

July 2010

Geoff Huston

Background Radiation in IPv6

To what extent is the IPv6 Internet polluted by "background radiation"?

In earlier work we set up a number of "black hole" experiments in the Internet, where traffic can enter the experimental setup, but the setup generates no packets in response. All received packets are recorded. So far we've used this setup to test a number of empty address blocks that have been allocated to APNIC in recent months, including 1.0.0.0/8, 14.0.0.0 and 223.0.0.0/8. Reports on these dark traffic experiments can be found at:

<http://www.potaroo.net/studies/1slash8/1slash8.html>

<http://www.potaroo.net/studies/14-223-slash8/14-223-slash8.html>

It's clear that these days the IPv4 Internet is now heavily polluted with various scanners and probes that attempt to detect the presence of vulnerable systems. This traffic is "dark" traffic in that it exists irrespective of whether it solicits a response from a remote system or not. The average level of dark traffic in the IPv4 Internet is an average of around 20kbps per /16, or the equivalent of a single incoming packet per address every 50 minutes. Individually, that does not appear to be so bad, but if one were to look at the IPv4 network as a whole, we have now managed to get to the point where there is in excess of 2Gbps of various forms of probing and scanning traffic across the entire IPv4 network. Background levels of traffic associated with scanning, backscatter, misconfiguration and leakage from private use contexts contribute to this traffic volume. If IPv4 is so polluted, then what about IPv6?

A paper published in 2006 reported on a 15 month experiment using a single /48. That experiment received 12 packets in the period, all ICMP6 [Matt Ford et al, 2006; "Initial Results from an IPv6 Darknet", In Proceedings of International Conference on Internet Surveillance and Protection (ICISP'06)], corresponding to a dark traffic rate of less than 1 packet per month.

Is result this "reasonable"? If we use a larger address block four years later, would we expect to see a different picture of the dark traffic in the IPv6 Internet?

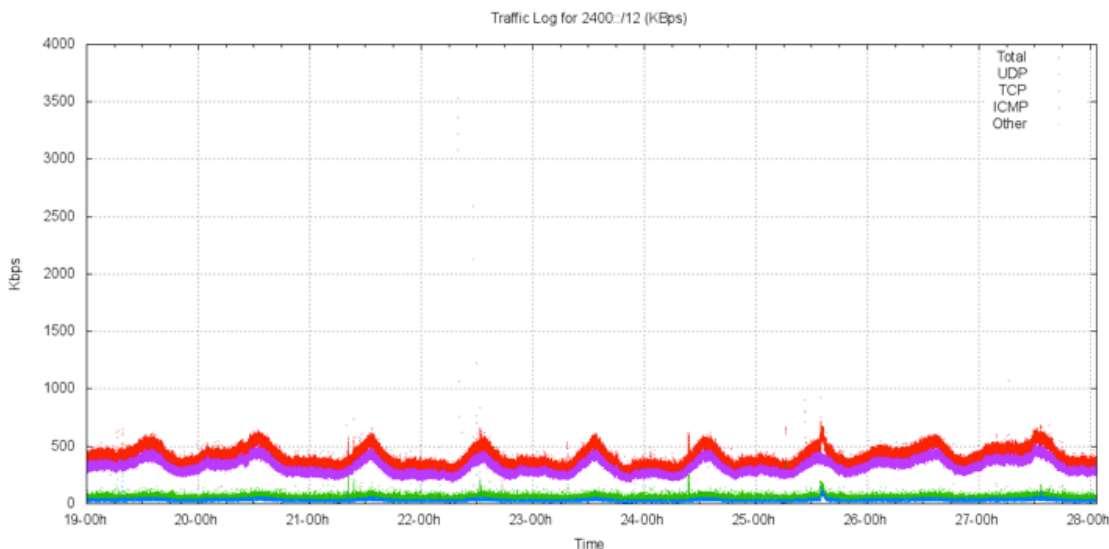
It seems to be a very low packet rate in comparison to IPv4, particularly when you consider that a /48 has 2^{80} , or some 1.2 billion billion IPv6 addresses in a single /48 prefix. However, we need to factor into this comparison the observation that the one critical difference between IPv4 and IPv6 is more addresses in IPv6, and this has been used to IPv6's advantage in the addressing architecture. IPv6 uses the lower 64 bits as an interface identifier within the context of a single subnet. In IPv6 each subnet there are 2^{64} possible addresses (that's 18 quintillion if you use the American system of the names of big numbers, or 18 trillion if you use the British system). Even if you scanned at a rate of 1 million addresses per second, a /64 subnet will take 584 billion years to scan, or a little under 50 times the age of the universe!

