual-stack nodes. They find that IPv6 network quality can be improved by fixing a limited amount of erroneous settings. ARIN and RIPE have links to various measurement studies, [8,9].

6 References

[1] S. Narayan, P. Shang & N. Fan, "Performance Evaluation of IPv4 and IPv6 on Windows Vista and Linux Ubuntu," Inter. Conference on Networks Security, Wireless Communications and Trusted Computing, 2009

[2] Y.N. Law, M.C. Lai, W. Tan & W. Lau, "Empirical performance of IPv6 vs. IPv4 under a dual-stack environment," IEEE Inter. Conference on Communications, 2008.

[3] X. Zhou, M. Jacobsson, H. Uijterwaal & P. Mieghem, "IPv6 delay and loss performance evaluation," Inter. J. of Communication Systems, Vol. 21, 2008.

[4] W.L.Siau, Y.F. Li, H.C. Chao & P.Y.Hsu, "Evaluating IPv6 on a large-scale network," Computer Communications, Vol. 29, No. 16, 2006.

[5] X. Zhou & P. Mieghem, "Hopcount and E2E Delay: IPv6 versus IPv4," Passive and Active Network Measurements, 2005.

[6] Y. Wang, S. Ye & X. Li, "Understanding Current IPv6 Performance: A measurement study," IEEE Symposium on Computers and Communications, 2005.

[7] K. Cho, M. Luckie & B. Huffaker, "Identifying IPv6 Network Problems in the Dual-Stack World," ACM SIGCOMM Workshop on Network Troubleshooting, 2004.

[8] labs.ripe.net/Members/mirjam/content-ipv6-measurement-compilation.

[9] www.getipv6.info/index.php/IPv6_Penetration_Survey_Results

7 Appendix A. Distribution of Packet Loss

The following plots show the complementary distribution function of packet loss (a.k.a. the complementary cumulative distribution function, i.e. 1 minus the cumulative distribution function, i.e. the probability the random variable is *greater than* a given value). The complementary distribution function is often used when the interest is in the tail behavior. There is one plot for each of the six geographic areas. Note that range on the y-axis varies.

To interpret these plots, consider the first one, for North America, and the red line representing v6 packet loss to tunneled v6 interfaces. The point (5, 0.2) on the plot means that 20% of the measurements had packet loss of 5% or more.

These plots emphasize the observations made in Section 2. The higher v6 packet loss to tunneled interfaces is clearly evident.

