Investigating the Impact of Service Provider NAT on Residential Broadband Users

Shane Alcock
University of Waikato
Hamilton, New Zealand
salcock@cs.waikato.ac.nz

Richard Nelson
University of Waikato
Hamilton, New Zealand
richardn@cs.waikato.ac.nz

David Miles
Alcatel-Lucent
David.Miles@alcatel-lucent.com.au

ABSTRACT

Service Provider NAT (SPNAT) allows ISPs to share a single public IPv4 address amongst multiple subscribers by placing NAT devices within the carrier network, potentially extending the lifetime of the IPv4 address space. However, there has little research aimed at evaluating the potential impact of SPNAT on residential broadband users. This paper presents a study of the behaviour of DSL subscribers for an ISP that has not implemented SPNAT, examining the number of outgoing connections created and utilised by each user. The results show that some subscribers utilise significant quantities of short-lived UDP flows that cannot be dealt with efficiently using the current NAT best practice. We also propose a new technique for expiring UDP flows more efficiently and demonstrate that it is effective in reducing session table usage.

Categories and Subject Descriptors

C.2.3 [**Computer Communications Networks**]: Network Operations

General Terms

Experimentation, Measurement, Performance

Keywords

SPNAT, Passive measurement, Residential broadband, Flows

1. INTRODUCTION

IPv4 address exhaustion is arguably the most pressing problem facing the Internet today. Demand for IP addresses is growing and most estimates now predict that the remaining IPv4 address space will be exhausted before the end of 2012 [1] [2]. As the exhaustion date approaches, it is becoming increasingly difficult for Internet Service Providers (ISPs) to acquire new address allocations, placing pressure on the addresses they already possess. The long-term solution to this problem for ISPs is to migrate to IPv6 [3]. Although IPv6 is widely supported in both hardware and software, the transition process is still occurring very slowly. It is now likely that many ISPs will exhaust their allocated address space before completing the migration to IPv6.

One technique designed to create more time to transition to IPv6 is Service Provider Network Address Translation (SPNAT), also known as Carrier-Grade NAT. SPNAT involves the widespread use of routers located inside the carrier network that are capable of Network Address Translation (NAT) [4] to share individual IPv4 addresses amongst multiple subscribers, allowing the ISP to serve the same customer base using a much smaller quantity of addresses.

One of the limitations to this approach is that NAT devices must maintain state in a session table for each active connection to ensure a consistent translation for packets belonging to the same flow [5]. The maximum rate at which entries can be added to and removed from the session table is defined by the processing power of the NAT device. Also, each transport protocol is limited to 65535 concurrent sessions per public IP address, due to the limited number of available ports. If either of these limits is reached, a subscriber behind SPNAT will experience lengthy delays when attempting to create a new outgoing connection. Therefore, it is important for an ISP utilising SPNAT to provision NAT devices such that the peak session demand seldom exceeds the capabilities of the NAT device.

As SPNAT is a relatively new technology, few ISPs have deployed it within their access networks thus far. One reason for this is that both hardware vendors and ISPs are uncertain how to suitably dimension SPNAT devices and networks such that any adverse effects caused by the limitations described above are minimised. This uncertainty is due to a lack of recent measurement data that examines the connection creation behaviour of residential users. It has been suggested that previous experience with stateful firewalls could be useful but our enquiries among national ISPs have revealed that stateful devices are rare, especially in network cores or on networks serving consumer broadband customers.

To address this, we have used packet traces captured at the border of a New Zealand ISP to examine broadband user behaviour within the context of SPNAT. Outgoing connections were examined to determine the likely requirements for NAT devices to enable current usage patterns to be supported. To do so, we have measured both the session creation and expiry rates, as well as the session table usage for each subscriber. The results can be used by hardware vendors to dimension SPNAT devices in terms of both processing power and memory. ISPs can also use the measurements to estimate an appropriate ratio of subscribers to public IPv4 addresses.

The principal contribution of this work is that it is an in-depth study based entirely on measurements of real-world broadband users, i.e. the consumers who will be most directly affected by the adoption of SPNAT. It is notoriously difficult for network researchers to gain access to such data from ISPs, primarily due to concerns about privacy and

commercial sensitivity, hence the lack of similar studies in the past. Furthermore, the packet traces that were used for this research have been captured within the past 18 months. This ensures that the results presented in this paper are not only relevant and timely but they can also be used to evaluate whether conventional wisdom regarding user behaviour is still applicable.

