Kingdom of Saudi Arabia
The Long-Term Comprehensive National
Plan for Science,Technology and Innovation
General Secretariat

مدينة الملك عبدالعزيز
للعلوم والتقنية KACST

Project # 08-INF97-4

FINAL REPORT

INTERNET ACCESS DENIAL BY INTERNATIONAL INTERNET SERVICE PROVIDERS: ANALYSIS AND COUNTER MEASURES

الحرمان المتعمد من قبل مزودي خدمات الإنترنت الدوليين للوصول لشبكة الإنترنت: تحليل وتدابير مضادة

Principal Investigator, ***Dr. Marwan Abu-Amara, Assistant Professor***

Department of Computer Engineering (COE)

Date: 17/11/1432
Date: 15/10/2011

NATIONAL SCIENCE, TECHNOLOGY & INNOVATION PLAN UNIT
KING FAHD UNIVERSITY OF PETROLEUM & MINERALS

PROJECT INFORMATION

Project #	**08-INF97-4**
Project Title	**Internet Access Denial by International Internet Service Providers: Analysis and Counter Measures**
Principal Investigator	**Dr. Marwan Abu-Amara**
Institution	**King Fahd University of Petroleum & Minerals**
Strategic Technology Area/Track/Sub-Track	**Information Technology – Networks and Security**
Award Period (Start Month/Year – End Month/Year)	**September 1, 2009 – August 31, 2011**
Extensions (if any)	0 months

Research Team

Senior Personnel				
No.	**Name**	**Rank**	**Role**	**Area of Specialization**
1	**Dr. Marwan Abu-Amara**	**Assistant Professor**	**P I**	**Computer Engineering**
2	Dr. Ashraf Mahmoud	Associate Professor	CO- I	Computer Engineering
3	Dr. Farag Azzedin	Associate Professor	CO- I	Information & Computer Science
4	Dr. Mohammed Sqalli	Assistant Professor	CO- I	Computer Engineering
Other Personnel				
5	Eng. Khawar S. Khan		Engineer	
Consultant				
6	Dr. Hesham Bin-Abbas (Saudi Arabia)			

المُلخص:

بما أن الإنترنت أصبح ذو أهمية بالغة للمملكة العربية السعودية، فقد تناول المشروع مشكلة الحرمان المتعمد من قبل مزودي خدمات الإنترنت الدوليين للوصول لشبكة الإنترنت. تجدر الإشارة بأن هذه المشكلة يمكن أن تحدث على مستوى التطبيق أو على مستوى التوجيه. عندما يقوم المستوى الأعلى لنظام أسماء النطاقات (DNS) بحرمان دولة أو منطقة معينة من الوصول لخدمات DNS، فإن هذه الدولة أو المنطقة سوف تفقد الوصول إلى تطبيقات الإنترنت العديدة التي تعتمد بشكل كبير على خدمات DNS مثل بروتوكول تصفح المواقع (HTTP)، وتحدث في هذه الحالة مشكلة الحرمان المتعمد للوصول لشبكة الإنترنت على مستوى التطبيق. من ناحية أخرى، عندما يتعمد مزود خدمات إنترنت دولي كيدي أن يصفي حركة المرور العابر بغرض إسقاط الحزم التي تنتمي إلى دولة أو منطقة معينة، فإنه تحدث في هذه الحالة مشكلة الحرمان المتعمد للوصول لشبكة الإنترنت على مستوى التوجيه.

وبناء على ذلك فإن المشروع قد وضع حلولا لكلا النوعين من الحرمان المتعمد للوصول لشبكة الإنترنت. وعلى وجه التحديد، فقد تم تطوير حل لمشكلة الحرمان المتعمد للوصول لشبكة الإنترنت على مستوى التطبيق مستندا على شبكات الند للند (P2P). ونظرا لإعتماد الحل المطور على إستخدام شبكات الند للند (P2P) فإن الحل المطور متين وقابل للتدرج بشكل عالي. وعلى هذا فإن الحل مناسب ليتم نشره في أي دولة أو منطقة. وبالمثل، فقد تم تطوير ثلاثة حلول مختلفة لمشكلة الحرمان المتعمد للوصول لشبكة الإنترنت على مستوى التوجيه مستندة على أساس ضبط بروتوكول بوابة الحدود (BGP)، وبروتوكولات الإتصال النفقي، وموجهات ترجمة عنوان الشبكة (NAT). فيقوم الحل القائم على أساس ضبط بروتوكول بوابة الحدود (BGP) بتوجيه حركة المرور المنتمية إلى دولة أو منطقة معينة بعيدا عن مزود خدمات الإنترنت الدولي (IISP) الكيدي، وبالتالي يحمي حركة المرور من إسقاطها من قبل ذلك المزود الكيدي. في المقابل، فإن هدف الحلين الآخرين المعتمدين على بروتوكولات الإتصال النفقي وعلى موجهات ترجمة عنوان الشبكة (NAT) هو إخفاء هوية حركة المرور المنتمية إلى دولة أو منطقة معينة بحيث يتم تضليل المزود الكيدي إلى توجيه حركة المرور بشكل إعتيادي حيث أنه قد تم إخفاء هويتها.

لقياس أهمية المشروع فيجب علينا أن نلاحظ أن الخطة الاستراتيجية للمملكة تتجه نحو اقتصاد قائم على المعرفة. وتبعا لذلك فإن نظم الحاسب والشبكات هي إحدى أولويات برنامج تقنية المعلومات والمتفرع عن الخطة الوطنية الشاملة للعلوم والتقنية والإبتكار [1]. وبالتالي فقد حقق هذا المشروع الهدف من هذه الأولوية والتي تهدف إلى ضمان مرونة الإنترنت في المملكة ضد الأنشطة الضارة وغير الضارة، بما في ذلك الحرمان المتعمد من قبل مزودي خدمات الإنترنت الدوليين للوصول لشبكة الإنترنت. من ناحية أخرى، فقد عالج المشروع نوعا جديدا من أنشطة الإنترنت الكيدية التي لم يتم تناولها من قبل الباحثين الآخرين، وفتح بذلك اتجاهات بحثية جديدة. بالإضافة إلى ذلك، فإن نتائج هذا المشروع ستسمح لمشغلي شبكة الإنترنت في السعودية، وفي المنطقة، وفي أي منطقة أخرى بأن يكونوا أقل اعتمادا على مقدمي الخدمات الدوليين. وبالمثل، فإن الحلول التي يوفرها المشروع ستكون مفيدة لمنظمي ومقدمي خدمات الإنترنت مثل مدينة الملك عبد العزيز للعلوم والتكنولوجيا، وهيئة الإتصالات وتقنية المعلومات، ولشركات النقل عبر الشبكات مثل شركة الإتصالات السعودية، وللشركات والمؤسسات المحلية. وأخيرا فقد أنتج هذا المشروع فريق متخصص من الباحثين في هذا المجال بالذات.