Figure 7. Complementary Distribution Functions of Packet Loss

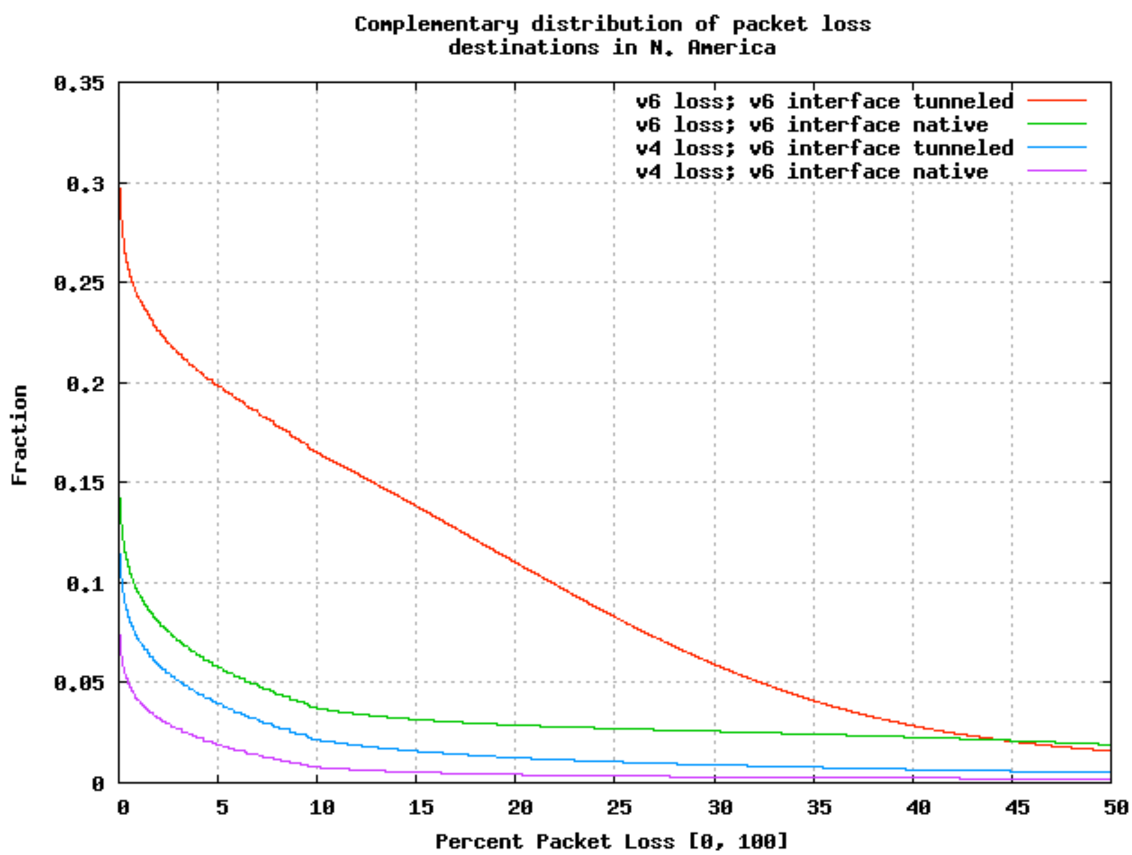

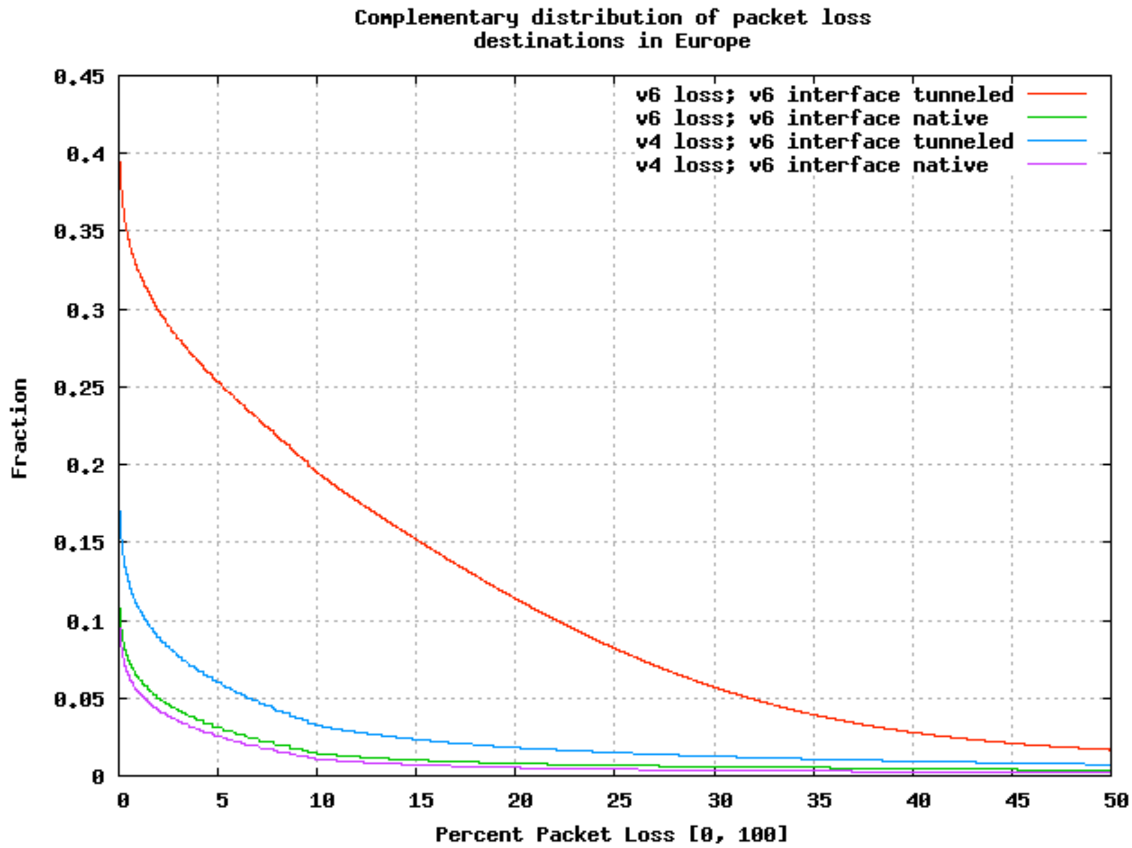
Complementary distribution of packet loss
destinations in Europe
v6 loss; v6 interface tunneled
v6 loss; v6 interface native
v4 loss; v6 interface tunneled
v4 loss; v6 interface native
Fraction
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
0
5
10
15
20
25
30
35
40
45
50
Percent Packet Loss [0, 100]

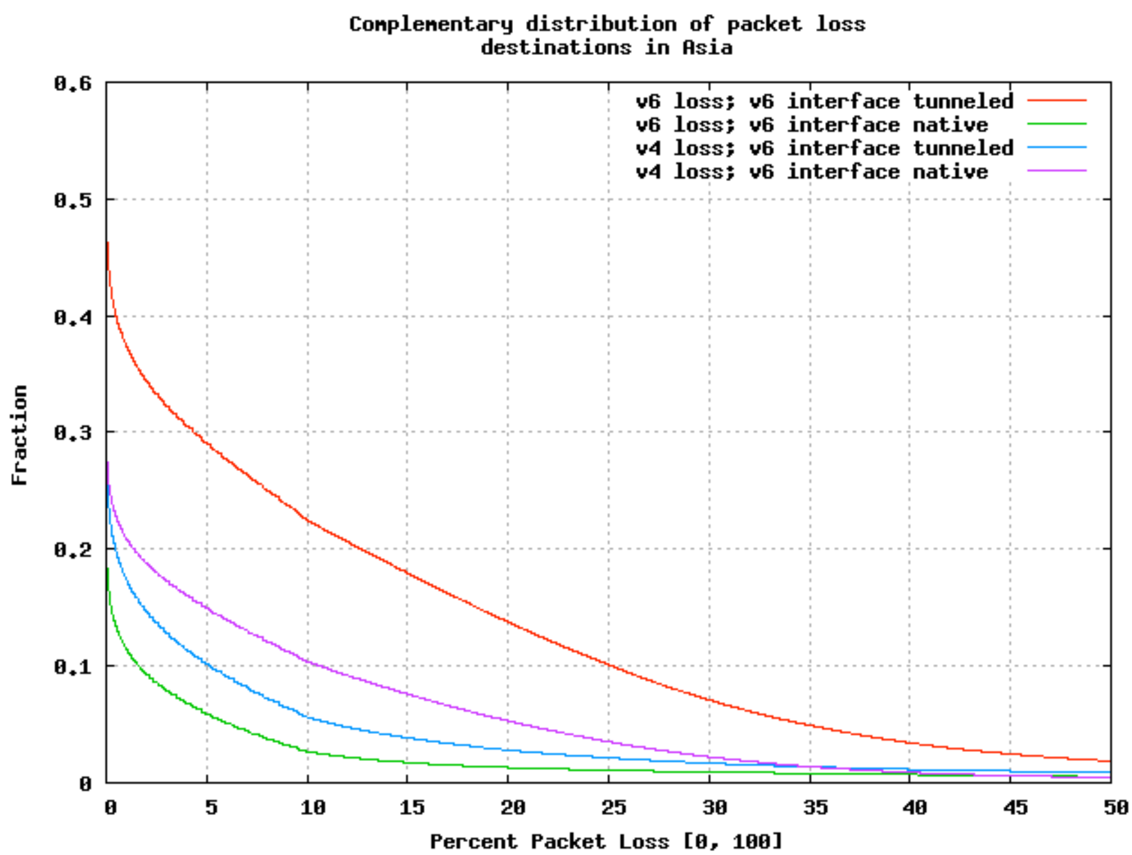
Complementary distribution of packet loss
destinations in Asia
v6 loss; v6 interface tunneled
v6 loss; v6 interface native
v4 loss; v6 interface tunneled
v4 loss; v6 interface native
Fraction
0.6
0.5
0.4
0.3
0.2
0.1
0
0
5
10
15
20
25
30
35
40
45
50
Percent Packet Loss [0, 100]