Much of the virus traffic in IPv4 is the result of the virus performing a scan of the address space, probing for new victims. However this rather simple blind probe technique of virus propagation is infeasible in IPv6. So when we look at "dark" address space, we don't expect to see probe traffic performing a scan of the IPv6 address space.

If this is not probe traffic performing a scan of the address space, then what form of traffic is out there in the dark nooks and crannies of IPv6 address space, if anything?

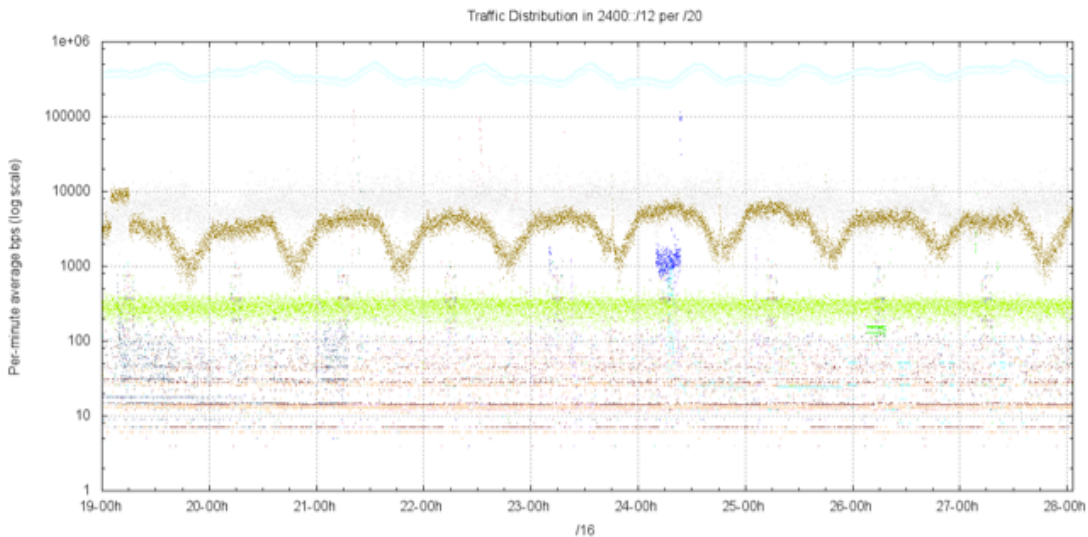
Examining Dark Traffic in 2400::/12

To answer this question we performed a darknet traffic experiment using 2400::/12 from the 19th to the 28th June 2010. Unlike the earlier IPv4 experiments, this is not a vacant address block as it was allocated to APNIC in October 2006. In the intervening period APNIC has performed 709 address allocations, spanning the equivalent of 16,629 /32's, or 1.59% of the /12 address block. Address space for private networks is handled differently in IPv6, and there is no direct equivalent of the private use RFC1918 space that is a part of the IPv4 address plan. In IPv6 what we see are both allocations that are being used in the context of the public Internet, and allocations that are used in private address contexts. At present, of the 709 IPv6 allocations that have been made from this address block, some 323 allocations are visible in the public IPv6 Internet. Of the /12 address block 0.91% of the address space is advertised in the inter-domain routing space, and 0.68% of the address space is allocation, but not advertised. 98.41% of the /12 is unallocated and unadvertised address space. It should be noted that this darknet experiment did not in any way alter traffic to advertised destination addresses within 2400::/12. Because 2400::/12 is a covering aggregate, the change in behaviour induced by this advertisement is to send only that traffic to the darknet sensor that would otherwise follow a default route and be discarded at the point where the traffic entered the default-free part of the network.