SUMMARY

As the Internet is becoming critically important to the Kingdom of Saudi Arabia (KSA), the project addressed the problem of *Internet access denial* by international Internet service providers (IISPs). The Internet access denial problem can occur at the application level and/or at the routing level. When a higher-level domain name system (DNS) server denies a specific country or region access to DNS services, then that specific country or region will lose access to the many Internet applications that are highly dependable on DNS services such as HTTP, and an **application level Internet access denial** takes place. On the other hand, when a malicious IISP filters transit traffic for the purpose of dropping packets that belong to a specific country or region, then a **routing level Internet access denial** occurs.

As a result, the project devised solutions for both types of Internet access denial. Specifically, one solution based on the concept of peer-to-peer (P2P) networks was developed to solve the application level Internet access denial problem. The developed solution is highly scalable and robust as a direct result of using a P2P approach for the solution. Hence, the solution is suitable to be deployed in a country or a region. Likewise, three different solutions based on border gateway protocol (BGP) tuning, tunneling protocols, and network address translation (NAT) routers were proposed to bypass the routing level Internet access denial problem. The BGP tuning-based solution directs the traffic belonging to a specific country or region around the malicious IISP, and thus protects the traffic from being dropped by that malicious IISP. Alternatively, the tunneling protocol-based solution and the NAT-based solution aim to hide the identity of the traffic belonging to a specific country or region so that the malicious IISP will be misled into routing that traffic normally since its identity is hidden.

To measure the importance of the project, we note that it is the strategic plan of KSA to move towards a knowledge-based economy. Accordingly, one of the priority technology areas for the Information Technology program of KSA's National Science, Technology and Innovation plan is ***computer systems and networks*** [1]. Hence, the project achieved the goal of this particular priority area that aims at ensuring Internet resiliency to KSA against non-malicious and malicious activities, including malicious Internet access denial by IISPs. Moreover, the project addressed a new type of malicious Internet activities that has not been addressed previously by other researchers, and opened new research directions. Furthermore, the results of the project allow the Internet network operators in KSA, in the region, or any other region to be less reliant on international service providers. Likewise, the solutions provided by the project are beneficial to Internet providers and regulators such as King Abdulaziz City for Science and Technology (KACST) and the Communications and Information Technology Commission (CITC), carriers such as Saudi Telecom Company (STC), and local businesses and institutions. Finally, the project produced a specialized team of researchers in this particular field.

ACKNOWLEGEMENT

The investigators acknowledge the support provided by King Abdulaziz City for Science and Technology (KACST) through the National Science & Technology Unit at King Fahd University of Petroleum & Minerals (KFUPM) for funding this work under project # 08-INF97-4 as part of the National Science, Technology and Innovation Plan.

TABLE OF CONTENTS

Section/Details Page

ARABIC Summary **3**
ENGLISH Summary **4**
ACKNOWLEGEMENT **5**

1.0 INTRODUCTION **9**
2.0 OBJECTIVES **12**
3.0 LITERATURE REVIEW **13**
4.0 RESEARCH METHODOLOGY **19**
5.0 RESULTS AND DISCUSSION **24**
6.0. CONCLUSIONS **52**
7.0 PROJECT OUTCOMES **53**
8.0 ADDITIONAL ACHIEVEMENTS **57**
9.0 VALUE TO THE KINGDOM **58**
10.0 BROADER IMPACTS OF THE STUDY **61**
11.0 OTHER CONCERNS **62**
12.0 REFERENCES **63**

LIST OF FIGURES

Figure 1: Impact of Internet access denial by different tiers of malicious ISPs 21
Figure 2: The lab testbed resembling a typical network configuration. 22
Figure 3: Block diagram of the research methodology. 23
Figure 4: Lookup failure. 26
Figure 5: (a) Load, and (b) load fairness. 27
Figure 6: Message path length. 27
Figure 7: Message timeout. 28
Figure 8: (a) Traffic, and (b) traffic fairness. 28
Figure 9: Control of outgoing traffic using BGP tuning. 29
Figure 10: AS Path shortening. 30
Figure 11: More specific prefixes announcement. 31
Figure 12: Control the traffic through the use of communities. 32
Figure 13: OPNET simulation model. 34
Figure 14: Convergence time for BGP-based solutions with 0.1 s average Internet delay. 35
Figure 15: BGP-based solution laboratory testbed setup 36
Figure 16: Verification of the malicious activity and the BGP-based solution 37
Figure 17: Virtual peering-based solution. 38
Figure 18: Virtual transit scheme. 39
Figure 19: IP tunnel traffic received and sent by routers R2 and R5. 40
Figure 20: FTP relative increase in end-to-end delay. 41
Figure 21: FTP absolute increase in end-to-end delay. 42
Figure 22: Video conferencing relative increase in end-to-end delay. 42
Figure 23: Video conferencing absolute increase in end-to-end delay. 43
Figure 24: Relative increase in throughput overhead. 43
Figure 25: Absolute throughput overhead. 44
Figure 26: Tunnel-based solution laboratory testbed setup 45
Figure 27: Extended NAT design using load-balancing over a number of NAT routers. 47
Figure 28: Simulated scenario to measure the effect of NAT delay on network performance. 48
Figure 29: End-to-end delay for high UDP and TCP traffic. 48
Figure 30: Relative increase of end-to-end delay for high UDP and TCP traffic. 49
Figure 31: Throughput of high UDP and TCP traffic. 49
Figure 32: Relative decrease of throughput for high TCP and UDP traffic. 50
Figure 33: NAT-based solution laboratory testbed setup 51

LIST OF TABLES

Table 1: Mapping between project phases and project objectives 20
Table 2: Default parameters 26
Table 3: Comparison between BGP tuning methods 33

1.0 INTRODUCTION

The Internet has become an important means of communication which allows us to interact, retrieve and propagate information easily with the rest of the world. Many businesses and governments are highly dependent on the availability of the Internet for conducting their daily affairs. Accordingly, the economic impact of the unavailability of the Internet on KSA will be severe. Furthermore, since most of the means of communication such as telephone, mobile, and email are merging with the Internet, virtually all people are directly or indirectly affected by the Internet. Therefore, we can realize the social impact and importance of the availability of the Internet to KSA.

