INTERNET ACCESS DENIAL BY HIGHER-TIER ISPS: A NAT-BASED SOLUTION

*Abdulaziz Al-Baiz, Marwan Abu-Amara, Ashraf Mahmoud, Mohammed H. Sqalli, Farag Azzedin**

Computer Engineering Department, *Information and Computer Science Department
King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
{g200240560, marwan, ashraf, sqalli, fazzedin}@kfupm.edu.sa

ABSTRACT

The Internet is an interconnection of Autonomous Systems (ASes) of which many are controlled by Internet Service Providers (ISPs). ASes use Border Gateway Protocol (BGP) to communicate routing information to each other. BGP does not allow a network to control how its traffic is routed. As a result, traffic belonging to a specific network can be intentionally dropped as it is routed by BGP through a malicious ISP; a behavior we define as *Internet access denial*. The impact of Internet access denial, especially when performed by higher-tier ISPs, can be severe. In this paper, Network Address Translation (NAT) is used as a solution to overcome the Internet access denial problem by hiding the traffic identity. The proposed solution is scalable to fit large networks, by using pools of IP addresses across several NAT routers. Under high network load, the performance degradation of introducing NAT on the end-to-end delay and throughput is at most 0.2% and 0.3%, respectively.

Index Terms— Internet access denial, traffic identity hiding, NAT, higher-tier ISP, OPNET

1. INTRODUCTION

The global Internet is an interconnection of Autonomous Systems. An Autonomous System (AS) is a network that is under a single administrative control. Most ASes are operated by Internet Service Providers (ISPs). ISPs are loosely classified into 3 tiers, based on their size and interconnections, with tier-1 ISPs forming the Internet core and tier-3 ISPs providing Internet service to end-users. Moreover, Border Gateway Protocol (BGP) is the dominant inter-AS routing protocol that interconnects ISPs. Routes in BGP are described as a sequence of ASes that traffic will traverse to reach its destination. BGP suffers from many security weaknesses [1].

One of the issues with BGP is the inability to control how traffic is routed through ASes. The received reachability paths can only be considered as "promises". There is no way to ensure that traffic will actually be routed through these paths [2]. BGP allows the network to control only which neighbor AS will receive the packet, but not how that neighbor AS, or any other AS in the remainder of the path, will handle that packet. Thus, the traffic may go through different paths than the advertised ones, and may go through ASes that the traffic originator is not aware of.

Consequently, many security concerns are raised because of this behavior. For example, the presence of a malicious ISP in any path to the destination results in the potential risk of routing the packets through that malicious ISP.

A malicious ISP may, for example, monitor, record, or even perform man-in-the-middle attacks. It may also blackhole the traffic that belongs to a specific network. Hence, it denies providing routing services for that particular network, preventing it from accessing many destinations.

In this paper, we tackle the problem of Internet access denial by malicious higher-tier ISPs. We define *Internet access denial* as the process of filtering transit traffic to drop packets that belong to a specific network. The ISP configures its routers to drop, or blackhole, some or all the traffic that is originated from or destined to one or more IP prefixes. We assume that the ISP will use the source IP address and the destination IP address to determine whether a packet belongs to the targeted network.

The impact of Internet access denial depends on the location, size, and connection topology of the malicious ISP. Lower-tier ISPs can only cause Internet access denial if they exist in the route of the traffic, while higher-tier ISPs may cause a larger impact.

The idea of malicious higher-tier ISPs seems unlikely at first, since ISPs that perform Internet access denial are risking their reputation, and eventually their business, as they will lose customers. However, there are several reasons that may force an ISP to become malicious and perform Internet access denial against a specific organization or country. For example, Internet access denial can be driven by political motivations, as governments may force ISPs to block Internet access to a specific region or country in an attempt to establish an Internet embargo on the targeted region. Many large services and networks have been attacked recently for political motivations. On December 2009, Gmail, for example, had many attacks targeting email

CCECE 2011 Niagara Falls, Canada
978-1-4244-9789-8/11/$26

accounts of Chinese human rights activists [3]. Twitter, a popular social network, has also been attacked during 2009 by hackers from Iran [4]. Another prime example of political motivations of a service provider to deny Internet access to an organization are the recent attempts by many governments to pressure service providers to block access to *WikiLeaks* [5]. These types of attacks are driven by political forces. Moreover, ISPs' routers may be hacked by attackers and reconfigured to drop traffic, which causes Internet access denial. Although the latter case might be temporary, it will still have an impact on the blocked network. Moreover, malicious BGP path advertisements can redirect traffic to malicious ASes, an attack technique known as BGP hijacking [1]. This type of incident has actually taken place many times in the past, where an AS, mistakenly or intentionally, advertises BGP routes that do not belong to it, and hence redirect all the traffic to a different route [6].

