Inferring The Extent of Network Address Port Translation at Public/Private Internet Boundaries

Grenville J. Armitage
Centre for Advanced Internet Architectures. Technical Report 020712A
Swinburne University of Technology
Melbourne, Australia
garmitage@swin.edu.au

Abstract-**This technical report describes a relatively simple method for inferring the percentage of public/private internet boundaries that utilize network address port translation (NAPT, often colloquially referred to as NAT). Estimates were obtained from the IP address/port pairs seen in the server logs of three well-used, online game servers between May 2001 and June 2002. The report concludes that NAPT may be in use at approximately 17 to 25% of public/private internet access boundaries in the online gaming community. Estimates of NAT deployment can help provide context for discussions about the need for IPv6 and other techqniues for scaling the Internet.**

Keywords-Internet, Network Address Translation, NAT, NAPT, IP Address, UDP Port, Quake III, fingerprint, inference, measurement

I. Introduction

It is becoming increasingly common for access routers and gateways to offer Network Address Port Translation (NAPT) - a superset of Network Address Translation (NAT) [1] - in both the home/consumer and business markets. Typically NAPT is used to allow a small number of publicly routable IP addresses to be shared by a far larger number of hosts inside a private intranet. In addition, using NAPT at the public/private network boundary allows private networks to change their public service provider without having to re-assign all their internal IP addresses. NAPT functionality is usually embedded into multi-purpose edge routers, gateways, or firewalls.

Unfortunately, NAPT is not entirely transparent [2]. The address and port translation process at the NAPT boundry (from now on the 'NAT box') imposes an asymmetry on the resulting IP communication model. For example, hosts inside the private network must usually initiate communication with hosts on the public side of the NAPT boundary (so the NAT box can automatically assign the appropriate address/port mappings). Hosts on the private network cannot act as well-known public servers without prior manual configuration of the NAT box, and generally only one host on the private network can act as a public server for each type of service being offered. (For example, only one internal http server can be mapped to port 80 on the public side of the NAT box.)

NAPT also creates problems for protocols that assume both ends of a connection see the same IP address and TCP/UDP port numbers. For instance, any Quality of Service (QoS) schemes that rely on well known (or signalled) IP address/port pair information may need to make explicit allowances for NAT boxes being introduced into the communication path.

Finally, NAPT allows IPv4 to support growth of the Internet far beyond the point where we had previously expected to need IPv6 [3]. It is, therefore, of some interest to discover and quantify how pervasive NAPT has become across the Internet. Such knowledge can help provide more focussed context for discussions about next generation IP architectures, migration to IPv6, and end-to-end IP QoS schemes.

This paper describes a relatively simple approach to inferring the percentage of locations using NAPT from the client IP address and port information logged by three public Quake III Arena game servers over a the period from May 2001 to June 2002. We infer that roughly 17 to 25% of internet access is through a NAPT-enabled gateway, router, or firewall.

The rest of this paper briefly summarizes how Quake III helps detect NAPT, the details of our QuakeIII testbeds, and the conclusions we can draw from our results.

II. Data Aquisition

A. The NAPT Finger print

Developing an active protocol for detecting NAPT between two hosts is relatively easy, and was not the goal of this research. Rather, our interest lay in passively discerning the degree to which people use NAPT by monitoring only one end of an IP packet exchange. This requires first that people will voluntarily establish IP communication with our monitored host, and that we know what the packet exchanges would look like in the presence and absence of NAPT.

NAPT can be detected when one host sees packets arriving from another host and the IP address or TCP/UDP port numbers don't match what the destination host expected to see. Our solution was to run a number of online "Quake III Arena" game servers [4], and monitor the UDP ports from which clients connected. Quake III uses UDP/IP for all client/server exchanges, and clients use a default source UDP port of 27960. We assume NAPT exists between a client and our server when the server sees a client connecting from a source UDP port other than 27960.

B. *Quake III Arena Testbeds*

Three servers cover most of the period from May 2001 to June 2002 - two in Palo Alto, California (PA1 and PA2) and one at the University College London (UCL). PA1 ran from May 2001 to February 2002, UCL ran from June 2001 to September 2001, and PA2 ran from January 2002 to June 2002.

In each case the Quake III server was modified to log each client's IP address and UDP port number. Clients were ignored if they didn't at least stay joined for a minute or two and pick up at least one item per minute. It wasn't necessary to quantify each client's actual ability or willingness to play Quake III since the game itself was merely a method for attracting volunteers we could analyse for their NAPT fingerprint.