It should be noted that the unavailability of the Internet could be due to several causes that can be divided into two main categories: **non-malicious** and **malicious**. The non-malicious causes include misconfigurations to Internet equipment, hardware and/or software failures, and congestion due to increase in traffic demand. On the other hand, malicious causes include denial of service attacks, terrorist attacks, security breaches and attacks such as viruses and worms, purposeful hardware failures, and deliberate denial of Internet access by service providers. Whether the cause is non-malicious or malicious it can affect the availability of the Internet at the infrastructure level, the routing level, and/or the applications level. The resilience of the Internet has been widely studied in the literature where different causes of Internet isolation were identified and many solutions were proposed. However, the malicious activities that had the least attention in the literature are purposeful hardware failures, terrorist attacks, and the deliberate denial of Internet access by Internet Service Providers (ISPs). Although the idea of malicious and intentional denial of Internet access by ISPs seems unlikely at first, there are several reasons that may force an ISP to become malicious and perform intentional denial of Internet access against a specific organization or country. For example, intentional denial of Internet access by ISPs can be driven by political motivations, as governments may force the ISPs to block Internet access in order to establish an Internet ban on the targeted region. Many large services and networks have been attacked recently for political motivations. For example, on December 2009, Gmail had many attacks targeting email accounts of Chinese human rights activists [2]. Twitter, has also been attacked during 2009 by hackers from Iran [3]. Another prime example of political motivations of a higher name server to perform intentional denial of service against an organization are the recent attempts by many governments to pressure service providers to block access to WikiLeaks [4].

As such, the project focused on devising solutions to improve the resilience of the Internet access when an ISP maliciously denies access to the Internet to any particular region. The project considered solutions for different scenarios of deliberate denial of Internet access by international service providers that could affect the application level and the routing level. The physical level was not considered by the project as it is considered to be the easiest to countermeasure amongst all three levels from the point of view of deliberate denial of Internet access by the ISPs.

For the applications level, the main focus was on the situation when a service provider blocks a region from accessing either the root domain name system (DNS) servers or the top level DNS (TLD) servers and thus causing inaccessibility to major Internet applications such as web surfing (HyperText Transfer Protocol, HTTP), email (Simple Mail Transfer Protocol, SMTP), and file transfer (File Transfer Protocol, FTP). Because of the DNS central role to the Internet availability and to the functionality of other Internet applications, the project intentionally focused only on the DNS denial of Internet access by a service provider. On the other hand, the denial by a service provider to access other Internet applications will only have a limited impact of denying that particular service but not an impact on the Internet availability. To solve the intentional DNS blocking, the project proposed a solution that combines the use of a peer-to-peer (P2P) network with the use of a round-robin approach. Using the P2P network creates a scalable and stable network, while the round-robin approach provides a fast path to resolve queries. Thus, the proposed solution is scalable. Moreover, the proposed solution was evaluated by means of simulations and it was found to be resilient. Finally, the proposed solution allows the simultaneous submission of DNS queries to both the standard path (i.e., through root/TLD DNS servers) and the P2P network.

For the routing level, the project addressed the case when an ISP refuses to route traffic from or to a particular region. To solve this problem, the project identified that the Internet access denial at the routing level takes place when two conditions are met: packets are routed through a malicious ISP, and the malicious ISP drops these packets. Hence, this problem can be resolved by eliminating one or both of these conditions. Therefore, two classes of solutions have been considered by the project: solutions to control the traffic path, so that it does not pass through the malicious ISP, and solutions to prevent traffic from being dropped by the malicious ISP by concealing the traffic identity.

The first class of solutions to the Internet access denial problem proposed by the project at the routing level is based on the use of Border Gateway Protocol (BGP) tuning techniques that included the use of autonomous system (AS) path shortening, more specific prefixes, and communities.

The second class of Internet access denial solutions proposed by the project at the routing level 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 misled into routing the traffic without filtering it. The project proposed two different solutions that are based on hiding the traffic identity.

The first proposed solution based on identity hiding utilizes network-layer encapsulation and tunnels where the traffic is carried through a tunnel created between two tunnel endpoints that use IP addresses that are different from the blocked ones. The proposed solution requires at least two cooperating networks as the endpoints of the tunnel. One of them should be located before the malicious network in the route path, and the other is located after it, so that the tunnel is established through the malicious ISP. As such, the intermediate routers of the malicious ISP will only see the two tunnel endpoints as the source and destination addresses, and therefore will not drop the traffic belonging to the blocked region.

The second proposed solution related to identity hiding is based on Network Address Translation (NAT) which 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 using its public IP address. Accordingly, 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 belonging to the blocked region will then use these non-blocked IP addresses when it is sent through the Internet. The proposed solution requires enabling NAT at the gateway-level of the blocked network.

The project's proposed solutions are directly related to the ***computer systems and networks*** priority technology area of the Information Technology program of KSA's National Science, Technology and Innovation plan [1]. Specifically, the proposed solutions are perfect match for the plan's requirement to conduct research that addresses problems in computer systems-related areas that are needed by the kingdom which include computer networking, and IT security and privacy [1]. Hence, the project aimed at ensuring Internet resiliency to KSA against malicious activities, namely the malicious act by ISPs to deny access to the Internet. Moreover, the project addressed a new type of malicious Internet activities that has not been addressed previously by other researchers, and opened new research directions. Furthermore, the results of the project allow the Internet network operators in KSA, in the region, or any other region to be less reliant on international service providers. Likewise, the solutions provided by the project are beneficial to Internet providers and regulators such as King Abdulaziz City for Science and Technology (KACST) and the Communications and Information Technology Commission (CITC), carriers such as Saudi Telecom Company (STC), and local businesses and institutions. Finally, the project produced a specialized team of researchers in this particular field.

2.0 OBJECTIVES

The main objective of this project was to investigate Internet unavailability due to the malicious act of denial of Internet access by international Internet service providers and devise proper countermeasures. The specific objectives are as follows:

1. Investigate the various methods an international Internet service provider may employ to deny Internet access to an entity including DNS access denial, and traffic routing refusal.
2. Explore existing and novel approaches that improve resilience of Internet access.
3. Devise solutions that follow the standards as close as possible.
4. Propose solutions that enhance Internet access resilience but with minimized deployment cost.
5. Propose solutions that are suitable for typical local and global needs and specifications, and that can be applicable to Saudi Arabia, the Gulf States, the entire Arab region, or any other region.
6. Carry out experimentations of the devised solutions, explore their feasibility, and evaluate their performance.
7. Propose an appropriate mechanism to realize the results and provide recommendations of the research.

3.0 LITERATURE REVIEW

The growing importance of the Internet has motivated many studies on the Internet resilience against different types of outages, failures, and attacks. Internet unavailability takes place due to either accidental or malicious causes at the physical, routing, and/or application levels. Some of the malicious activities that may cause Internet unavailability include Denial-of-Service (DoS) attacks, security breaches, terrorist attacks, intentional hardware failures, and deliberate Internet access denial by service providers. Most of the research that has been done in this area targets DoS attacks and security breaches [5][6][7][8]. Only few research efforts targeted terrorist attacks and intentional hardware failures [9][10]. A comprehensive list of the different causes of the unavailability of the Internet can be found in [11].