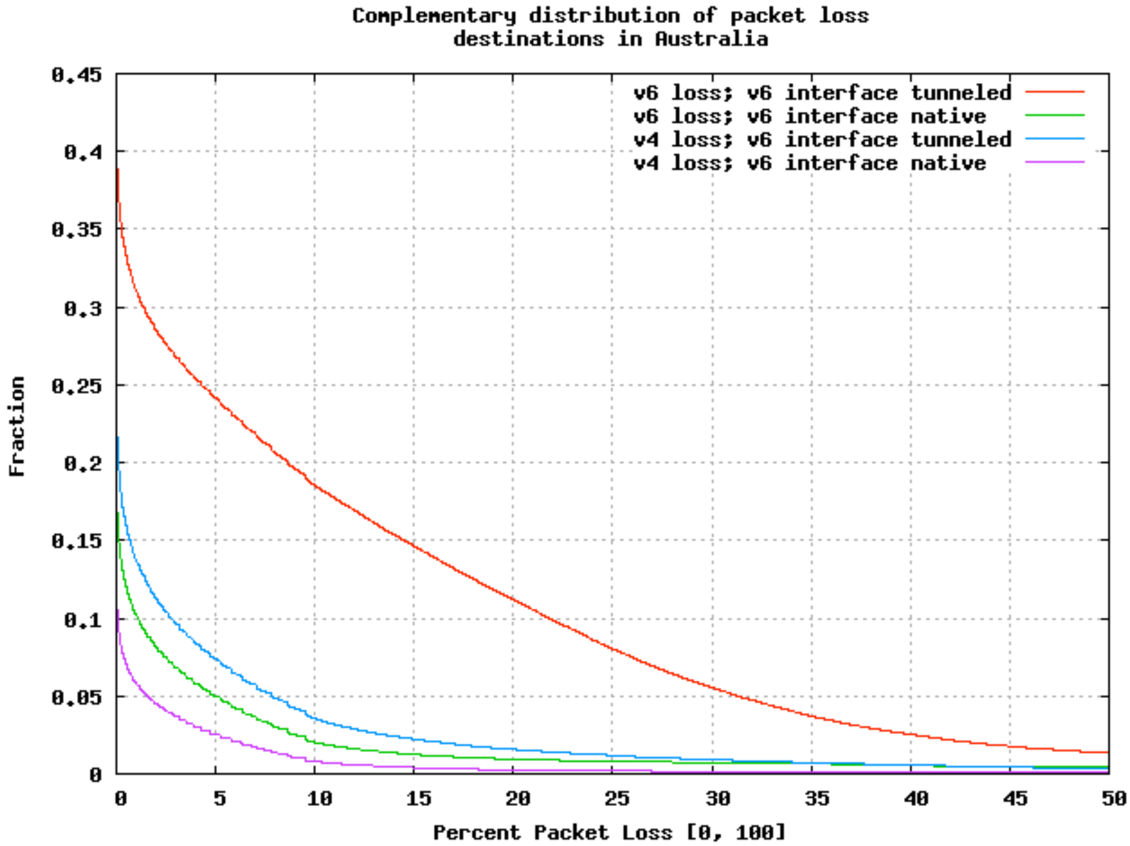
Complementary distribution of packet loss
destinations in Australia
v6 loss; v6 interface tunneled
v6 loss; v6 interface native
v4 loss; v6 interface tunneled
v4 loss; v6 interface native
Fraction
Percent Packet Loss [0, 100]
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
0.4
0.45
0
5
10
15
20
25
30
35
40
45
50