This figure shows a 9 day period from 0000 hours of the 19th June, UTC until 2400 hours of the 27th June, recording the incoming traffic rate in bits per second. The figure shows the total traffic rate (red), and the breakdown into ICMP6 (violet), UDP (green) and TCP (blue). Over the period the average traffic rate was 407Kbps, with a per second peak rate of 3.5Mbps. The traffic is predominately ICMP6, with an average of 611 packets per second (pps) of ICMP6 traffic. The UDP packet rate is 68 pps, and TCP 30 pps.

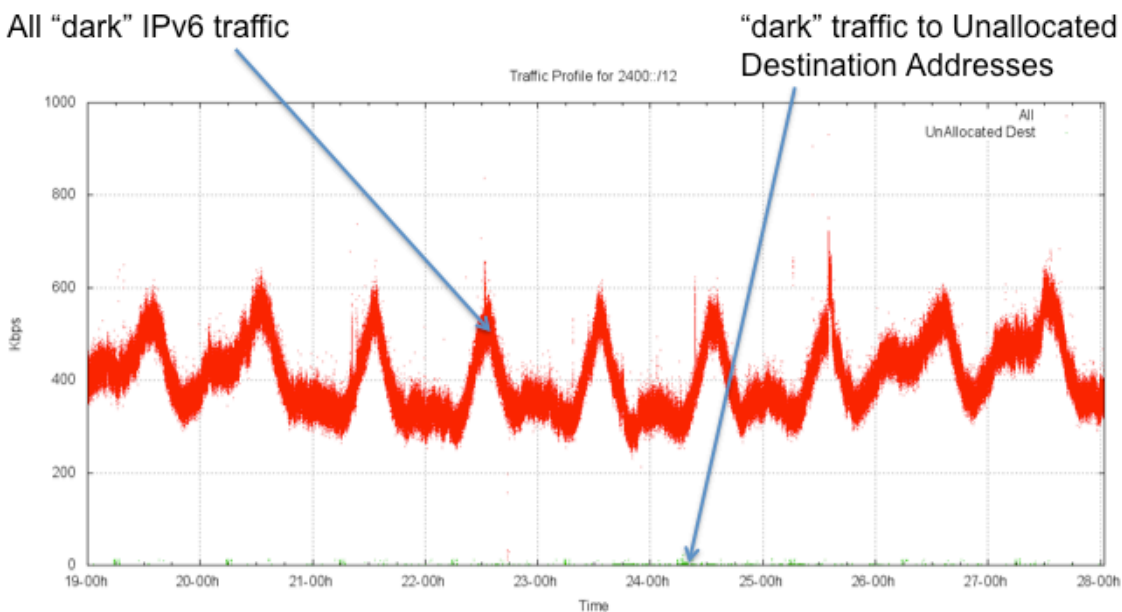
Its possible to look at the profile of this traffic by dividing the /20 into 256 component /20 blocks, and looking at the traffic profile per /20. This is shown in the following figure.



Interestingly, the traffic is not evenly spread across the /12. Only 21 /20s recorded an average traffic in excess of 1 bit per second. Evidently most of this /12 is pretty much devoid of traffic, and there are a small number of "hot spots" in the set of /20s. Of these /20s that have traffic, Table 1 lists the top 5, that have the highest average traffic levels.