Of a particular interest, the unavailability of the DNS system has a direct impact on the Internet unavailability at the **application level**. The DNS system can be attacked in many different ways as described by Cheung [12]. Cheung divided the DNS system into three sections; the low level servers (resolvers), the communications between the servers and the clients, and the top level servers (nameservers). In the resolvers, the attackers try to get access to the DNS system through the available vulnerabilities which could lead to damaging the resolver servers or misbehavior of the resolver servers. By flooding the DNS servers with a massive number of forged queries, the normal DNS queries could be dropped by the routers. This could also happen if the attacker was able to take advantage of a critical router along the DNS server path to cause it to misbehave. Finally, the nameservers could be damaged by directly accessing the nameservers or through damaging other necessary services to the nameservers. Other malicious attacks that take advantage of the DNS caching aspect include cache poisoning. In such an attack, the attacker inserts an incorrect DNS record or modifies an available DNS record in the DNS cache server so that it redirects the client to the attacker's website [13][14].

To countermeasure malicious attacks on the DNS, several solutions were proposed. The robustness of the name servers could be increased by using "Anycast routing" that counters the flooding attacks [15]. Another countermeasure uses the DNS security extensions to provide a secure integrity and authenticity to the DNS traffic [16]. Researchers in [15] and [17] discussed how to manipulate the time-to-live (TTL) value of the DNS records to improve the availability of the DNS system against DoS attacks. A solution provided in [18] increases the efficiency of the DNS caching by suggesting two policies. The first policy is renewal cache refreshment where the DNS record is refreshed upon the expiration of the TTL. The second policy is simultaneous validation where the expired DNS record is refreshed whenever a DNS lookup requests this record.

To ensure the availability of the root DNS servers, a secondary DNS lookup service through a cooperative cache sharing network that is based on Pastry and Chord P2P protocols was proposed in [17]. It was found that the effort to attack a domain becomes harder when the resource records of the domain are distributed over different cache resolvers.

A new platform was proposed in [19] that implements a Cooperative Domain Name System (CoDoNS) based on a structured P2P Pastry protocol with analytically-informed proactive caching used to distribute the popular records in the nodes near the home node. An alternative

solution uses Chord P2P protocol to replace the existing DNS [20]. In addition to implementing all the functionality of the existing DNS, the new domain record requires authentication from the higher level. Although the design had more availability and better fault-tolerance, it suffered from high latency.

Of a particular interest also is the fact that BGP can be considered to have a responsibility for the Internet unavailability at the **routing level** as one of the issues with BGP is the inability to control how traffic is routed through ASes. The received prefix reachability paths can only be considered as "promises". There is no way to ensure that traffic will actually be routed through these paths. Practically, routers may provide the list of ASes that propagated the BGP update messages, which are not necessarily the same as the list of ASes traversed by data packets [21]. 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. Moreover, many networks use load-balancing and multihoming techniques to distribute traffic over multiple links. Thus, the traffic may go through different paths other than the advertised ones, and may go through ASes that the traffic originator is not aware of.

This issue does not normally affect the delivery of traffic, as packets will eventually reach their destinations regardless of the used path. However, many security concerns are raised because of this behavior. Packets may go through ASes that the traffic originator is unaware of, as they do not appear in the AS path. 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, or a hacked into ISP can, for example, monitor, record, or even modify packets that are routed through it, performing man-in-the-middle attacks. It may also *blackhole* the traffic that belongs to a specific network (referred to as the *victim* network), i.e., drop all the packets originated from or destined to the victim network. Hence, it denies providing routing services for that particular network, preventing it from accessing many destinations, namely the ones that are reachable through paths that go through the malicious ISP. Accordingly, the Internet access denial at the routing level by malicious or hacked into ISPs is defined as the process of filtering transit traffic to drop packets that belong to a specific network.

Accordingly, it is important to explore BGP further while noting that BGP was not designed with protection and security in mind and that makes it vulnerable to many attacks. The study in [22] identifies three main security-related drawbacks of BGP. The first is that BGP does not check the integrity and the freshness of BGP messages, and it does not provide authentication of the origin. The second drawback is that BGP does not check the validity of the AS-Path announced by a specific AS. Thirdly, BGP speakers receiving announcements do not typically check for the genuineness of path attributes announced by a specific AS [22][23]. A more recent and comprehensive survey of BGP security issues can be found in [24]. The survey clearly states that BGP while being the dominant, if not the only, inter-domain routing protocol, it fails to adequately address security. The study refers most of the security problems in BGP to one or more of the following reasons: 1) uncertainty of the mapping between the IP prefixes and the AS numbers for the ASes that manage them; 2) the

utilization of the Transmission Control Protocol (TCP) as the underlying transport protocol; and 3) the potential to produce false or incorrect route announcements to undermine a specific BGP routing policy. Another comprehensive survey that focuses on the various security techniques that can be implemented for BGP, as opposed to the vulnerabilities and their root causes, can be found in [25].

Nordstrom and Dovrolis [26] discuss the types of BGP attacks and point out four main goals for BGP attacks. The first goal is *blackholing*, which is to drop the traffic that arrives to the router. *Redirection* is the second goal, where the attacker sends the traffic to a different destination for analysis of the contained data. The third goal is *subversion* which is similar to the redirection attack, but with the intention to eavesdrop or modify the data and then forwarding the packet to the original destination. Finally, the fourth goal of BGP attacks is to *cause instability in the network*, which may happen by sending successive advertisements and withdrawals. This attack may also occur by sending false update or by prefix hijacking through announcing a prefix that the hijacker does not own or advertising a path it does not have. Another similar attack is to announce link flapping, i.e., the announcement of link failure then followed by the announcement of the recovery of the same link several times, to trigger route flap dampening [27]. An incident of network instability due to prefix hijacking happened in April 1997 when AS7007 advertised most of the routes of the Internet causing an Internet outage for more than two hours [26][28][29]. Another incident of network instability occurred in April 2001 when AS3561 forwarded a huge number of wrong advertisements from one of its downstream customers causing problems in connectivity [28][29]. There are many proposed counter measures against these types of attacks identified in [26], but the paper discussed only two of them. The first one is the use of route filtering in order to enable ASes to filter out malicious or faulty updates, however this requires ASes to know what to filter. In order to get ownership information of prefixes, Internet routing registries (IRR) databases can be consulted; however these databases are always not up-to-date [26]. The second solution is the use of a Secure Border Gateway Protocol (S-BGP) [30]. Although this method can provide high security against attacks, it will add high overhead on the Internet. Wang and Wang [31] improve on BGP by employing a verification mechanism referred to therein by the Assignment Track (TA) where all ASes are to provide assignment and attestations of their announced prefixes. The study shows that this TA-based scheme is superior to S-BGP. A summary of ongoing efforts to secure BGP on the standardization front is also briefly presented in [32].