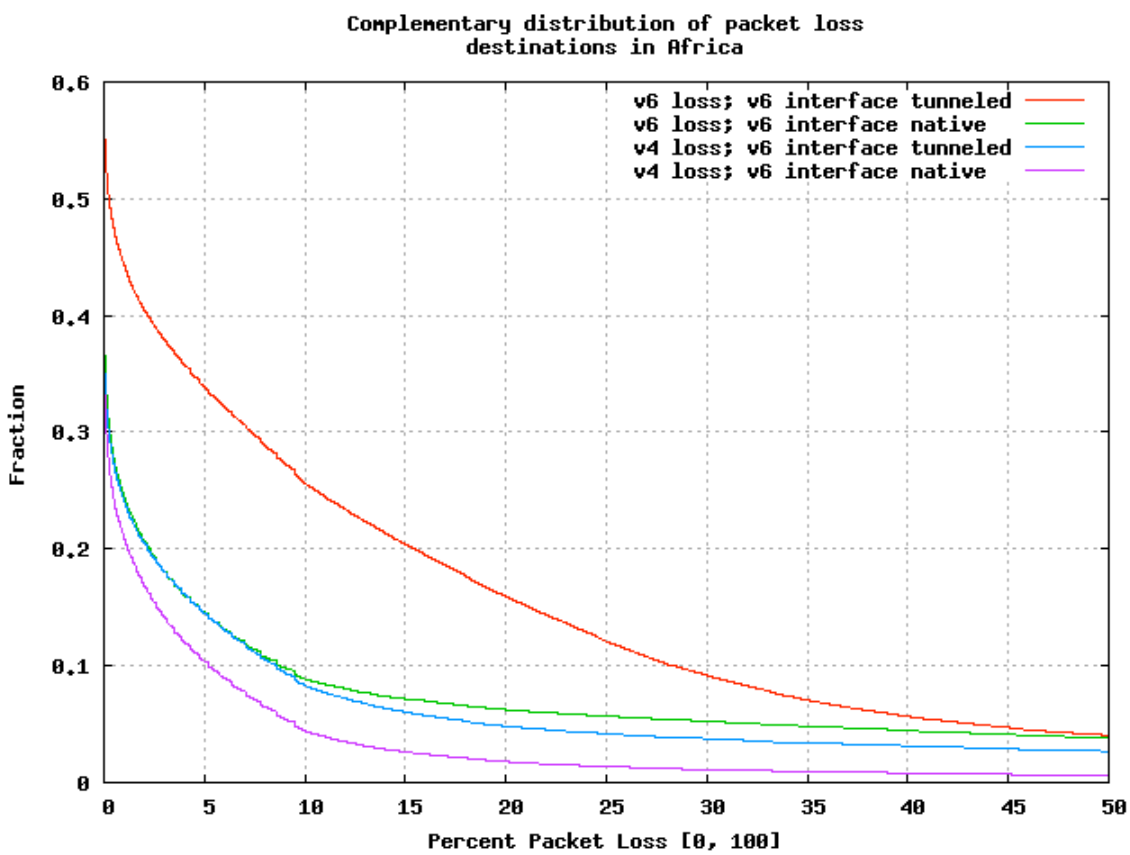
Complementary distribution of packet loss
destinations in Africa
v6 loss; v6 interface tunneled
v6 loss; v6 interface native
v4 loss; v6 interface tunneled
v4 loss; v6 interface native
Fraction
Percent Packet Loss [0, 100]
0
0.1
0.2
0.3
0.4
0.5
0.6
0
5
10
15
20
25
30
35
40
45
50

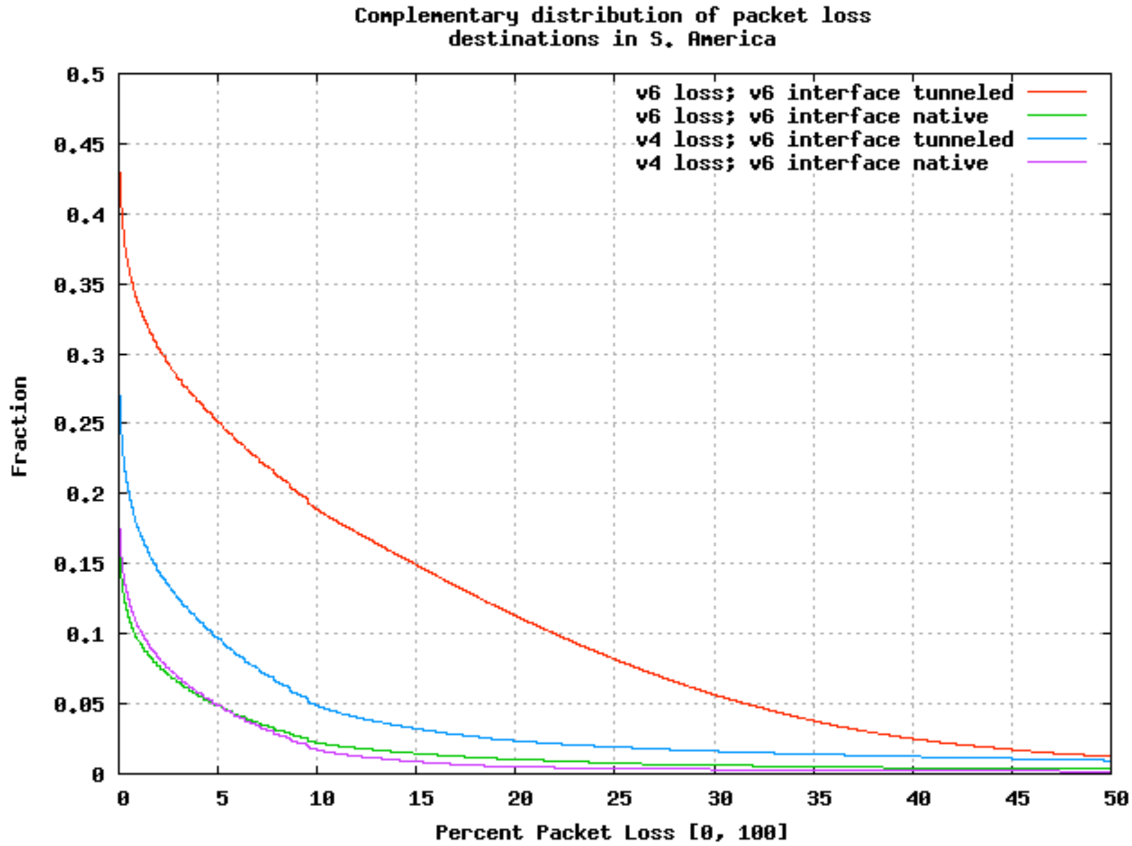
Complementary distribution of packet loss
destinations in S. America
v6 loss; v6 interface tunneled
v6 loss; v6 interface native
v4 loss; v6 interface tunneled
v4 loss; v6 interface native
Fraction
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
0
5
10
15
20
25
30
35
40
45
50
Percent Packet Loss [0, 100]