Defining a unique client posed a minor challenge, since the Quake III server does not export a unique client ID to the logfile. Thus we're left with each player's IP address, port number, and self-assigned 'player name' with which to differentiate clients. We do not want to overcount clients whose public IP address changes from day to day (for example, because their broadband ISP uses DHCP [5] to re-assign addresses every so often, or they're using a dial-up connection) since they really represent the same instance of a public/private internet access boundary.

Our solution was to perform a reverse-DNS lookup on every IP address, and use the ASCII playername and non-hostname part of the fully-qualified domain name as a *sufficiently unique* client identifier. For example, consider the hypothetical player 'InspectorKluSo' being seen over a period of time from three IP addresses that resolved to dsl1.bigprovider.com, dsl9.bigprovider.com, and dialup45.bigprovider.com respectively. Our algorithm counted that as the single unique client <*InspectorKluSo, bigprovider.com*>. Where no domain name could be resolved, we created a fake '.unresolved' domain - InspectorKluSo from address w.x.y.z would be client <*InspectorKluSo, w-x.unresolved*>. (In other words, clients with unresolvable IP addresses are treated as though their ISPs assign addresses from a class B sized address pool.) This approximation seemed to work well.

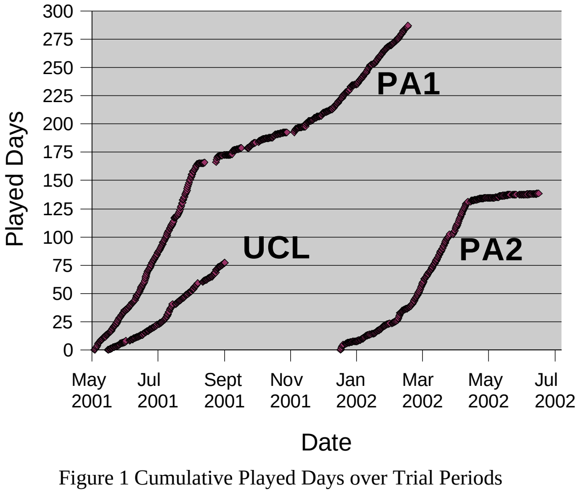

Figure 1 Cumulative Played Days over Trial Periods

III. Usage Patterns

Before looking at the NAPT estimates it is worth noting that the servers were reasonably well used over their lifetimes, although the usage did fluctuate during the year due to changes in game popularity, some periods of server downtime, and other unknown factors. Fortunately, NAPT estimation doesn't require that people enjoyed the server. It is sufficient that lots of people chose to visit the servers, even if only for a few minutes.

Fig. 1 illustrates the cumulative played time for each server (a count of how many player-days have been played at each point in the May 2001 to June 2002 period - two players playing half a day each counts as one player-day). The distinct slow-downs in both PA1 and PA2 occur when each server was upgraded to a higher patch level. (Initially version 1.17, PA1 went to 1.29h in late August 2001 and 1.30 in November 2001, while PA2 started at version 1.30 and went to 1.31 in early 2002). Each server lagged the latest patch level by a few months, making it popular with players who were late to upgrade. Then each upgrade would lose us our recurring players until they too upgraded.

The server logs also reveal regional-bias in player choice of servers. Fig.2 shows the player-hours per hour of the week, and quite clearly reveals the most popular time to play as the afternoon and evening at each server's location. The plot for UCL was brought back 8 hours, the time zone difference between Palo Alto and London, to line up the plots in Fig.2. In addition, the UCL server's logfiles showed more European and east-coast US players while the Palo Alto servers showed a bias towards mid-west and west coast US players. (Regional bias has been explored further in a study on latency sensitivity of Quake III players, which showed players actively preferring server's within 150-180ms [6].)

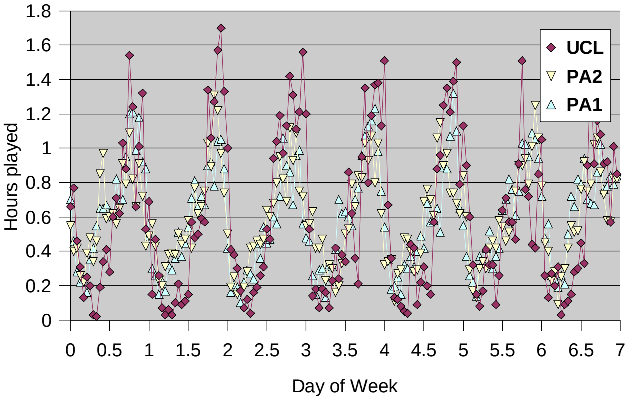

Figure 2 Played time per hour of week

We believe the servers were regularly used and attracted a reasonably large sample space. In total PA1 saw 11454 unique clients and accumulated 287 player-days over 256 calendar days, PA2 saw 7254 unique clients and accumulated 138 player-days over 170 calendar days, and UCL saw 4157 unique clients while accumulating 77 player-days over 98 calendar days. Fig.3 shows the an average of 15 to 45 minutes is played per client each month.