The prefix hijack attack can be addressed by using a prefix hijack alert system (PHAS) [33]. PHAS is an email notification system that alerts a prefix owner whenever there is a change in the origin AS that owns the prefix. Every day there are a number of prefix changes and most of them are valid. However, only the AS that owns the prefix is capable of differentiating between a valid origin change and a prefix hijack [33]. The system examines the data collected in RouteViews [34] and notifies the prefix owner about any possible hijack.

Similarly, Zheng et al. [35] build an IP hijacking detection system. Their system depends on two observations noticed when there is no hijacking. The first observation is that the number of hops from a source to the prefix generally does not change or is stable. The second

observation is that the path from a source to the prefix covers the path from the source to a reference point along the original path that is topologically close to the prefix. The study focuses on two types of hijacking; the first type is referred to by the *imposture* where the attacker imitates the behavior of the victim by responding to the sender of the hijacked traffic. The second is the *interception* type where the attacker spies on the traffic and records its content and then forwards it to the correct destination.

In [36], Hu and Mao implement a prefix hijacking identification system. The utilized technique is based on collecting data from the control plane, i.e., passively collected BGP updates, and on data collected from the data plane. The latter process is referred to as *fingerprinting*. Fingerprinting removes the ambiguity about an expected IP hijacking occurrence because it is based on information such as host operating system properties, IP identifier, TCP timestamp, and ICMP timestamp to identify the hijacker. The authors note that it is not possible for an attack to affect the whole Internet; more specifically routers which are close to the legitimate prefix owner most likely will not be affected. In addition, the authors devise ways to counteract a number of IP prefix attack types.

The study by Quoitin [37] proposes the design and implementation of a BGP modeling tool called C-BGP that is used to compute routes in a large scale network topology [38]. The study considers approaches of controlling outgoing and incoming traffic when the Regional ISP is multi-homed. Usually, this control is either for the purpose of having a backup route or for load balancing. In his dissertation, he proposed a cooperative approach, called virtual peering, in order to provide a deterministic approach of controlling the incoming traffic [39]. The proposed approach modifies slightly the BGP protocol by automating the establishment of virtual peering and adding BGP messages specifically to accomplish that. Only end systems need to have these modifications and none are required for intermediate systems. Virtual peering is used to achieve load balancing for the incoming traffic and for selecting the path with the lowest delay.

Given the previous review of BGP and its impact on the Internet access denial, it should be noted that the Internet access denial takes place at the routing level 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. Therefore, two classes of solutions can be considered: solutions to control the traffic path, so that it does not pass through the malicious ISP, and solutions to prevent traffic from being dropped by the malicious ISP by concealing the traffic identity.

The first class of solutions to the Internet access denial problem depends on preventing the traffic from being sent through the malicious ISP. As stated earlier, although BGP provides reachability information that includes the AS-path, it does not allow a network to control the actual routing path of its traffic. A network can only select which neighbor ASes will route its packets, but does not know how that neighboring AS is going to handle them [21].

Controlling the outgoing and incoming traffic requires modifications or adjustments of the routing protocols. *Source Routing* [40], which allows the traffic originator to specify the path its traffic will travel through, is a solution to control the outgoing traffic so that it avoids the malicious ISP. However, the existing Internet protocols do not implement this type of

routing. Modification of BGP is needed on all routers in the Internet to achieve this type of traffic control.

The other class of Internet access denial 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 misled into routing the traffic without filtering it.

One solution is to change the IP addresses of the victim network to different ones. The victim network can just register a new IP block and use it instead of its current one. This, however, only provides a temporary solution, as the malicious ISP can easily detect the new IP block and will simply block it again. Hence, this solution is not robust.

Network-layer encapsulation and tunnels are other methods of hiding the identity. Traffic is carried through a tunnel created between the two tunnel endpoints. First, packets are routed as usual until they reach the first tunnel endpoint. At the first tunnel endpoint, each packet is encrypted (optionally) and encapsulated, as payload, into another packet, and 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 endpoint of the tunnel, and sent to their destination.

Such usage of tunneling protocols is common to achieve better Internet resiliency. For example, Kini et al. [41] used IP-in-IP tunneling protocol [42] to enhance the robustness of a network to dual link failures. Similarly, Wu et al. [43] considered a failure scenario that breaks an AS into two or more isolated parts and disrupt the connectivity among these AS partitions. Wu et al. proposed the use of tunneling techniques by the neighbors of the affected AS partitions to provide extra connectivity to bypass the failure. Thus, the AS partitions can communicate with each other through the use of tunneling protocols.

There are many tunneling protocols, such as IP-in-IP, Internet Protocol Security (IPSec) [44], and Generic Routing Encapsulation (GRE) [45]. Additionally, anonymous routing protocols, such as Onion Routing [46], Cashmere [47], Crowds [48], and Hordes [49], provide means to hide the content of the packet, as well as the identities of the source and destination, from the routers that carry the traffic. Moreover, virtual peering, proposed by Quoitin [37] and that was discussed earlier, attempts to automate the establishment of tunnels (i.e., virtual peering) by using multi-hop BGP sessions. Remote ASes establish virtual peering tunnels to control the traffic destined to the local AS. However, this solution is not scalable as it requires all remote ASes to implement virtual peering and establish tunnels for all communications. The project enhanced the virtual peering method to make it more scalable. The enhanced method is referred to as *virtual transit* [50].

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 network in the route path, and the other is located after it, so that the tunnel is established through the malicious ISP. It should be noted that the use of anonymous routing protocols as a solution for Internet access denial results in a very high performance degradation [46][52].

Network address translation (NAT) [53][54] 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 using 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 then use these non-blocked IP addresses when it is sent through the Internet.

The project considered solutions to resolve the Internet access denial problem at both the application level and the routing level. Specifically, the project solved the problem of the Internet access denial at the application level by proposing a P2P solution for the intentional blocking of the DNS system. Furthermore, the project addressed the problem of Internet access denial at the routing level by providing one solution that is based on preventing the traffic from being sent through the malicious ISP, and two solutions that are based on hiding traffic identity from the malicious ISP.

4.0 RESEARCH METHODOLOGY

To achieve the objectives outlined earlier, the team used a combination of theoretical, experimental, and developmental approaches. Accordingly, the team proposed to first conduct an extensive literature survey to identify and review any additional scenarios where a service provider is the main cause of the denial of Internet access, and to identify and review approaches and solutions proposed that address this type of problems. Also, the team reviewed and studied proposed approaches and solutions in the literature that address other non-malicious and malicious causes of Internet isolation, and that address the robustness in the Internet in general. The purpose of this step was to assess the suitability and the adaptability of such approaches and solutions in solving the problem of Internet denial of access caused by a service provider.