The remainder of the paper is organized as follows. Section 2 provides a summary of related work on Internet access denial and potential solutions. Sections 3 and 4 describe how Network Address Translation (NAT) can be adapted to provide a solution for Internet access denial. The performance of the proposed solution is evaluated in section 5. Finally, section 6 concludes the paper.

2. RELATED WORK AND SOLUTIONS

Internet access denial takes place when two conditions are met: packets are routed through a malicious ISP, and the malicious ISP drops these packets. Hence, the Internet access denial problem can be resolved by eliminating one or both of these conditions. Thus, two classes of solutions can be considered: solutions to control the traffic path so that the traffic does not pass through the malicious ISP, and solutions to prevent traffic from being dropped at the malicious ISP by concealing the traffic identity.

The first class of solutions depends on preventing the traffic from being sent through the malicious ISP by controlling the outgoing and incoming traffic. Modification of BGP is needed at all routers in the Internet to achieve this type of traffic control. Accordingly, Quoitin et al. [7] proposed BGP Tuning techniques that use AS-Path pre-pending, prefix splitting, and the use of Community to influence the path selection process of remote ASes. Virtual Peering, also proposed by Quoitin et al. [8], is a technique that uses multi-hop BGP sessions to control incoming traffic. Remote ASes establish virtual peering tunnels to control the traffic destined to the local AS. This solution is not scalable as it requires all remote ASes to implement virtual peering and establish tunnels for all communications. Alternatively, Virtual Transit, proposed by Mahmoud et al. [9], is a modification of virtual peering. With virtual transit the remote ASes advertise the virtual-peering tunnel reachability information to their neighbor ASes allowing them to use the same established tunnel to transmit traffic to the local AS, and which results in better scalability.

The other class of solutions is based on hiding traffic identity from the malicious ISP so that it does not identify the traffic's origin or destination. These techniques use IP addresses that are different from the blocked ones. Therefore, the malicious ISP will be mislead into routing the traffic without filtering it.

Network-layer encapsulation and tunnels are methods of hiding the identity where the traffic is encapsulated and carried through a tunnel created between the two tunnel endpoints. Hence, the intermediate routers will only see the two tunnel ends as the source and destination addresses.

There are many tunneling protocols, such as IP-in-IP [10], Internet Protocol Security (IPSec) [11] and Generic Routing Encapsulation (GRE) [12] that provide means to hide the identities of the source and destination from the routers that carry the traffic.

Implementing tunneling as a solution to bypass Internet access denial requires at least two cooperating networks as the endpoints of the tunnel. One of them should be located before the malicious ISP network in the route path, and the other is located after it, so that the tunnel is established through the malicious ISP. However, the performance degradation of using tunnels is expected to be significant due to the added overhead.

Network Address Translation (NAT) [13] is a technique that allows a large number of hosts (i.e., private network) to use a small set of IP addresses to communicate with other hosts on the Internet (i.e., public network). NAT can be used as an identity hiding technique, by using a set of non-blocked IP addresses as the NAT's external IP addresses. All traffic will carry these non-blocked addresses when it is sent through the Internet. The solution we adopt in this paper is based on using NAT as an identity-hiding technique at the gateway-level of the blocked network.

3. NAT-BASED SOLUTION FOR INTERNET ACCESS DENIAL

NAT provides a level of security for the private network by hiding its internal addressing structure and topology. Hence, NAT can be used as an identity hiding technique to bypass Internet access denial. The blocked network uses NAT routers as gateways to connect to their ISPs, and uses a set of non-blocked IP addresses as the NAT routers' external public IP addresses. These addresses can be obtained from a neighboring network. The outgoing packets, therefore, will not be blocked by the malicious ISP, as they will not be recognized as part of the blocked network.

Implementing the NAT solution requires enabling the NAT functionality on the gateway routers. Once NAT is enabled and configured properly, clients within the blocked network can send requests and receive responses even if traffic passes through the malicious ISP.

Although entities in the private network behind NAT are recommended to have IP addresses from the reserved private address blocks, they can still work with different IP address blocks if the NAT routers are configured properly. Therefore, for the NAT solution of Internet access denial, entities within the blocked network do not need any modifications to adapt with the NAT solution. The only modification needed is at the gateway routers.