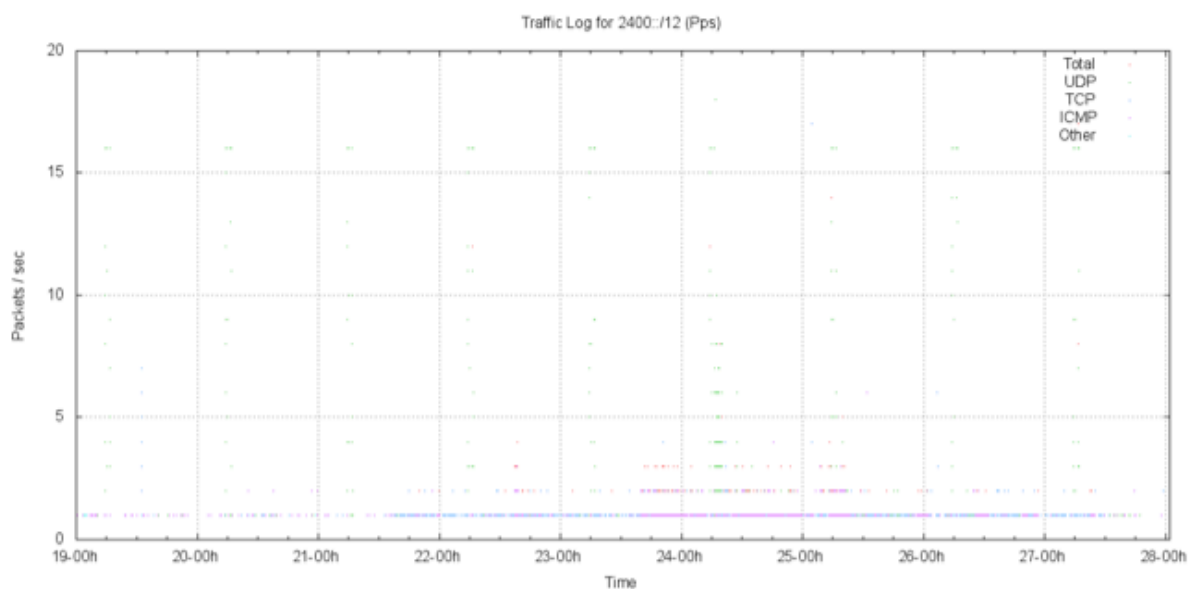
/20 Prefix	Average Traffic	Allocations
2408:0000:/20	197Kbps	Allocated: 2408::/22 - NTT East, JP
2401:d000::/20	7Kbps	8 x /32 allocations in this block
2403:8000::/20	4Kbps	4 x /32 allocations in this block
2404:0000::/20	1Kbps	29 allocations in this block
2405:b000::/20	0.3Kbps	4 x /32 allocations in this block

Closer inspection of this dark traffic reveals that the overall majority of this traffic is directed to addresses that are allocated, but not advertised. The following figure shows this traffic divided into traffic that appears to be leakage from private contexts and truly dark traffic.



A reasonable supposition is that most of the traffic in 2400::/12 is attributable to leakage from private use contexts. So now let's focus on what's left after we remove all this leakage from private use networks.

The level of the dark traffic is one packet every 37 seconds across the entire /12. Of the 21,166 packets collected, 7,881 (37%) were ICMP6, 7,660 were TCP (36%) and 5,609 (26%) were UDP. The traffic pattern appears to be sporadic, and very light. The following figure shows the packet rate of traffic directed to this unallocated dark IPv6 address space in 2400::/12 over the 9 day period.



There are few enough packets in this period to be able to perform a relatively intense scrutiny. What we are looking for here is a greater level of understanding of the likely cause of this traffic. Are we already seeing the signs of virus propagation and malign traffic in IPv6? Or is this traffic attributable to more benign causes? Let's look at the TCP, UDP and ICMP6 traffic in turn from this packet capture.

Dark TCP

The TCP traffic was further analysed to understand the likely cause of the traffic. TCP starts with a 3-way handshake of a SYN packet, an answering SYN+ACK packet and an ACK packet to complete the handshake. If the initiator of the connection is using an incorrect destination address for the intended service point it will send a SYN packet to the dark traffic collector. If the initiator has managed to misconfigure itself and is using a dark address at its local IPv6 interface then the SYN packet will be successfully sent, but the responding SYN+ACK packet will be sent to the dark traffic trap.