At the same time, to ensure a thorough understanding of all aspects of the problem, the team proposed to analyze the problem by investigating the various ways a service provider may use to deny access to the Internet on each of the two major areas that can be the cause of Internet isolation, i.e., applications and routing. In addition, the team identified a typical network topology for consideration of the problem. The typical network topology was chosen to resemble what is usually deployed in countries such as KSA.

Subsequently, based on the previous two steps, the team proposed to investigate and study the effectiveness, the suitability, and the applicability of existing solutions proposed in the literature to the identified topology.

Consequently, the team proposed to investigate ways to enhance, extend, and adapt existing approaches and solutions to the service provider denial of Internet access problem. Further, the team investigated new approaches that are specific to this problem. To complete this step, the team defined and formulated the hardware and software requirements (e.g., need for redundant servers, protocol changes, etc.) as well as the techniques to be used in the proposed approaches.

Moreover, to consider the feasibility and the adequacy of the proposed solutions, the team proposed to devise solutions that follow standards as close as possible, and that do not require many changes to existing protocols. In addition, the team considered providing several typical deployment scenarios and possible solutions to each scenario so that the solutions can be used locally as well as globally. Accordingly, the team identified solutions to each of the two major types, i.e., applications and routing, of the service provider denial of Internet access problem that require a minimal deployment cost in terms of the amount of effort needed to implement such solutions.

Finally, the team developed a testbed for the routing level of the service provider denial of Internet access problem. To verify, validate, and evaluate the performance, the team built an experimental setup and carried out experiments for the proposed solution. Moreover, the team evaluated the performance of the proposed solution in terms of end-to-end delay, throughput, and drop rate.

Throughout the project, the team documented all findings, and recommended appropriate and generic solutions, for the local and global needs, to the problem of the denial of Internet

access by service providers. All findings and recommendations were communicated on a regular basis in the form of progress reports as well as conference and journal publications.

In order to carry out the methods and approaches listed earlier, the project was divided into the following eight phases:

1. Literature review update.
2. Analysis of the problem.
3. Identification and study of existing solutions.
4. Investigation of new approaches.
5. Devising of solutions that follow standards and provide required functionalities.
6. Assessment of deployment cost of devised solutions.
7. Prototyping and performance analysis.
8. Reporting and documentation.

Accordingly, Table 1 provides a mapping between the different phases of the project and the corresponding project objectives achieved.

Table 1: Mapping between project phases and project objectives

Phase	Achievements
Phase 1: Literature review update	**Objective 1** (partially)
Phase 2: Analysis of the problem	**Objective 1**
Phase 3: Identification and study of existing solutions	**Objective 2** (partially)
Phase 4: Investigation of new approaches	**Objective 2**
Phase 5: Devising of solutions that follow standards and provide required functionalities	**Objectives 3 and 5**
Phase 6: Assessment of deployment cost of devised solutions	**Objective 4**
Phase 7: Prototyping and performance analysis	**Objective 6**
Phase 8: Reporting and documentation	**Objective 7**

As such, the team started with phase 1, the literature review update, to identify the additional scenarios, approaches, and solutions that could be applicable to the service provider denial of Internet access problem. As a result of phase 1 the team identified that the Internet access denial problem can occur at the application level and/or at the routing level. When a higher-level domain name system (DNS) server denies a specific country or region access to DNS services, then that specific country or region will lose access to many Internet applications that are highly dependable on DNS services such as HTTP, and an application level Internet access denial takes place. On the other hand, when a malicious ISP filters transit traffic for the purpose of dropping packets that belong to a specific country or region, then a routing level Internet access denial occurs.

With respect to the routing level, the team characterized the effect of the location, size, and connection topology of the malicious ISP on the Internet access denial problem. The team noted that ISPs are loosely classified into 3 *tiers*, based on their size and interconnections.

Tier-1 ISPs own large networks that cover one or more than one continent, and they form the Internet core. Tier-2 ISPs are smaller networks that mostly cover one or few countries. Tier-3 ISPs are the smallest, covering a country or a metropolitan area of a country. Tier-3 ISPs provide Internet service to end-users, and connect to one or few larger ISPs for the delivery of their customers' traffic to destinations outside their networks. Higher-tier ISPs, i.e., tier-1 and tier-2 ISPs, carry not only traffic that belongs to their networks, but also traffic that is originated from or destined to one of the networks they are connected to. Thus, packets that are sent from one end-user to another are carried over multiple different tier ISPs. As such, the team identified that lower-tier ISPs can only cause Internet access denial if they exist in the route of the traffic, while higher-tier ISPs may have a larger impact.

Because tier-3 ISPs do not act as transit for other networks, they only carry traffic that belongs to their networks. Therefore, a malicious tier-3 ISP can only block access to its own network. Hence, the impact of this type of ISP is limited to only a small set of hosts and services. On the other hand, malicious higher-tier ISPs can have more impact as they can block not only traffic that belongs to their networks, but also all other traffic that passes through them in transit. For example, a malicious tier-2 ISP blocks access to its own network, and to all its customer ISPs' networks. Furthermore, Internet access denial by tier-1 ISPs presents a more critical problem. A malicious tier-1 ISP can isolate the victim network and block it from accessing a large portion of the Internet. Figure 1 shows a simplified network of ISPs of different tiers, and how Internet access denial by higher-tier ISPs results in a larger inaccessibility to other parts of the network.

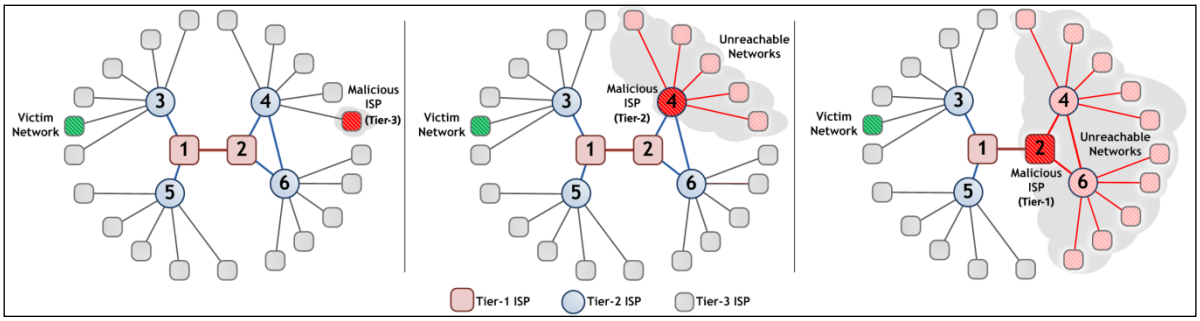

Figure 1: Impact of Internet access denial by different tiers of malicious ISPs

With the aid of the outcomes of phase 1, an investigation was carried out in phase 2 to identify the various ways a service provider may use to deny access to the Internet as well as to identify a typical network topology for consideration of the problem. Towards the end of phase 2, the team identified that the Internet access denial at the routing level takes place when two conditions are met: packets are routed through a malicious ISP, and the malicious ISP drops these packets. Hence, the team recognized that the Internet access denial problem at the routing level can be resolved by eliminating one or both of these conditions. As such, the team identified that there are two classes of solutions that can be considered: solutions to control the traffic path, so that it does not pass through the malicious ISP, and solutions to prevent traffic from being dropped by the malicious ISP by concealing the traffic identity.