This research is also relevant to other NAT-based technologies where many subscribers are translated onto a smaller number of addresses, where NAT64 is an obvious example [6]. The suggested behaviour of the NAT64 device, especially the session expiry rules, are very similar to those that were used to produce the measurements presented in this paper. Therefore, many of the results and conclusions reported here can also be applied to NAT64 implementations.

Our results showed that large quantities of short-lived UDP connections would be retained in the simulated session table long after the connection itself had concluded. We noticed that the efficiency of an SPNAT implementation could be significantly improved by expiring the sessions for those connections quickly. Based on these results, we developed and evaluated a new expiry method that shortens the timeout for UDP sessions that have not sent more than one outgoing packet.

The paper is organised as follows. The data set used in the course of this study is introduced in Section 2. Our approach for analysing the data is described in Section 3 and the results are presented in Section 4. In Section 5 we propose and evaluate a new technique for expiring UDP sessions. Related work is discussed in Section 6 and Section 7 concludes.

2. DATA SET

The trace set used for this analysis was captured from a New Zealand ISP between January 7 and January 11, 2009. The ISP had not implemented SPNAT at the time of the capture. The traces were captured using a passive monitor located between the border routers and the rest of the carrier network, recording both inbound and outbound traffic for all customers. Upon capture, each packet was tagged with the appropriate direction using MAC address information provided by the network operator. In addition, each captured packet was truncated four bytes after the transport header, retaining a small amount of application payload without significantly compromising the privacy of the network users. This payload can be useful for identifying application protocols using a limited form of deep packet inspection. Aside from the truncation, no alterations were made to the contents of the packet, e.g. IP addresses and port numbers were not anonymised.

The trace set is entirely contiguous and no sampling was performed during the capture process, ensuring that every packet observed by the monitor was captured, truncated and recorded to disk. However, to exclude corporate and wireless clients from the analysis, only packets with IP addresses that were within the DSL subscriber subnets were examined during the course of the study. BPF filters based on known IP ranges for DSL subscribers were used to facilitate this.

Due to space considerations, we have only analysed connections that were initiated by the DSL subscriber, i.e. outgoing connections. Unlike incoming connections, where the application behaviour may change to successfully connect through to the customer on the other side of a NAT, these sessions are unaffected by the NAT process and do not require NAT traversal techniques to pass through the NAT device. As a result, the measured traffic should be representative of what would be observed if the subscribers were all placed behind an SPNAT device.

3. METHODOLOGY

There are two limiting factors that define the performance of a NAT device: the rate at which entries are created and removed from the session table and the total size of the session table. A NAT device creates a new session for each connection and both the maximum session creation rate and the size of the session table are finite. The principal concern for any SPNAT implementation is provisioning NAT devices efficiently, i.e. how many subscribers can share the same NAT device without experiencing a significant degradation in performance.

To investigate this, we developed a software analysis tool that simulates the session table for a NAT device. Packets are read from the trace set and assigned to connections using the 5-tuple (source and destination IP addresses and ports and the transport protocol). If an entry for that connection does not already exist in the session table, a new entry is added. Otherwise, the statistics (e.g. byte and packet counts) and expiry time for the connection are updated but the session table itself remains unchanged. After processing each packet, entries in the session table that have expired are removed and details about the expired sessions are written to an output file for subsequent analysis. Throughout this process, a series of counters track the number of new, expired, currently active and peak active outgoing sessions for each observed subscriber IP address.

A connection is defined as outgoing if the first observed packet is sent from within the ISP network, i.e. the packet was tagged as outbound during the capture. For TCP connections, the initial packet is also required to be a SYN packet, removing any malformed or invalid TCP connections from the analysis. This also has the effect of ignoring any in-progress connections at the beginning of the trace set, but we believe that our trace set is of sufficient length that the overall effect of this would be minimal. A TCP connection is defined as established once a SYN ACK packet has been observed in response to the initial SYN. By contrast, the first observed UDP packet is always treated as the start of a valid UDP connection, because the protocol does not provide any obvious indication of the start of a connection.