Of the 7,660 TCP packets some 1,126 packets were SYN packets being sent to unallocated (and presumably invalid) addresses. These packets may well be the result of configuration errors in the DNS zone files, where the AAAA resource record value is incorrectly specified with an invalid IPv6 address.

A much larger set of TCP packets, 6,392 in total, have the SYN and ACK flags set. One potential explanation of this behaviour is that a system has a badly configured local address, and while it can send IPv6 packets, the incorrect source address in these packets mean that the system never gets to see the response.

There were 141 other TCP packets, which are neither SYN or SYN+ACK packets, which are perhaps a little more curious because, if everything else is operating conventionally the dark packet collector show not see such fragments of a TCP conversation where one party to the conversation is using an invalid IPv6 address. For this to occur in the normal course of events the initial three way handshake would need to complete naturally, which contradicts the constraint that in order for these packets to be captured by this dark traffic capture experiment one party to this TCP session has managed to configure itself with an invalid IPv6 address.

Lets look at these anomalous TCP PACKETS in a little more detail.

The first of these TCP "oddities" is a TCP DNS response.

```
2001:468:1802:102::805b:fe01.53 > 2401:1a19::123:108:224:6.49121,  
Length: 1319 ACK: 1672186592 WIN 49980  
Query: A? finlin.wharton.upenn.edu.  
Response: finlin.wharton.upenn.edu. A 128.91.91.59
```

This looks like a perfectly normal DNS response using TCP. The odd aspect of this packet is that there is no evidence of the initial TCP handshake, and given that the IPv6 address of 2401:1a19::123:108:224:6 falls in the dark address space, there is no reason to believe that the TCP handshake could've ever taken place in any case. However, 2001:468:1802:102::805b:fe01 is the IPv6 address of a DNS server, and this server will respond to the above query as shown in this packet.

The second TCP "oddity" is this sequence of 9 TCP packets:

```
13:12:56.528487  
2001:250:7801:a400::x:y.33729 > 2402:e968:6000::d27e:4ed:fb5b.2273: .,  
3207301626:3207303066(1440) ack 3706857348 win 63916  
  
01:47:00.122909  
2001:250:7801:a400::x:y.57777 > 2402:2b75:2100:0:42:dc34:e8f3:52a4.3113  
272892761:272892761(0) ack 2064800132 win 64800  
  
01:50:47.197265  
2001:250:7801:a400::x:y.57777 > 2402:2f2a:179:341f:d6:dc34:e8f3:52a4.3113  
302360250:302360250(0) ack 2091174988 win 64800  
  
03:44:39.140290  
2001:250:7801:a400::x:y.57777 > 2402:a236:6000:0:4d8:dc34:e8f3:52a4.3113  
829577701:829577701(0) ack 2622550921 win 64800  
  
03:58:23.851708  
2001:250:7801:a400::x:y.57777 > 2402:9a23:100:2:d6:dc34:e8f3:52a4.3113  
829661294:829661294(0) ack 2702723699 win 64800  
  
05:02:52.568996  
2001:250:7801:a400::x:y.57777 > 2402:1123:1ba:ec05:ef:f2c6:ce35:c40f.1158  
1365702964:1365702964(0) ack 3293642040 win 64800  
  
05:50:43.706430  
2001:250:7801:a400::x:y.57777 > 2402:76d9:16b:7320:d8:f2c6:ce35:c40f.1158  
1409613792:1409613792(0) ack 3600529388 win 64800  
  
07:20:15.728521  
2001:250:7801:a400::x:y.57777 > 2402:6219:4100:0:2b0:dc34:e8f3:52a4.311  
830692465:830692465(0) ack 3672203022 win 64800  
  
08:37:57.505208  
2001:250:7801:a400::x:y.57777 > 2402:b54e:1cc:e14:52:dc34:e8f3:52a4.3113  
831214068:831214068(0) ack 4169603866 win 64800
```

The packets share a common source address and source port. Some of the packets share a common low order 64 bits in the destination address, and only a small set of destination port addresses are used. The packets are ACK packets without any data payload, except the first, which contains a data payload. Interestingly the ack sequence numbers are monotonically increasing. It is possible that these packets are part of some form of TCP probe, and because

there are no TCP handshake packets preceding these packets its reasonable to presume that these are not backscatter packets, but packets that are probing towards the nominated destination addresses.