In phase 3, the outcomes of both phase 1 and phase 2 assisted in the identification of the most suitable, effective, and applicable of the existing solutions to the service provider denial of Internet access problem. By the end of phase 3, the team identified that the concept of peer-to-peer (P2P) networks is suitable to devise a scalable and robust solution for the Internet access denial problem at the application level. Likewise, the team recognized that BGP tuning concepts that are used for traffic engineering can be extended to provide a solution that falls under the traffic controlling class of solutions to the routing level Internet access denial problem. On the other hand, the team identified that the concepts of Network Address Translation (NAT) and tunneling can be extended to provide a solution that falls under the traffic identity hiding class of solutions to the routing level Internet access denial problem.

To devise solutions to the service provider denial of Internet access problem in phase 4, the outcomes of both phase 1 and phase 3 were utilized. Once the devised solutions were identified at the end of phase 4 along with the hardware and software requirements, then both phase 5 and phase 7 commenced. Phase 5 determined how close the devised solutions are to the standards, and to some extent influenced phase 7. On the other hand, phase 7 prototyped the devised solutions for the routing level of the problem and measured the network performance through simulations using OPNET network simulator [55], and lab experiments using a testbed. The lab testbed, shown in Figure 2, consisted of a minimum of six routers and two switches interconnected to resemble a typical network configuration. The lab testbed has AS 100 as the blocked region, AS 200 as a neighboring cooperative region, AS 300 as the malicious international ISP, and AS 400 as a distant region. Both the OPNET simulations and the lab experiments measured the network performance in terms of end-to-end delay, throughput, and drop rate before and after applying each devised solution. Details of both the OPNET simulations and the lab experiments are provided in Section 5.

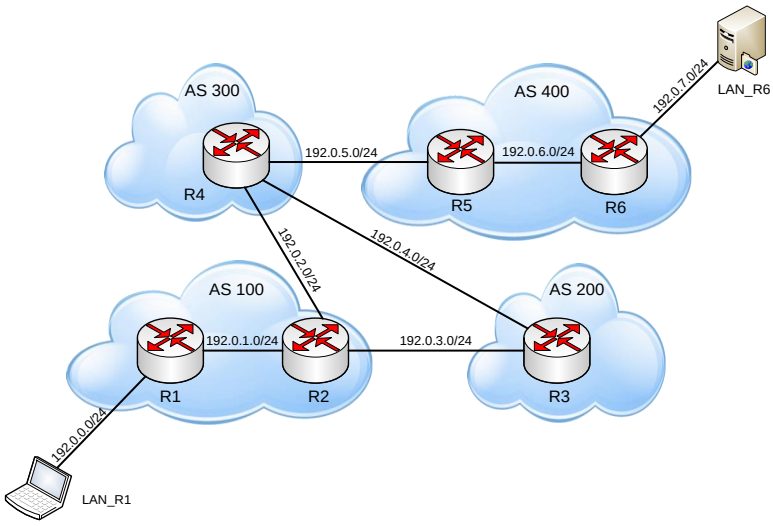

Figure 2: The lab testbed resembling a typical network configuration.

Both the OPNET simulations and the lab experiments measured the network performance in terms of end-to-end delay, throughput, and drop rate before and after applying each devised solution.

To consider the impact of the devised solutions that were determined in phase 5 on the cost of deployment effort, a comparison between the different devised solutions took place in phase 6. Throughout the project, all findings were documented and published as part of phase 8. Figure 3 provides a block diagram that summarizes the research methodology.

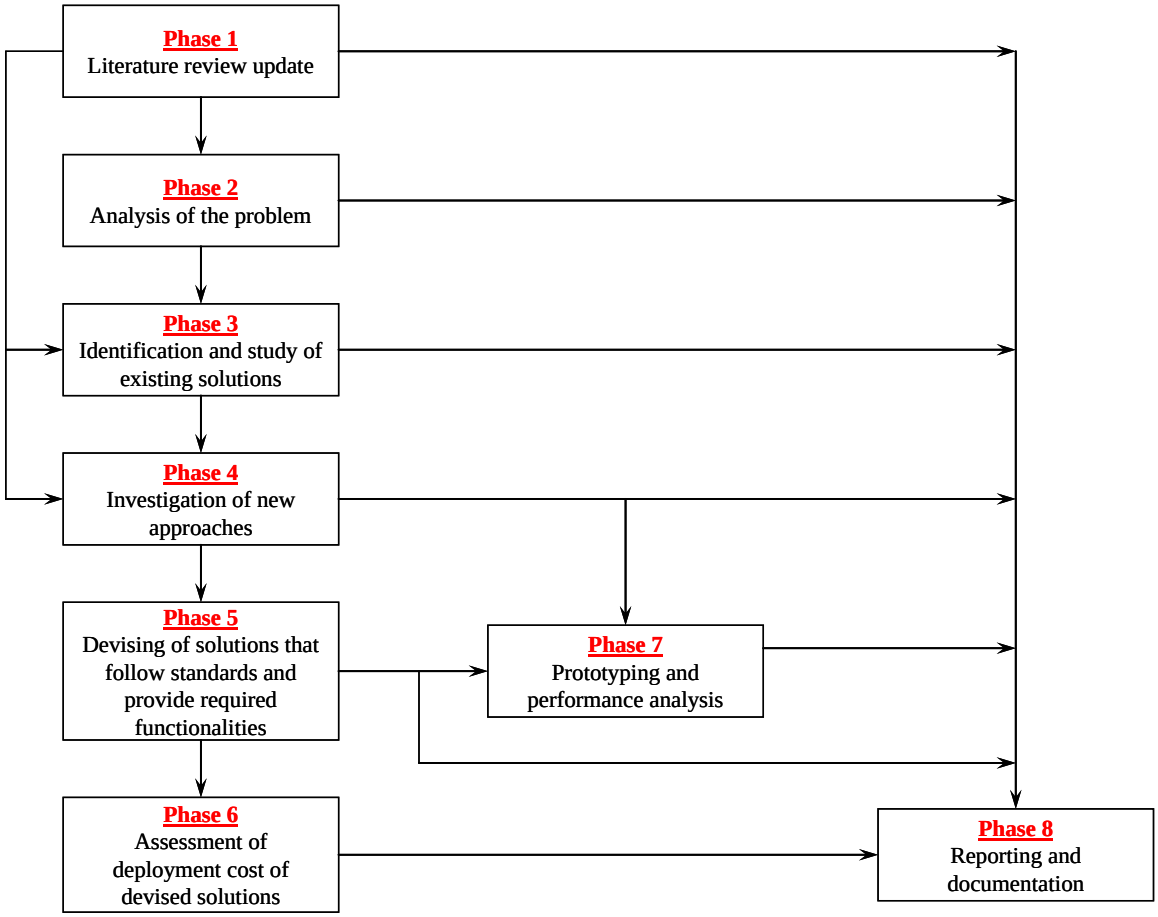

Figure 3: Block diagram of the research methodology.