8 Appendix B. Additional Latency Distributions

The following plots are analogous to those in Section 3 except in addition to displaying the case of all the destinations in a given geo-region, also shown is the partition of destinations into native and tunneled. The range on the axes is held constant across the plots, and is chosen to highlight the portion where there is the most action. (The distributions do increase to 1, if the full range on the x-axis were shown.)

Note that the distribution given native (blue line) lies above that given tunneled (red line). This implies that the amount that transport is better over v4 (in the sense of lower latency) is greater for destinations with v6 tunneled interfaces than for native. For example, consider the 0.5 value on the y-axis, the median. For North America, 50% of the time the v4 latency is at least 14 ms better than v6 for native destinations (blue line), and is significantly more, at least 40 ms better for tunneled destinations (red line). For South America, these times are 7 ms and 43 ms.

For another viewpoint on the same concept, consider South America (and looking at 20 ms on the x-axis, and taking the complement of the y-value), for destinations where v6 is tunneled (red line), the v4 latency is better by at least 20 ms 68% of the time, whereas for destinations where v6 is native, the v4 latency is better by at least 20 ms only 31% of the time. Looking at -20 ms on the x-axis, v6 is better by at least 20 ms 14% of the time for native interfaces and is better by at least 20 ms not as frequently, 5% of the time, for tunneled interfaces.

Figure 8. Distribution of difference in latencies, partitioned by v6 interface

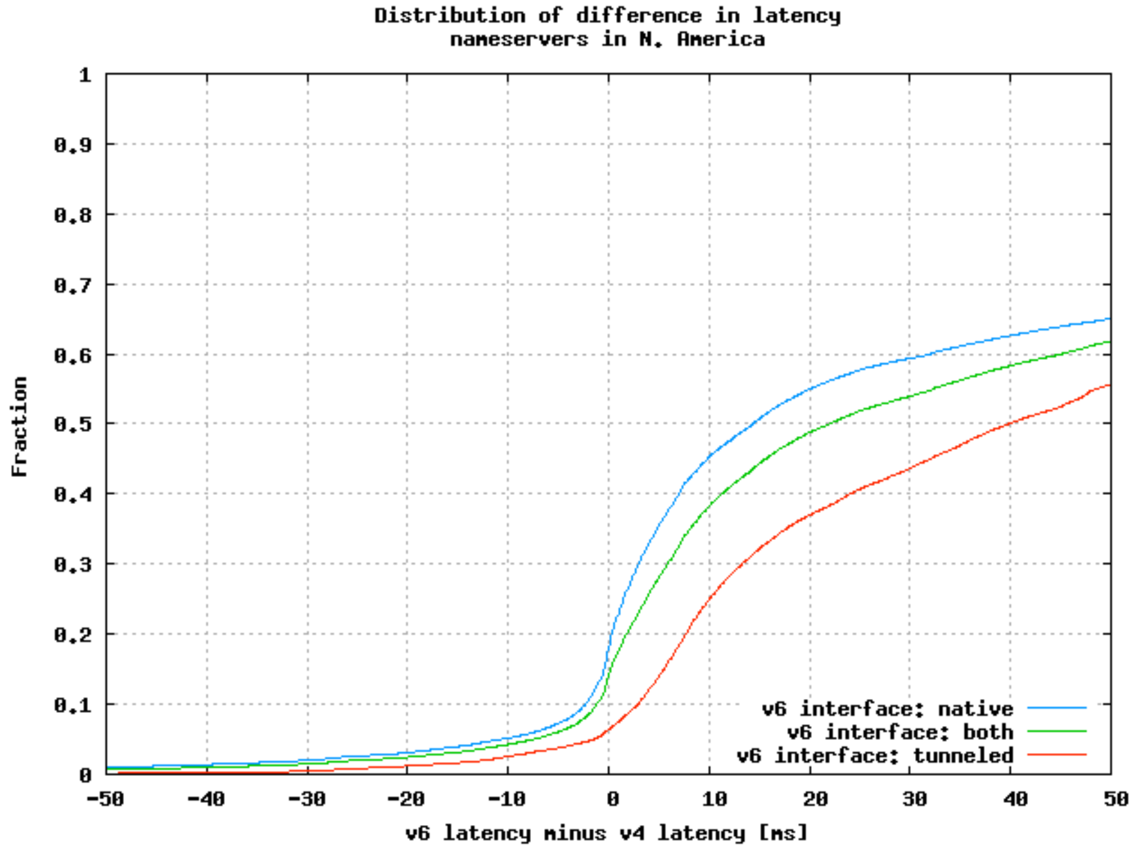
Distribution of difference in latency
nameservers in N. America
Fraction
v6 latency minus v4 latency [ms]
v6 interface: native
v6 interface: both
v6 interface: tunneled
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