Unestablished TCP connections are expired and removed from the session table after four minutes, or twice the maximum segment lifetime, of inactivity [7]. Established TCP connections are expired after two hours and four minutes of idle time, as defined in [8]. Any TCP connection for which a FIN is observed in both directions are expired immediately, as are connections that observe a RST packet. UDP connections are expired after two minutes of inactivity, which is the minimum described in [9]. Finally, an ICMP error message, such as the Destination Unreachable message, results in immediate expiration of the original session that triggered the ICMP response.

For each session, four bytes of application payload from the first payload bearing packet in each direction is retained, along with the original application payload length. When the session is expired, this data is used to attempt to determine the application protocol by comparing the observed

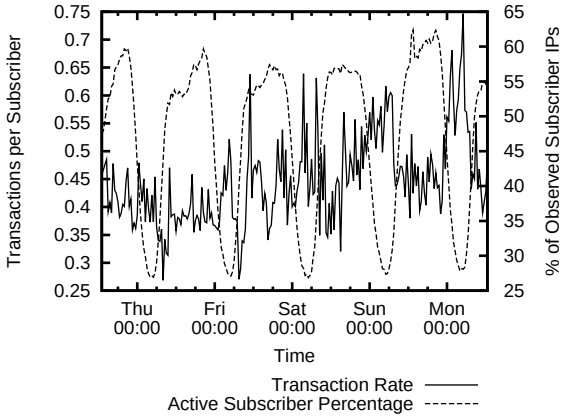

Figure 1: Transactions per active subscriber IP during the busiest second in each 30 minute period. The dashed line shows the percentage of subscribers that were deemed active during that period.

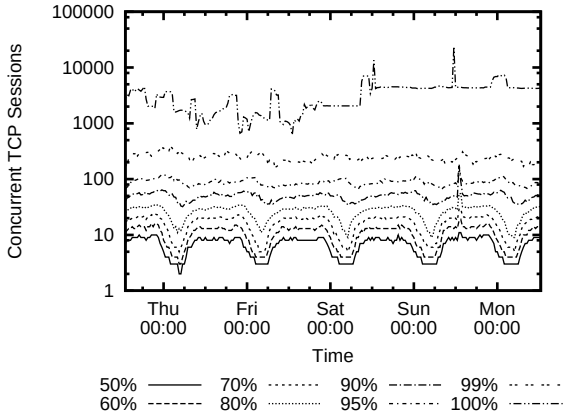

Figure 2: Percentile analysis of peak concurrent TCP sessions for each active subscriber IP.

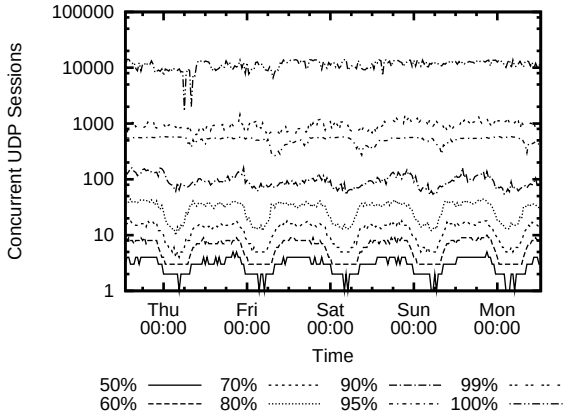

Figure 3: Percentile analysis of peak concurrent UDP sessions for each active subscriber IP.

payload against a set of common patterns, similar to the method presented in [10]. This technique enabled us to frequently identify protocols that could not be reliably classified using port numbers. For example, BitTorrent connections can be identified by matching either of the retained payloads with the character 0x13 followed by the string "Bit". [11] describes the method in further detail and lists all of the identifable protocols.

As we did not have access to any information regarding the allocation of IP addresses to subscribers during the measurement period, such as RADIUS logs, we have assumed that each unique IP address observed within the address range of the ISP represented an individual subscriber. This is not an ideal approach, as the subscriber addresses will have been allocated dynamically from a pool and can change whenever the subscriber reboots their DSL modem. This is highlighted in [12] where it was found that half of the IP addresses in the DSL pool was assigned to more than one subscriber over a 24 hour period.

To alleviate this, the results presented in this section are averaged across the number of *active subscribers* during each thirty minute period. We define an active subscriber as one that either creates a new connection during the period or had at least one outstanding unexpired entry in the session table when the period began. As it is unlikely that a significant portion of DSL subscribers will switch addresses over the course of thirty minutes, the results should not be notably affected by changes in the subscriber IP address.