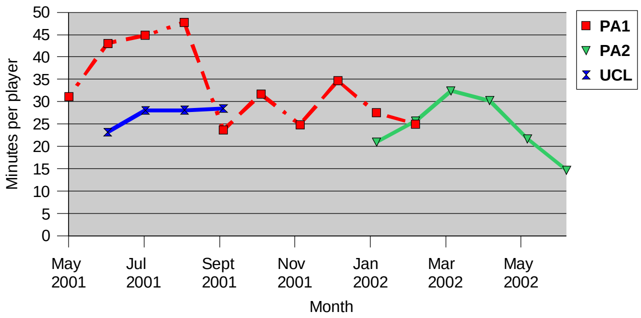

Figure 3 Average Minutes Played per Player

IV. NAPT Penetration

A. *NAPT over the trial period*

Fig. 4 summarizes the central results of our analysis. In general, NAPT appears to be present in 15% to 30% of the QuakeIII-playing population, with the results clustered primarily between 17% and 25% of the population. This provides some support for the notion that, although not dominant, there is a significant deployment of NAPT functionality across the internet. A weaker argument might also be made that Fig.4 suggests a slight drop in NAPT deployment from May 2001 to June 2002, but we don't believe the results are strong enough to draw that conclusion.

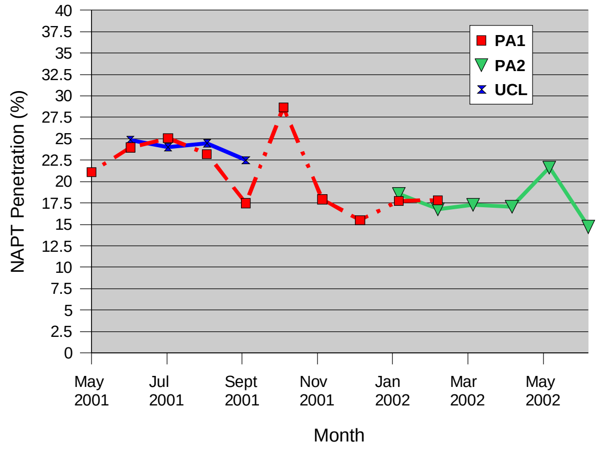

Figure 4 Overall NAPT Penetration

Interestingly, the UCL server's NAPT figures are similar to PA1. It could be that NAPT has penetrated Europe as it has the US. Alternatively, this may reflect the UCL server being probed by a lot of US-originated Quake III clients who didn't stick around long once they discovered the 'ping' time (or lag) to the London-based server. (Such clients still count towards the NAPT/no-NAPT statistics for UCL.)

A related question is: "How much of this reflects consumer/home use of NAPT-enabled gateways?" One might expect that many or most players would be connecting in from home, since it would be hard to hide a game of Quake III Arena from your boss in most corporate/work environments! Nevertheless, it is still worth putting a lower bound on the number of clients we think are connecting in from home.

B. *NAPT from Home Users*

Evidence of NAPT from a 'home user' probably represents a consumer firewall/gateway product or a software gateway such as Windows 'connection sharing'. However, inferring home users is an imprecise exercise. We are armed only with each client's IP address and therefore know only what the domain name and route database reveal to us about the source ISP.

.da.uu.net .pacbell.net .shawcable.net .dsl.gtei.net .uswest.net	.mediaone.net .earthlink.net .swbell.net .telus.net .popsite.net	.bellatlantic.net. Level3.net .dial-access.att.net .pub-ip.psi.net .inreach.net	.home.com .rr.com .aol.com .dialup.mindspring.com .dsl.mindspring.com	.lvcm.com .btinternet.com .telocity.com .rcn.com

TABLE I. Domains considered to represent consumer accounts

Table 1 is a list of domains that were common in the logfiles for PA1, and we believe are likely to represent consumer ISP accounts. PA1 overwhelmingly saw clients from the 'non-regional' top level domains (TLDs) of ".net" and ".com". Therefore we have calculated the percentage of home users as a ratio of clients falling under Table 1 to clients falling under the ".com", ".net", ".edu", and ".org" non-regional TLDs. (This simplification side-steps the problem of identifying consumer ISPs under regional TLDs such as ".jp", ".au", etc, which were hardly seen on PA1 anyway.) Clients who do not fall under one of the domains in Table 1 are considered to be non-home users, or 'work' users.