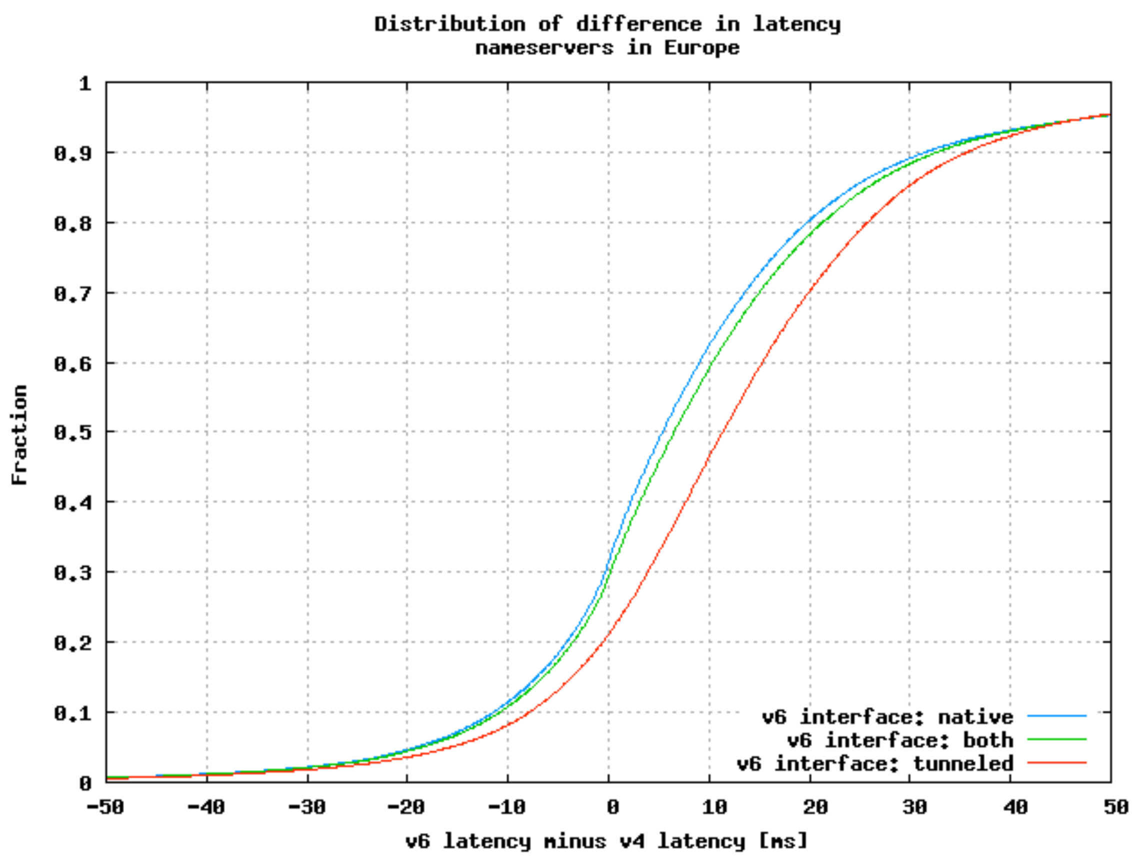
Distribution of difference in latency
nameservers in Europe
Fraction
v6 latency minus v4 latency [ms]
v6 interface: native
v6 interface: both
v6 interface: tunneled
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

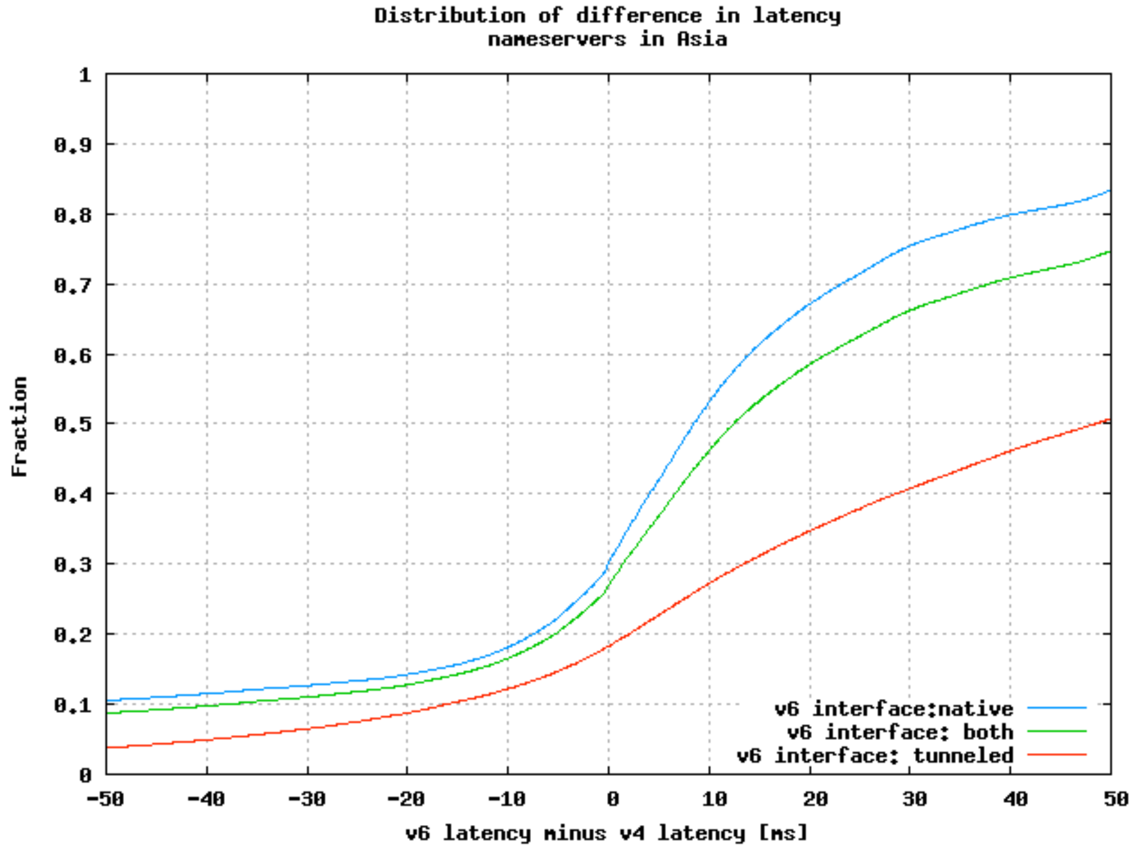
Distribution of difference in latency
nameservers in Asia
Fraction
v6 latency minus v4 latency [ms]
v6 interface:native
v6 interface: both
v6 interface: tunneled
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