4. RESULTS

Figure 1 shows the number of transactions per active subscriber during the busiest second from each half-hour period in the trace set, where a transaction is defined as the addition or removal of an entry from the session table. This metric can be used to determine the quantity of subscribers that can be assigned to the same SPNAT device without the total transaction rate ever exceeding the processing capabilities of the device. Selecting the busiest second, i.e. the second where the most transactions occur, means that the measurements indicate occasions where the simulated session table is under the heaviest workload. The values presented in Figure 1 were produced by dividing the number of transactions during the busiest second in a thirty minute period by the number of subscribers that were active during that entire period.

The results suggest that the transaction rate required to cope with the peak session demand observed in the trace set is 0.71 transactions per second for each subscriber. This peak occurs during the early hours of Monday morning at a time when less than 30% of observed subscriber IP addresses were active and therefore may have been skewed upwards by the behaviour of a relative minority of users. If periods where less than half of the observed addresses were active are ignored, a peak demand of 0.64 transactions per subscriber occurred at 11 a.m. on Friday.

The other principal concern for an SPNAT implementation is the amount of session table space that each subscriber will require. To investigate this, the peak concurrent session counts, i.e. the maximum number of session table entries for each period, for each subscriber were broken down into percentiles. The results are presented in Figures 2 and 3, which show the percentile analysis for TCP and UDP sessions respectively. The percentiles were calculated separately for

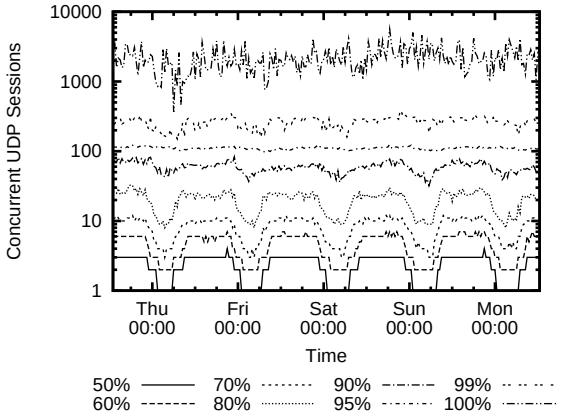

Figure 4: Percentile analysis of peak concurrent UDP sessions using the proposed 10 second expiry rule for UDP sessions.

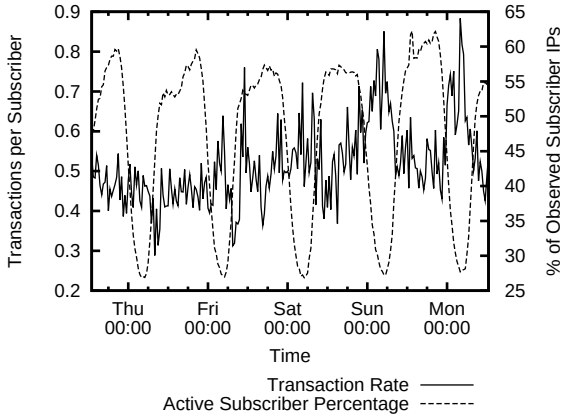

Figure 5: Transactions per active subscriber IP during the busiest second using the proposed 10 second expiry rule for UDP sessions.

each half-hour period to create the time series shown in the graphs. The median (the 50th percentile) is not particularly high for either protocol, usually ranging between one and ten concurrent sessions. The distribution is very long-tailed, suggesting that provisioning an SPNAT implementation based solely on the requirements of an "average" subscriber may not be the best approach.

We believe the most interesting result can be observed at the 100th percentile, i.e. the largest of the peak values observed for each half-hour period. In the case of TCP, the 100th percentile is regularly between four and five thousand concurrent sessions. The peak for UDP is even higher, frequently surpassing 10000 concurrent sessions. As each NAT device can only maintain a maximum of 65535 UDP port mappings per IP address, this is a significant quantity of session table space to allocate to a single subscriber.

To determine the underlying cause of the high session counts, we analysed in further detail 316 separate occasions where the peak UDP session count for a subscriber exceeded 10000 sessions in further detail. This was done by filtering the output from the simulated session table to show only sessions belonging to the subscriber that were present in the session table when the peak was achieved. The 316 cases were shared across 18 unique subscriber IP addresses, with the most prominent IP address accounting for 107 instances.