There are many advantages of keeping the same addresses. The NAT solution would be transparent to the clients within the blocked network, as they do not have to make any changes in their networks. Moreover, local Domain Name System (DNS) servers do not have to update their records with private IP addresses, since no changes are made internally. In addition, keeping the same addresses would prevent addressing conflicts in case there are existing NAT networks within the blocked network, an issue many NAT networks suffer from [14].

4. SOLUTION SCALABILITY

Because the proposed NAT solution is meant to solve the Internet access denial problem, the blocked network can range from a small Local Area Network (LAN) to an entire country. Therefore, the deployed solution must be scalable to fit the size and requirements of the blocked network.

For a small network, a single NAT router with an external IP address is used. The NAT router is used to connect to the Internet and all the traffic is translated into its public IP address.

As the size of the private network increases, scalability issues start to appear. The first issue is the limited number of possible port-mappings. NAT maps each session to a single external port number. TCP and UDP use 16-bit port numbers, providing 65,536 ports, out of which ports 1 through 1023 are reserved. That leaves 64,512 ports usable as source ports. Hence, a NAT router is limited to mapping up to 64,512 simultaneous sessions with a single public IP address. This issue can be resolved by using a pool of public IP addresses, with each added address using the complete port space for mapping.

Other NAT scalability issues include memory, bandwidth, and processing requirements. For each NAT mapping, an entry is added to the NAT table. Since a router can map up to 64,512 sessions with a single IP address, that many NAT entries are expected to be in the NAT table.

A NAT table entry requires about 160 bytes [15]. Therefore, a fully-utilized NAT table would require a little less than 10 megabytes of memory, which is much less than the available memory in routers nowadays. Hence, the growth of the NAT table is not an issue when a single public IP address is used. However, the use of pools of public IP addresses will significantly increase the required memory. Therefore, router memory may become a limitation on the design. Moreover, the NAT router has a limited processor power such that it may not be able to handle that much traffic. Bandwidth and processor limitations need to be considered as well.

To resolve these issues, load-balancing can be used by adding more NAT routers at the gateway level. Each NAT router handles a portion of the private network, and has its own pool of IP addresses.

5. PERFORMANCE EVALUATION

Enabling NAT in a router introduces a computational overhead that, theoretically, affects performance. NAT performs a number of added operations on packets (e.g. change the source IP address, replace the source and destination ports, etc.). However, many router vendors suggest that the extra delay added by enabling NAT is negligible because routers are designed to minimize the NAT computational overhead [16][17]. Nevertheless, we evaluate the effect of NAT on network performance by first modeling the NAT processing overhead, then describing the simulation setup, and finally presenting the results.

5.1. NAT processing overhead

For proper network performance evaluation, a correct delay model of NAT must be implemented in the simulator. Ramaswamy et al. [18] have studied the network processing delay that packets experience. They estimated that on a 1Gbps network, the processing delay of complex packet modifications, including NAT, firewall, and IPSec encryption, is 1,000μs. They modeled a simplified network processor to measure the end-to-end delay that a single packet experiences. They did not consider the effect on the overall throughput, as routers are designed to improve performance by processing many packets in parallel, and the processing overhead would have a significant effect only on the end-to-end delay of a single packet. Although the study shows that processing delay is not very small, we still can consider it negligible for the NAT-based Internet access denial solution. The reason is that the measured delay is much smaller than the Internet delay, which ranges from tens to hundreds of milliseconds. Also, the study's measured delays included not only NAT, but also more complex operations such as encryption and firewall. Hence, the NAT delay is only a small portion of the measured processing delay. Moreover, routers process traffic with high-level of parallelism and pipelining. This hides the processing delay for a flow of packets. Thus, the NAT processing delay is expected not to have any significant impact on the performance of the network.

In order to evaluate the impact of implementing the proposed NAT solution on the network, simulations are performed using the OPNET network simulator [19]. The objective of the simulations is to compare the network performance before and after implementing the NAT

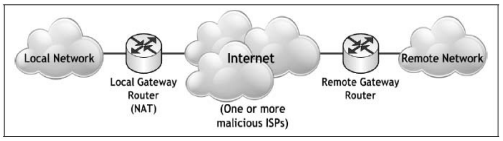

Figure 1: Simulated scenario to measure the effect of NAT delay on network performance.

solution. We select to simulate the range of NAT delay values between 10μs and 250μs. In reality, the range for real routers is between 10μs and 50μs. The remaining range, i.e. from 50μs to 250μs, does not reflect the real routers' performance. It is simulated only to see the effect of high processing delay on performance.

