Introduction

Internet Service Providers (ISPs) form the backbone of the Internet. They own large, worldwide networks to provide Internet connectivity to their customers. However, looking closely at how ISPs are structured to form the Internet, it can be seen that they have the most control over the Internet, at least in terms of connectivity and reachability. Large and medium-sized ISPs, often called tier-1 and tier-2 ISPs, respectively, are closer to the Internet core, and therefore, they carry most Internet traffic. The small and local ISPs, called tier-3 ISPs, are limited in carrying only traffic that belongs to their networks.

Because higher-tier ISPs control how Internet traffic is routed, the presence of one or more malicious ISPs among them can lead to many security concerns. Traffic can be monitored, critical data can be infiltrated, and packets can be modified. Even worse, traffic can be totally blocked from reaching its destinations.

The problem we are about to tackle is when a malicious ISP, usually a tier-1 or tier-2 ISP, blocks some or all the traffic that belongs to a specific network. The victim network, which may range from a single user to an entire continent, will not be able to reach some portions of the Internet, specifically the networks that are accessible through the malicious ISP. We assume that the malicious ISP uses the Internet Protocol (IP) address to identify the source or destination of a packet, and drops that packet if it belongs to the blocked victim network.

Thus, we are interested in devising solutions to the problem of blocking Internet access by ISPs that maliciously drop traffic that belongs to the victim network. The Internet in Saudi Arabia may become a victim of such type of problem; specially that none of the ISPs in Saudi Arabia is an international, higher-tier ISP. Therefore, a solution to increase Internet resilience against Internet denial is very crucial to address.

Solutions to Internet Denial

Two classes of solutions can be considered: (1) solutions to control the traffic path, so that it does not pass through the malicious ISP, and (2) solutions to prevent traffic from being dropped at the malicious ISP, by concealing the traffic identity.

In this project we are interested to examine the second class of solutions. These techniques use IP addresses that are different from the blocked ones. Therefore, the malicious ISP will be mislead into routing the traffic without altering it. Specifically, we consider two such possible solutions: (1) tunneling protocol based solution, and (2) NAT based solution.

Tunneling Protocol Based Solution

Network-layer encapsulation and tunnels are other methods of hiding the identity. Traffic is carried in a tunnel between the two tunnel endpoints. Packets are sent normally until they reach

the first tunnel endpoint. Then, each packet is optionally encrypted then encapsulated as payload into another packet, then sent to the other tunnel endpoint. The intermediate routers will only see the two tunnel ends as the source and destination addresses. Packets then are decapsulated at the other end of the tunnel, and sent to their destination. The tunneling protocols that we will consider in this project are:

(1) IP-in-IP,
(2) Generic Routing Encapsulation (GRE), and
(3) GRE with checksum (GRE-CS)

To implement tunneling as a solution to bypass Internet denial, at least two cooperating networks are needed as the endpoints of the tunnel. One of them should be located before the malicious network, and the other is located after it, so that the tunnel is established through the malicious ISP.

Network Address Translation

Network Address Translation (NAT) is a technique that allows a large number of hosts to use a small set of IP addresses to communicate with other hosts on the Internet. A NAT router separates the network into two subnetworks, a private network, where the hosts are given private IP addresses; and the public network, where the NAT router is connected to the Internet by its public IP address. 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 nonblocked addresses when it is sent through the Internet.

Implementing the NAT solution requires setting the gateway routers of the affected region to use NAT to translate all traffic into the non-blocked public IP addresses. Once NAT is enabled and configured properly, clients within the victim network can send requests and receive responses. Even if traffic passes through the malicious ISP, it will not be recognized as traffic that belongs to blocked networks, and the ISP will route it normally through its network. Although entities of 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 denial, entities within the victim network, including hosts and routers, do not need any modifications to adapt with the NAT solution. The only modification needed is at the gateway routers.

Prototyping and Evaluation of Solutions

It is desired to prototype and to evaluate the two proposed identity hiding solutions; tunneling protocol based solution, and the NAT based solution. Accordingly, the following network topology will be considered for the purpose of prototyping and evaluation.