Our analysis showed that the high numbers of concurrent UDP sessions were primarily caused by a disparity between the connection duration and the time spent in the session table. On all except one of the examined occasions, at least half of the UDP connections had a duration (where duration is measured as the time between the first and last observed packet) of less than a second but remained in the session table for over two minutes. 99% of these sessions transmitted only a single outgoing UDP packet.

5. UDP SESSIONS

The results presented above suggest that the performance of an SPNAT implementation can be improved by detecting and expiring UDP sessions that last less than a second quickly rather than waiting for the full two minutes. This approach is mentioned in [9] but it is suggested only for protocols that use well-known port numbers. However, there was no consistent port number being used by the sub-second duration UDP sessions that were examined, meaning that approach will fail for almost all of the problematic sessions. Payload analysis using the techniques described in [11] revealed that BitTorrent DHT (Distributed Hash Table) requests accounted for over 95% of the sessions in all except one of the cases that we encountered, with the Steam server browsing protocol being responsible for the one exception.

We therefore propose that a shorter session expiry timeout be used for UDP sessions where only a single outgoing packet has been observed (we will henceforth refer to these sessions as short-lived UDP sessions). Once a second outgoing packet is observed, the current standard timeout of two minutes can be applied instead. This will remove short-lived UDP sessions from the session table quickly without affecting long-running UDP exchanges. This can be done without examining application payload and the NAT device will only need to maintain a small amount of additional state for UDP sessions. This could be encoded using a single bit acting as a boolean flag. The first outgoing packet would create the session with the flag set to be false and the next observed outgoing packet would then set it to true. As long as the flag is set to false, the shorter timeout would be used.

To evaluate the viability of the proposed solution, we adjusted the session expiry rules in the simulated session table to expire short-lived UDP sessions after only 10 seconds of idle time. UDP connections that had sent two or more outgoing packets continued to be expired after 120 seconds of inactivity, as before.

The percentile analysis for peak UDP session counts under the new expiry rules is shown in Figure 4. Although the 100th percentile consistently remains above 1000 concurrent sessions, the values for the higher percentiles have decreased significantly with the value of 100th percentile falling by 79% on average. By contrast, the peak session counts for subscribers closer to the median have changed little. This suggests that expiring short-lived UDP flows faster will significantly decrease session consumption by the heaviest users without having any noticable effect on the majority of subscribers.

However, this improvement in session table utilisation also

Table 1: Performance of different idle expiry thresholds for short-lived UDP sessions

Threshold	Total Sessions Expired	Peak Transaction Rate	Peak Transaction Rate (>50% active)	Average 50th Percentile	Average 100th Percentile	Repeated UDP 5-tuples
120 sec	104.5 million	0.7125	0.6365	2.85	10101.4	27.8 million
60 sec	105.9 million	0.7916	0.6274	2.56	6721.0	29.1 million
30 sec	107.5 million	0.8236	0.7130	2.42	4693.8	30.7 million
10 sec	111.5 million	0.8994	0.7591	2.33	2126.4	34.5 million
1 sec	122.3 million	0.9066	0.7791	2.12	287.6	44.2 million

had an impact on the transaction rate during the busiest seconds, as shown in Figure 5. The total number of sessions created across the entire trace set grew by 6.7% and a peak transaction rate was observed approaching 0.9 transactions per second per subscriber, an increase of 26% over the previous expiry scheme. A similar increase was observed for the Friday 11 a.m. period mentioned earlier, which grew from 0.64 to 0.76 transactions per second per subscriber.

Table 1 shows the results from extending the analysis to include several different possible expiry thresholds ranging from the current suggested practice of 120 seconds through to a single second timeout. The first statistic shown is the total number of sessions expired over the course of the analysis which gives an indication of the workload for the simulated SPNAT device across the entire measurement period. This is followed by two values describing the peak transaction rate. The first is the peak transaction rate per active subscriber across the entire traceset, regardless of the number of subscribers that were active during that time period. This is equivalent to the largest of the peaks observed in Figures 1 and 5. The second is the peak transaction rate per active subscriber that was observed during periods when at least half of the subscriber IPs were active.

The percentile values in Table 1 show the mean of the 50th and 100th percentiles for peak concurrent UDP sessions from each half-hour period. The principal aim of the experiment is to reduce the impact of the subscribers that utilise UDP heavily, thus a significant decrease in the 100th percentile is desirable. Ideally, such a decrease would also have a minimal effect on the 50th percentile representing the average subscriber.