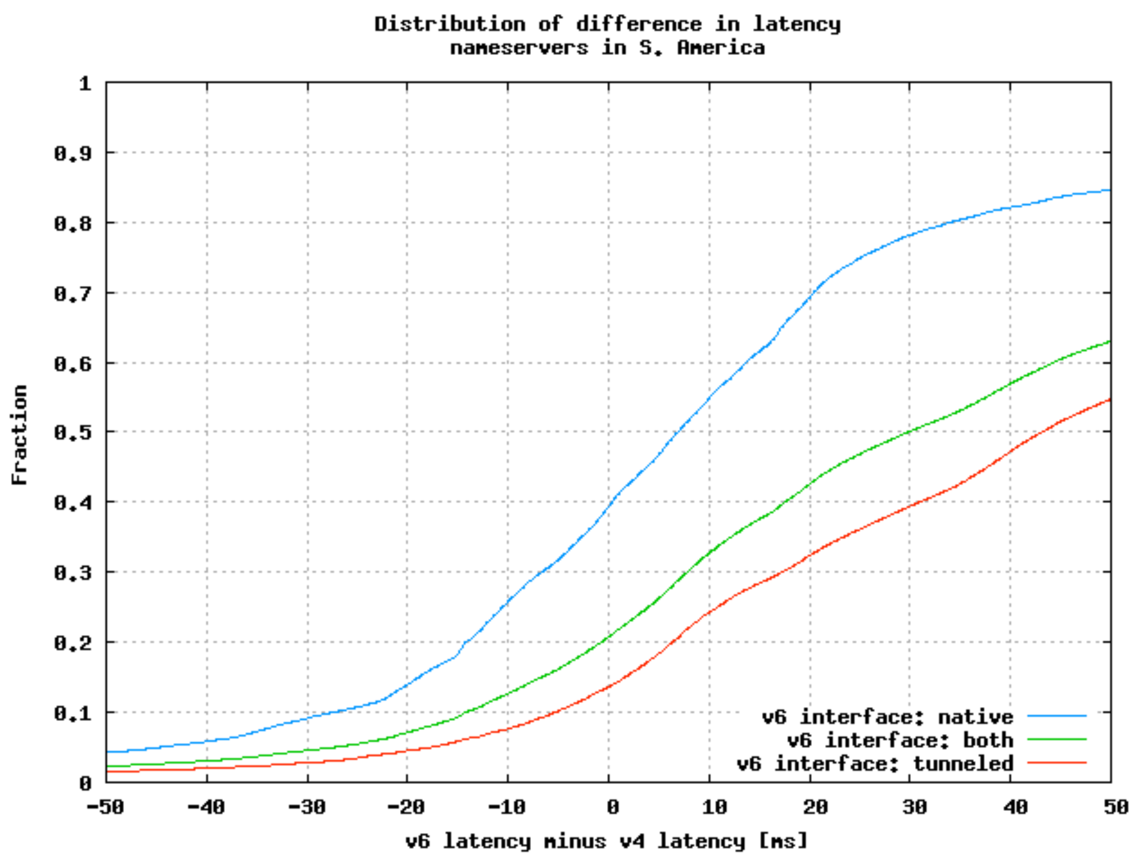
Distribution of difference in latency
nameservers in S. America
Fraction
v6 latency minus v4 latency [ms]
v6 interface: native
v6 interface: both
v6 interface: tunneled
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-50
-40
-30
-20
-10
0
10
20
30
40
50