5.2. Simulation setup

The simulated network shown in Figure 1 consists of two networks, local and remote. Each network consists of a LAN and a gateway router. NAT is enabled in the local network's gateway router. An IP cloud, representing the Internet, is connecting the two gateway routers.

The local and remote networks are set to 100Mbps *Fast Ethernet* networks. Each network has 10 connected hosts that will serve as clients and servers for each application. The gateway routers are based on the generic router model in OPNET that supports BGP and NAT. Both routers are connected to the central Internet cloud using DS-1 links, providing a data rate of 1.544 Mbps.

Two applications are simulated: FTP which runs over TCP, and video conferencing which runs over UDP. Each application is simulated under high traffic that utilizes about 1,200 kbps (i.e. about 75% of the bandwidth). Each simulation is run 5 times, and the average of the 5 results is taken. The performance is evaluated for the end-to-end delay, traffic throughput, and packet drop rate metrics.

5.3. Results and analysis

Each simulation measures the end-to-end delay. End-to-end delay refers to the amount of time that a packet takes to travel from the client to the server, and includes transmission times, queuing delays, and added NAT delay.

The effect of NAT delay on the total end-to-end delay for UDP and TCP traffic is shown in Figure 2. When NAT is not enabled, the NAT delay is not taken into account. Hence, the end-to-end delay is constant for the NAT-disabled case. However, when NAT is enabled, the delay

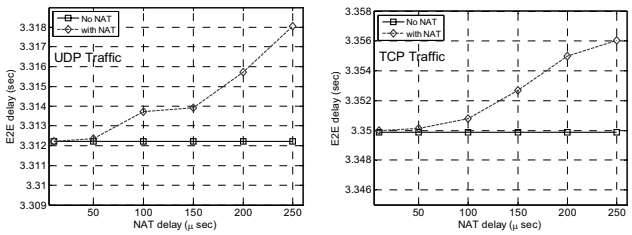

Figure 2: End-to-end delay for high UDP and TCP traffic.

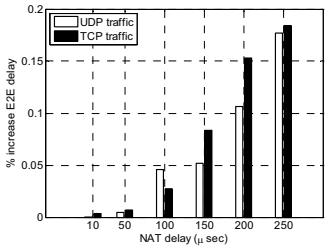

Figure 3: Relative increase of end-to-end delay for high UDP and TCP traffic.

packets suffer to reach the destination increases linearly.

The relative increase of the end-to-end delay for UDP and TCP traffic is shown in Figure 3. The relative increase is computed as $(Delay_{NAT} - Delay_{NoNAT})/Delay_{NoNAT}$. We can see that for small NAT delays, specifically below 100μs, the effect of NAT does not reach 0.02% of the total end-to-end delay. Larger values of the NAT delay cause a relatively higher increase in the end-to-end delay. However, the maximum end-to-end delay still does not exceed 0.2% of the total delay. It can be concluded that NAT does not have any significant impact on the end-to-end delay, especially for the reasonable range of NAT delay (i.e. between 10 and 50μs).

Throughput is another performance measure that is evaluated to study the impact of NAT on the amount of transmitted and received traffic. Throughput is measured as the amount of application traffic sent and received by the hosts per second. The simulation is set to measure the throughput at the client side.

NAT only starts to affect the throughput when the NAT delay is very high, i.e., more than 150μs. Figure 4 shows the throughput for UDP and TCP traffic. The degradation of throughput is due to the high NAT delay which slows down the processing of packets, and causes the router queue to be filled with waiting packets.

The relative decrease of throughput, which is computed as $(Throughput_{NoNAT} - Throughput_{NAT})/Throughput_{NoNAT}$, is shown in Figure 5. It can be noticed that the degradation of throughput starts earlier in TCP traffic as a NAT delay of 150μs causes a small decrease in the throughput. The maximum relative decrease is less than 0.3% of the total throughput, which is insignificant. Moreover, in the realistic NAT delay range, the throughput is not affected at all. We can conclude that NAT effect on the throughput of the network is negligibly small.

As for the drop rate, the simulation results show

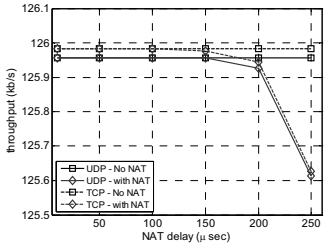

Figure 4: Throughput of high UDP and TCP traffic.