Finally, the table shows the number of 5-tuples that were observed in the session table on multiple occasions where the protocol was UDP. A sharp increase in this value indicates that UDP sessions had been expired prematurely by our simulation. These sessions then must be re-entered into the session table at a later point, increasing the workload of a NAT device. In addition, if a session is expired early, any incoming packets for that session can no longer be forwarded to the subscriber that they were intended for, which is obviously a very undesirable outcome.

The results in Table 1 suggest that notable improvements in session table usage can be obtained even with an expiry threshold of 60 seconds. In that case, session table utilisation at the 100th percentile fell by approximately a third, while the peak transaction rate rose by 11%. Interestingly, the peak transaction rate when the majority of subscribers were active decreased slightly. We suspect this was due to transactions that had previously happened during the busiest second were now occurring at other times. With a 30 second expiry threshold, the improvements were even greater: the maximum session table usage decreased by over 50% with a corresponding increase in the two peak transaction rates of 16% and 12% respectively.

Using a single second expiry timeout for short-lived UDP sessions resulted in a 97% decrease in the average 100th percentile. The increase in the peak transaction rates were relatively small compared with the 10 second threshold. However, using a one second timeout is not a realistic option for SPNAT implementations because latency exceeding one second is not uncommon, particularly on home DSL connections. Should the initial response packet for an outgoing UDP session be delayed by more than a second, the session would no longer exist in the session table by the time the response arrives and the packet would not be forwarded to the subscriber. This would have a very detrimental effect on the UDP application. The large increase in repeated UDP 5-tuples compared to the ten second threshold suggests that this is likely to be a widespread problem.

In summary, decreasing the expiry timeout for short-lived UDP sessions to 60 seconds would result in a significant improvement in session table utilisation for an SPNAT device whilst having a very minimal impact on other SPNAT performance metrics. We therefore recommend that this change be strongly considered for adoption by SPNAT implementors. Decreasing the expiry threshold to 30 or 10 seconds instead will result in further improvements, but implementors would be advised to consider the corresponding increases in transaction rate and prematurely expired sessions carefully before doing so.

Finally, we note that the overall number of repeated UDP 5-tuples was surprisingly high, even when using the standard 2 minute expiry threshold. This is due to some BitTorrent clients reusing the same local port to send DHT requests to remote peers. Peers would be probed on multiple occasions several minutes apart, such that the gap was long enough to ensure the original session will always have timed out, resulting in a repeat 5-tuple.

6. RELATED WORK

There has been some previous discussion about the potential flaws in SPNAT, for example [13], but there has been little in-depth analysis using real-world measurements to assess the potential impact of SPNAT. Most attention in the area has been focused on standardising NAT device behaviour, and defining implementation details through RFCs [8] [9] [6]. This work is primarily aimed at ensuring that the practical requirements of SPNAT can be achieved and there has been little consideration towards the resulting impact on network users, aside from [14] which summarises the poten-

tial problems with shared address solutions to the problem of IPv4 exhaustion, such as SPNAT.

As part of a broader study of residential broadband traffic, [12] used passive data captured from a European ISP to assess the potential viability of SPNAT. The authors concluded that SPNAT could produce a signficant reduction in IP address utilisation. However, as the authors themselves admit, this analysis was not performed in any depth and the results presented appear to be based purely on the average of the concurrent sessions per subscriber. [12] does not mention the large quantities of concurrent UDP sessions that we encountered during our analysis, for instance, that could rapidly overwhelm a SPNAT device. Given that the primary source of the UDP sessions that we saw were BitTorrent DHT requests, we would have expected this behaviour to have also been prominent in their data. Another weakness of this analysis was that a fixed timeout was applied to each outgoing session rather than an expiry policy matching the IETF standards such as the one we used.

Another relevant work that we encountered was [15] which studied residential broadband utilisation in Japan by using aggregated traffic logs and sampled NetFlow records to analyse raw traffic volumes. The authors noted the existence of "heavy-hitters" who account for a disproportionate amount of traffic and the popularity of peer-to-peer file-sharing amongst heavy users. However, the results presented in the study are of little direct use in assessing SPNAT where connection counts, rather than traffic volumes, are the most critical metric.