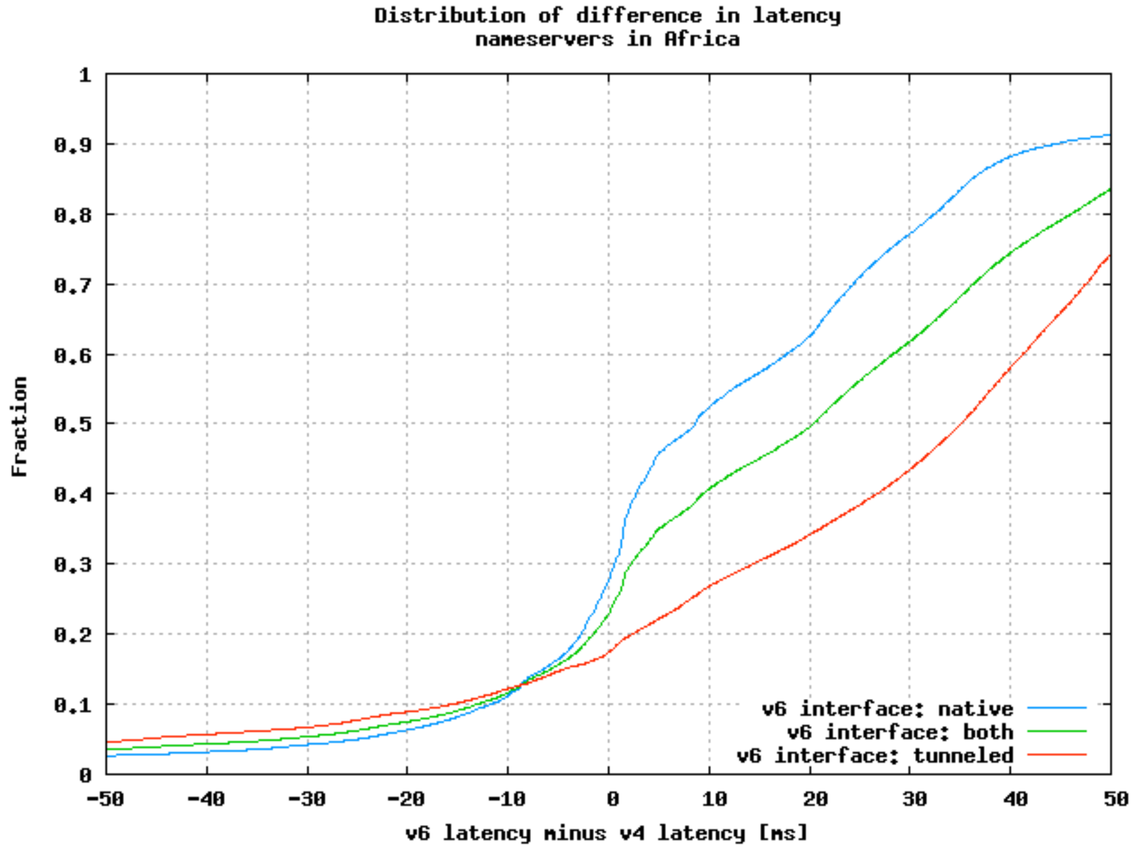
Distribution of difference in latency
nameservers in Africa
Fraction
v6 latency minus v4 latency [ms]
v6 interface: native
v6 interface: both
v6 interface: tunneled

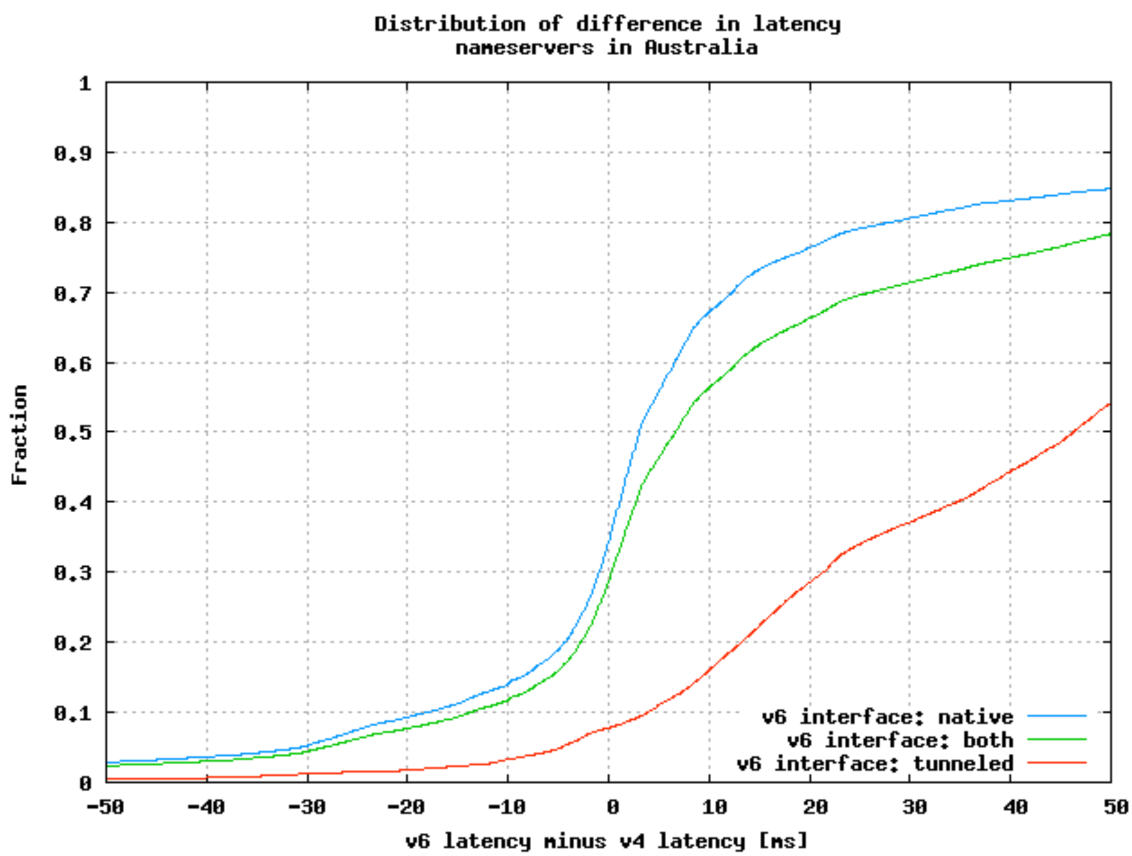
Distribution of difference in latency
nameservers in Australia
Fraction
v6 latency minus v4 latency [ms]
v6 interface: native
v6 interface: both
v6 interface: tunneled