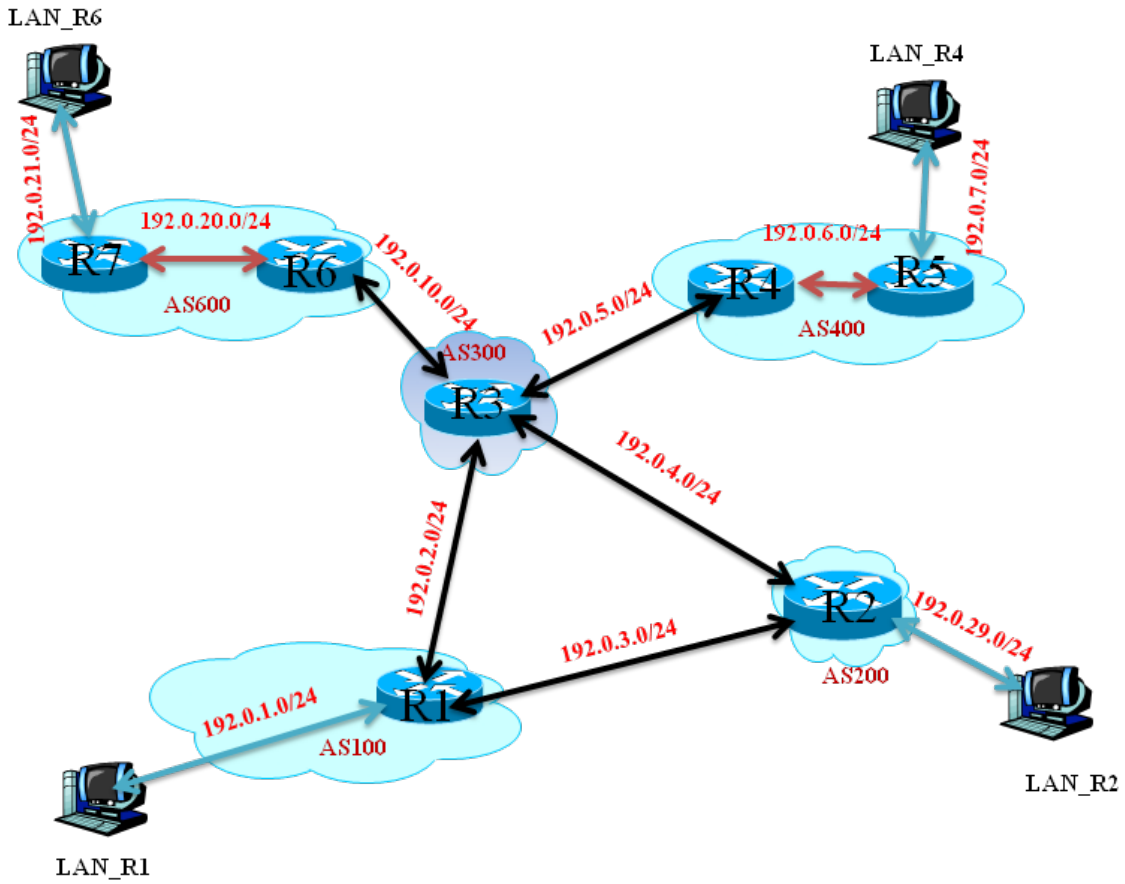

In the network topology above, AS100 acts as the affected region with R1 being the BGP gateway router for AS100. On the other hand, AS200 acts as a neighbor region to AS100 that is willing to help the affected region (i.e. AS100) by either allowing the establishment of a tunnel through it that will be used for the tunneling protocol based solution, or by providing a subset of IP addresses that will be used for the NAT based solution. The malicious ISP is represented by AS300, and the distant regions to communicate with the affected region (i.e. AS100) are represented by AS400 and AS600.

For the purpose of the tunneling protocol based solution, two tunnels must be established from AS100 to communicate with AS400 and AS600. The first tunnel must be established between R1 and R4 so that AS100 can communicate with AS400, and the second tunnel must be established between R1 and R6 so that AS100 can communicate with AS600. Both tunnels are created with the help of the neighboring region (i.e. AS200) through the malicious ISP (i.e. AS300).

Tunneling Protocol Based Solution Test Cases

The following activities must be verified before applying the solution:

(1) Traffic to and from AS100 goes through R3 when R3 is not malicious
(2) The ability to create a malicious ISP router (i.e. R3 is malicious)

(3) The malicious ISP router responds properly to BGP messages received from all ASes
(4) The malicious ISP drops traffic from/to AS100

The following activities must be verified after applying the solution:

(1) Traffic to/from AS100 from/to AS200 does not go through the established tunnels
(2) Traffic to/from AS400 from/to AS200 does not go through the established tunnels
(3) Traffic to/from AS600 from/to AS200 does not go through the established tunnels
(4) Traffic to/from AS100 from/to LAN_R4 in AS400 flows properly through the associated tunnel whether LAN_R4 is connected to R4 or R5
(5) Traffic to/from AS100 from/to LAN_R6 in AS600 flows properly through the associated tunnel whether LAN_R6 is connected to R6 or R7

NAT Based Solution Test Cases

The following activities must be verified before applying the solution:

(1) Traffic to and from AS100 goes through R3 when R3 is not malicious
(2) The ability to create a malicious ISP router (i.e. R3 is malicious)
(3) The malicious ISP router responds properly to BGP messages received from all ASes
(4) The malicious ISP drops traffic from/to AS100

The following activities must be verified after applying the solution:

(1) Traffic originated outside AS100 and destined for AS100 in the form of requested services must use the external IP address of AS100 NAT router, and that the requested services from AS100 are successfully delivered to the requester. NAT Port mapping must be used for such operations to be successful.
(2) More than one service of the services provided by AS100 can be simultaneously requested and fulfilled from outside AS100.

Project Deliverables

A full report detailing the following:

(1) Each router's configuration
(2) Verification's logs before and after applying the solution
(3) Baseline (i.e. no malicious router) performance statistics in terms of end-to-end delay, throughput, and drop rate
(4) Solution (i.e. with existence of malicious router) performance statistics in terms of end-to-end delay, throughput, and drop rate