The packet capture also contains a second set of possible of TCP probes, this time as a cluster of SYN probes:

```
12:44:54.038234
  2001::4137:9e76:28ae:355f:x:y.80 > 240a:f000:1405:6001:1cbc:f191:1384:7cde.1597
    Flags [S.], seq 3889176058, ack 2381452531, win 8192, length 0
12:44:54.038358
  2001::4137:9e76:28ae:355f:x:y.80 > 240b:f000:1685:6001:1cbc:f191:1384:7cde.1597
    Flags [S.], seq 3889176058, ack 2381452531, win 8192, length 0
12:44:54.038613
  2001::4137:9e76:28ae:355f:x:y.80 > 240c:f000:1905:6001:1cbc:f191:1384:7cde.1597
    Flags [S.], seq 3889176058, ack 2381452531, win 8192, length 0
12:44:54.914216
  2001::4137:9e76:28ae:355f:x:y.80 > 240c:f000:1905:6001:1cbc:f191:1384:7cde.1597
    Flags [.], seq 1, ack 220, win 17080, length 0
12:44:54.914341
  2001::4137:9e76:28ae:355f:x:y.80 > 240a:f000:1405:6001:1cbc:f191:1384:7cde.1597
    Flags [.], seq 1, ack 220, win 17080, length 0
12:44:54.914466
  2001::4137:9e76:28ae:355f:x:y.80 > 240b:f000:1685:6001:1cbc:f191:1384:7cde.1597
    Flags [.], seq 1, ack 220, win 17080, length 0
12:49:52.061661
  2001::4137:9e76:28ae:355f:x:y.80 > 240b:f000:1685:af01:b469:173f:8bc8:3411.3991:
    Flags [.], seq 536162733, ack 2327619384, win 16621, length 0
12:49:52.061785
  2001::4137:9e76:28ae:355f:x:y.80 > 240c:f000:1905:af01:b469:173f:8bc8:3411.3991:
    Flags [.], seq 536162733, ack 2327619384, win 16621, length 0
12:49:52.061915
  2001::4137:9e76:28ae:355f:x:y.80 > 240a:f000:1405:af01:b469:173f:8bc8:3411.3991:
    Flags [.], seq 536162733, ack 2327619384, win 16621, length 0
```

In this case the packets are sent in two bursts, and the packets vary in the destination address in only bit positions 32 – 47. Again, this traffic pattern could be consistent with some form of experimental TCP probing. In this case the source address is a Teredo source address.

However, despite these oddities, most of the TCP traffic is the probably result of a very small number of configuration errors.

One of these forms of configuration errors is in incorrectly entering the local IPv6 address into an end system.

```
10:56:20.719296
  2001:470:1f04:815::2.25 > 2402:5000::250:56ff:feb0:11aa.37839:
    SYN, 2261394238:2261394238(0) ack 2082559012 win 64768
    <mss 1420,sackOK,timestamp 128287793 3737661225,nop,wscale 11>
```

What appears to happening here is that a mail server within Hurricane Electric is correctly responding to a client's attempt to connect to the SMTP port. The client has managed to mistype their IPv6 local address, and has attempted to connect to the mail server using the invalid source address of 2402:5000::250:56ff:feb0:11aa. The mail server's response to the Initial TCP syn packet is being captured by the dark net.