7. CONCLUSION

The implementation of Service Provider NAT to extend the lifetime of IPv4 address allocations appears imminent. However, there has been little research based on real-world measurements to investigate how existing Internet users will be affected by SPNAT. In this paper, we have examined the outgoing connections for a sample of DSL subscribers to measure the impact that current broadband user behaviour would have under a simulated SPNAT implementation and to identify subscribers and applications that may be adversely affected by SPNAT.

In doing so, we found that the peak session demand in our data set was approximately 0.7 transactions per second per subscriber, assuming the current expiry standards are implemented. We also found that the suggested expiry policy for UDP sessions in a NAT device retains short-lived UDP connections in the session table for an excessive period of time, wasting session table space. The principal cause of the short-lived UDP sessions appeared to be BitTorrent DHT requests. In response to this, we proposed a new technique to expire short-lived UDP connections quickly by reducing the expiry threshold for sessions where only one outgoing packet had been observed.

We repeated our original analysis using several different expiry threshold values and demonstrated that peak session table utilisation could be decreased significantly using the proposed technique. Based on our results, we recommend that an expiry threshold for short-lived UDP sessions of no more than 60 seconds be adopted by SPNAT implementors, as this value provided a notable reduction in session table usage with a minimal impact on the other metrics, including transaction rate. Reducing the expiry threshold to a lower value, such as 30 or 10 seconds, was also shown to have merit, but we noted that the resulting increases in transaction rate and premature session expiration should be considered carefully before doing so. We believe this approach could also be adopted by other transitional technologies based on NAT, such as NAT64.

In the future, we aim to repeat our analysis on trace sets captured from other ISPs to confirm that the conclusions reached in this paper are applicable to consumer Internet traffic in general. We are also planning on conducting a study that will investigate the extent of residential users accepting inbound connections from other hosts, as these users will also be adversely affected by the loss of end-to-end connectivity that will result from the use of SPNAT.

8. REFERENCES

[1] G. Huston, "IPv4 Address Report," http://www.potaroo.net/tools/ipv4/index.html.

[2] I. van Beijnum, "IPv4 Address Consumption," *The Internet Protocol Journal*, vol. 10, no. 3, 2007.

[3] S. Deering and R. Hinden, "RFC 2460 - Internet Protocol, Version 6 (IPv6) Specification," 1998.

[4] K. Egevang and P. Francis, "RFC 1631 - The IP Network Address Translator (NAT)," 1994.

[5] M. Ford, A. Durand, P. Roberts, and P. Levis, "Address Sharing - Coming to a Network near You," *IETF Journal*, vol. 5, no. 1, 2009.

[6] M. Bagnulo, P. Matthews, and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers (Internet Draft)," http://tools.ietf.org/html/draft-ietf-behave-v6v4-xlate-stateful-08.

[7] R. Braden, "RFC 1122 - Requirements for Internet Hosts - Communication Layers," 1989.

[8] S. Guha, K. Biswas, B. Ford, S. Sivakumar, and P. Srisuresh, "RFC 5382 - NAT Behavioral Requirements for TCP," 2008.

[9] F. Audet and C. Jennings, "RFC 4787 - Network Address Translation (NAT) Behavioral Requirements for Unicast UDP," 2007.

[10] T. Karagiannis, A. Broido, N. Brownlee, k. claffy, and M. Faloutsos, "Is P2P Dying or Just Hiding," in *IEEE Globecom 2004 - Global Internet and Next Generation Networks*, 2004.

[11] S. Alcock, "Application Protocol Identification," http://www.wand.net.nz/~salcock/proto/.

[12] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On Dominant Characteristics of Residential Broadband Internet Traffic," in *IMC '09: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference*. New York, NY, USA: ACM, 2009, pp. 90–102.

[13] O. Maennel, R. Bush, L. Cittadini, and S. Bellovin, "A Better Approach than Carrier-Grade-NAT," Technical Report. 2008.

[14] A. Durand, M. Ford, and P. Roberts, "Issues with ISP Responses to IPv4 Address Exhaustion (IETF Draft)," 2009, http://tools.ietf.org/id/draft-ford-shared-addressing-issues-00.txt.

[15] K. Cho, K. Fukuda, H. Esaki, and A. Kato, "The Impact and Implications of the Growth of Residential User-to-User Traffic," in *Proceedings of SIGCOMM'06, Sept. 11-15 2006, Pisa, Italy*, 2006.