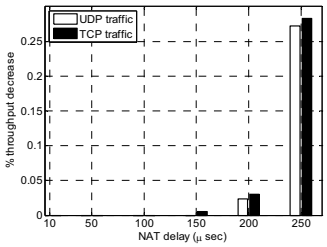

Figure 5: Relative decrease of throughput for high TCP and UDP traffic.

insignificant increase in the drop rate and, accordingly, were omitted from the paper for brevity.

It was shown that NAT does not have any significant impact on the performance of the network, and it occurs only when the NAT delay is set to very large and unrealistic values. Thus, deploying NAT as an Internet access denial solution would operate without any performance impacts.

6. CONCLUSION AND FUTURE WORK

This paper introduces the Internet access denial problem by malicious ISPs, and proposes a NAT-based solution that depends on hiding the blocked network behind a non-blocked IP address. The solution is shown to be scalable and has minimal performance impact on end-to-end delay, traffic throughput, and drop rate.

Future work would include examining the effect of NAT connectivity limitations such as private server reachability, and peer-to-peer application connectivity on the solution, and the use of different techniques, other than NAT, to bypass the Internet access denial problem.

7. ACKNOWLEDGEMENT

The authors acknowledge the support provided by King Fahd University of Petroleum and Minerals (KFUPM). This project is funded by King Abdulaziz City for Science and Technology (KACST) under the National Science, Technology, and Innovation Plan (project No. 08-INF97-4).

8. REFERENCES

[1] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," *Proceedings of the IEEE*, vol. 98, pp. 100-122, Jan. 2010.

[2] Z. Mao, J. Rexford, J. Wang, and R. Katz, "Towards an accurate AS-level traceroute tool," *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, Germany, pp. 365-378, August 2003.

[3] D. Drummond, "A new approach to china," *The Official Google Blog*, http://googleblog.blogspot.com/2010/01/new-approach-to-china.html, Jan. 2010.

[4] J. Finkle and D. Bartz, "Twitter hacked, attacker claims Iran link," *Reuters*, http://www.reuters.com/article/idUSTRE5BH2A620091218, December, 2009.

[5] "WikiLeaks," *Wikipedia*, http://en.wikipedia.org/wiki/WikiLeaks#cite_note-197, 2011.

[6] "Chinese ISP hijacks the Internet," *BGP Mon*, http://bgpmon.net/blog/?p=282, April 2010.

[7] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig, "Interdomain traffic engineering with BGP," *IEEE Communications Magazine*, vol. 41, no. 5, pp. 122-128, May 2003.

[8] B. Quoitin and O. Bonaventure, "A cooperative approach to interdomain traffic engineering," *Proceedings of Next Generation Internet Networks,* Rome, Italy, pp. 450- 457, 18-20 April 2005.

[9] A. Mahmoud, A. Alrefai, M. Abu-Amara, M. Sqalli, and F. Azzedin, "Qualitative analysis of methods for circumventing malicious ISP blocking," *Arabian Journal for Science and Engineering*, in press.

[10] C. Perkins, "IP encapsulation within IP," RFC 2003, Internet Engineering Task Force, Oct. 1996.

[11] R. Atkinson, "Security architecture for the internet protocol," RFC 1825, Internet Engineering Task Force", http://www.rfc-editor.org/rfc/rfc1825.txt, Aug. 1995.

[12] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina, "Generic routing encapsulation (GRE)," RFC 2784, Internet Engineering Task Force, Mar. 2000.

[13] K. Egevang and P. Francis, "The IP network address translator (NAT)," RFC 1631, Internet Engineering Task Force, May 1994.

[14] P. Srisuresh and B. Ford, "Unintended consequences of NAT deployments with overlapping address space," RFC 5684, Internet Engineering Task Force, Feb. 2010.

[15] J. Doyle and J. Carroll, *Routing TCP/IP, Volume II*. Cisco Press, 2005.

[16] "CISCO IOS Network Address Translation Q&A." http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6586/ps6640/prod_qas0900aecd801ba55a.html.

[17] "System Architecture Overview for the Juniper Networks SSG500 Line," Juniper Networks, http://www.juniper.net/us/en/local/pdf/whitepapers/2000177-en.pdf, February 2009.

[18] R. Ramaswamy, N. Weng, and T. Wolf, "Characterizing network processing delay," *Proceedings of IEEE GLOBECOM*, Texas, vol. 3, pp. 1629-1634, 29 Nov. - 3 Dec. 2004.

[19] "OPNET Modeler," http://www.opnet.com/.