Fig. 5 shows the results of applying Table 1 to both PA1 and PA2 logfiles. In both cases the NAPT penetration for nominally 'home' users is slightly lower than the overall NAPT levels shown in Fig. 4. The curves also seem to show that the percentage of home users (relative to all clients seen in each month) is dropping off as time progresses. An adequate explanation for this has not been found.

Fig. 5 also assumes that the set of unresolvable IP addresses has the same distribution of home and work connections as clients of resolvable IP addresses. However, it is conceivable that many IP addresses are unresolvable precisely because they come from a corporate entity who has chosen not to register reverse domain names for their corporate addresses. Therefore, a

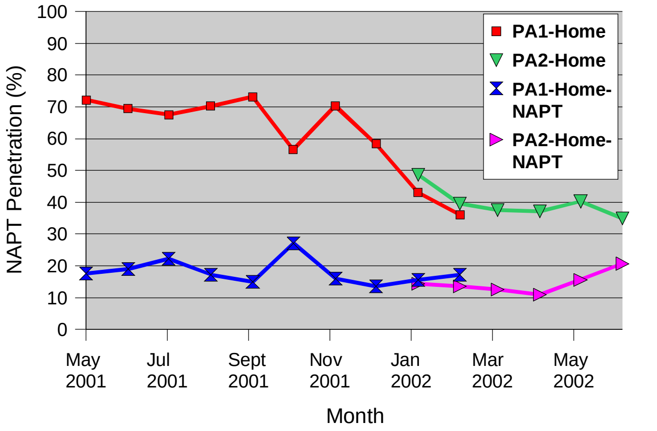

Figure 5 Percentage of Home Users over Trial Period

more pessimistic approach would treat all unresolvable IP addresses (roughly 15-20% of all clients seen in PA1's logfile) as non-home addresses, bringing the percentage of home users down by 20% in mid 2001 and 15% by early 2002. This would not, however, change the percentage of NAPT within the set of home users.

Finally, culling the regional domains means the estimate of home users will strongly reflect the distribution in North America, where the majority of non-regional domain names originate. (Most ISPs from other countries are registered under their country-specific TLD.) The UCL server was not included in Fig. 5 because we did not have sufficient insight into the consumer ISPs that exist under European TLDs. It would be interesting to investigate further whether Europe has a different pattern of home vs work users for games.

V. Discussion and Conclusion

This paper has described a relatively simple method for inferring the approximate extent to which network address port translation (NAPT) functionality has been deployed around the Internet. A recently-popular internet game, Quake III Arena, was used to attract attention of people around the world. By watching the UDP/IP address and port numbers from which people would connect, we were able to develop some estimates of NAPT deployment. Three servers were deployed - two in Palo Alto, California and one at University College London. Between them all, client address and port information has been collected from May 2001 to June 2002.

NAPT is estimated to exist in 17 to 25% of the client connections seen on our three Quake III Arena servers. Somewhere between 40 and 70% of clients are believed to connect from their home accounts, and NAPT is slightly less prevalent among clients identified as being 'home users'.

These results do not necessarily carry much validity outside the North American and European contexts. Areas such as Asia, the Middle-East, and the Pacific Rim are under-represented in my Quake III server logs. It would be interesting to measure the prevalence of NAPT in countries where ISPs have fewer IPv4 address blocks to offer customers.

Acknowledgments

Tristan Henderson provided a Linux host for the UCL server, Brian Reid provided rack space and a direct T1 connection to the Palo Alto internet exchange for my FreeBSD-based PA1 server, and Peter Allen provided a Linux host for server PA2. Their help at various stages in the past year was much appreciated. The data used in this report was gathered prior to the author joining the Centre for Advanced Internet Architectures at Swinburne University of Technology.

References

[1] P. Srisuresh, K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)," RFC 3022, January 2001

[2] T. Hain, "Architectural Implications of NAT," RFC 2993, November 2000

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

[4] Quake III Arena, http://www.quake3arena.com

[5] R. Droms, "Dynamic Host Configuration Protocol," RFC 2131, March 1997

[6] G. Armitage, "Sensitivity of Quake III Players to Network Latency and Jitter," SIGCOMM Internet Measurement Workshop (poster presentation), San Francisco, November 1, 2001