There are sequences of 8 such packets paced over 90 second periods with doubling periods between successive packets. This is a typical signature of a SYN handshake failure in a Dual Stack environment, where a client will attempt to perform a connection using IPv6 sending a sequence of SYN packets with a doubling of the inter-packet interval between each SYN attempt. The application will timeout after a number of unsuccessful attempts and fall back to IPv4. Conventionally, a client would cache this failure and not reattempt this IPv6 connection for some time. In this case the client has been unusually persistent and managed to send a total of 6,284 connection attempts over 9 days, corresponding to some 780 attempts to connect to the mail

server over IPv6. This single misconfiguration instance generated the majority of the "dark" TCP traffic in this period.

Dark UDP

What about UDP? Once more there is evidence of misconfiguration in addresses, and again this is likely to be the result of a DNS configuration error. 1,883 packets were detected performing a traceroute to the dark address 2405:a800::1.

There were 2,892 DNS queries in this 7 day period, originating from just 4 source addresses, and 30 DNS responses which appears to be backscatter from DNS servers attempting to respond to clients who have misconfigured the address of their system. These look a lot like configuration errors in dual stack environments that go largely unnoticed because the client has the ability to silently mask over the problem by falling back to IPv4.

There were some detected oddities in the DNS packet capture, possibly as an outcome of experiments in source address spoofing

```
2001:468:1802:102::805b:fe01 53 > 2407:ed24::113:23:133:101 40288
2607:f470:1003::3:3 53 > 2407:bde7::113:23:133:101 16554
2001:468:1802:102::805b:fe01 53 > 2407:df5c::113:23:133:101 47237
2607:f470:1003::3:3 53 > 2407:cd0d::113:23:133:101 44097
2001:468:1802:102::805b:fe01 53 > 2407:29ca::113:23:133:101 47145
2001:468:1802:102::805b:fe01 53 > 2407:f0e1::113:23:133:101 29201
2001:468:1802:102::805b:fe01 53 > 2407:72c4::113:23:133:101 53336
2607:f470:1003::3:3 53 > 2407:6a85::113:23:133:101 40444
```

In this packet set 2 DNS servers are responding to queries from dark addresses in 2400::/12, probably as an instance of backscatter from a queries with spoofed source addresses. What is not clear here is why the destination addresses appear to vary only in bit position 16 – 31.

Dark ICMP

There are 7,882 ICMP6 packets in this 9 day period. Of these ICMP6 packets 7,804 are ICMP echo request packets, as would be generated by a standard *ping* application. There are just 125 unique destinations of these *ping* packets.

63 ICMP6 packets are *destination unreachable* messages, presumably a backscatter effect from misconfigured systems attempting to make connections with systems that are denying the connection attempt:

```
08:44:12.679309
2001:470:1f04:815::2 > 2402:5000::250:56ff:feb0:11aa: ICMP6, destination unreachable,
unreachable port, 2001:470:1f04:815::2 tcp port 25
```

The remaining 15 packets are ICMP6 *parameter problem* messages, indicating that there are a small number of systems out there with both a misconfigured local IPv6 address and also an implementation of IPv6 that is generating malformed IPv6 packets.

Dark Traffic in IPv6

What can we say about the dark traffic in IPv6? Is it malign or mostly the outcome of generally benign mistakes in configuring equipment?

It is certainly true to state that what we see as dark traffic in IPv4 has no counterpart in IPv6. The nature of the massively sparse population of the low-end 64 bit interface identifier addresses in IPv6 makes address scanning pretty much impractical. There may be some small number of guess probes being directed to ::1 and ::2 addresses, but on the whole there is no evidence of any systematic scan of address space happening across all IPv6 addresses. So far there is no direct evidence of virus scanners probing into the dark address blocks in IPv6.

What we do see is some evidence of configuration errors in IPv6. The overwhelming volume of the traffic seen in this exercise is not truly dark packets, but leakage from private use contexts. Due to a failure in the local configuration, a sizeable amount of supposedly private network traffic is incorrectly sent out into the public IPv6 Internet. To a much lesser extent there is a small volume of dark traffic that is the result of transcription errors in editing DNS zone files with IPv6 addresses and local system configuration in manually setting up local IPv6 interface addresses.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

www.potaroo.net