5.0 RESULTS AND DISCUSSION

As stated earlier, the project devised solutions for the two types of Internet access denial. Specifically, one solution was developed to solve the application level Internet access denial problem. Likewise, three different solutions were proposed to bypass the routing level Internet access denial problem. Consequently, we present an explanation of the devised solutions and the associated results.

P2P Solution to Application Level Internet Access Denial Problem

The proposed solution uses both the Chord P2P protocol [56] and the round-robin approach. The Chord protocol arranges the peers (nodes) in a ring, and constructs a *finger table* that is made of a number of selective nodes so as to minimize the number of hops to the next responsible resolver node. Accordingly, the proposed solution requires each Domain Name System (DNS) resolver to construct a list of unique nodes from the finger table. Such a list is referred to as a *routing table*. In addition to using the normal DNS lookup procedure (i.e., root/TLD DNS servers), the routing table will be used to forward the query to one node in a round-robin fashion. When the query is forwarded to the next hop node, it will resolve the query through the existing DNS lookup procedure. If the next hop node is BLOCKED (i.e., the node is being intentionally denied DNS service), the query will not be resolved. To keep the routing table updated with ACTIVE nodes (i.e., nodes that can resolve queries through higher nameservers) and minimize the number of unresolved queries, there are two modifications to the normal round-robin approach; using Chord as underlying layer, and using an enhanced round-robin. Chord protocol is used as an underlay infrastructure to solve the scalability issue in the round-robin approach by updating (i.e., stabilizing) the finger table in the Chord protocol on a periodic basis. Also, the stabilization process is modified to support the scenario of a BLOCKED node. The other enhancement is made to the round-robin technique such that it has to go through the entire routing table if there are FAILED nodes, which guarantees a successful lookup.

Moreover, to enhance the performance of the devised solution we note that the Chord P2P protocol constructs a *Successor List* that is used to provide better lookup and high stability when there are failures in the network. Moreover, the use of *virtual nodes* was proposed in [56] to increase the fairness between the nodes. Also, the Chord P2P protocol uses two main procedures, *Find Successor* and *Closest Predecessor*, to submit a query to nodes. Accordingly, the devised solution made three modifications to the Chord P2P protocol introduced in [56].

The first modification introduced a new condition to the Find Successor procedure that checks if the entry in the Successor List is BLOCKED or DEAD. Since the BLOCKED node cannot resolve any query, it will be flagged in the Find Successor procedure. However, the BLOCKED node will not be removed from the Successor List so as to help in the Closest Predecessor procedure. Note that the BLOCKED node does not cause a timeout since it continues to participate in the stabilization process.

The second modification extended the condition in the Closest Predecessor procedure used by the Chord P2P protocol that checks if the entry is ALIVE or not to consider the BLOCKED node as an ALIVE node. Each entry in the finger table or the Successor List will be checked to see if the node is either ACTIVE or BLOCKED, and if not then the node will be removed from the routing table.

The third modification ensured that the feature of notifying the predecessor and successor when a node voluntary departs that was suggested in [56] will be used also when a node is BLOCKED.

To enhance the performance of the devised solution further, the devised solution enhanced the round-robin performance by checking a window of the routing table in each round. This window could be one entry resulting in the basic round-robin, portion of the routing table, or the complete routing table. This enhancement is applied only when the first entry in the window is DEAD or BLOCKED. Using the lookup window, the lookup procedure will start with the first entry in the window. If this entry is either DEAD or BLOCKED, the algorithm will remove that entry from the routing table. In addition, a timeout is generated when the entry is DEAD. If all entries in the window have been removed, the query message will be dropped. On the other hand, the message will be considered successfully resolved if there is at least one ACTIVE entry in the window. A pointer is used in order to help in determining the window's starting point in the next round. The pointer is incremented in a round-robin fashion after each check.

The proposed P2P solution was evaluated and compared against the modified Chord P2P protocol by means of simulations. The simulation environment was built using NetBeans IDE 6.8 with Java 1.6.0. The Chord behavior in the simulator was verified first by reproducing the results in [56]. In general, the simulator had the same behavior as the Chord protocol presented in [56].

The following describes the analysis metrics used for comparison between the modified Chord P2P protocol and the proposed round-robin P2P approach. Note that the unit of the analysis metric, if any, appears between parentheses after the name of the analysis metric.

1. *Load (keys/node)* – The average number of keys per node that only appear in the lookup (query) messages.
2. *Traffic (message/node)* – The average number of forwarded query messages per node.
3. *Fairness* – Reflects how uniformly the query resolving is distributed among the nodes (i.e., *load fairness*). Also, fairness is used to identify the uniformity of the number of messages forwarded by each node (i.e., *traffic fairness*). Note that the messages in the *load fairness* case refer to the resolved queries that only appear in the lookup messages, and in the *traffic fairness* case refer to any forwarded queries.
4. *Percentage of Failed Lookup* – Percentage of dropped queries out of the total queries sent.
5. *Message Path Length (hops)* – The number of hops needed to identify the resolver.
6. *Timeout (seconds)* – The average timeout when a node does not respond (i.e., failed or departed) in a successful lookup.

Two scenarios were simulated and analyzed; *Blocking* and *Failure*. The blocking scenario considers simulating the situation when a number of nodes become BLOCKED, whereas the failure scenario simulates when a number of nodes FAIL. Each scenario consists of two network configurations. The two configurations examined are:

1. *Chord-20:* it is the basis of the other configuration and is used to update the routing information of the other configuration. The Successor List size in Chord-20 is $2\times (\log_2 N)$, where N is the total number of virtual nodes in the network. So, with 1,000 nodes and one virtual node per real node, the size equals 20. Chord-20 was picked for implementation similar to what is used in [56].
2. *Round-10:* change the lookup process from Chord to round-robin with the lookup window size set to the size of the complete routing table. Thus, for a network with 1,000 nodes, the window size = $(\log_2 1{,}000) = 10$.

Table 2 summarizes the common parameters and the corresponding values that are used. Most of the values used are derived from [56].

Table 2: Default parameters

Parameter	*Value*
Iterations	10 iterations per sample
Initial Network Size	1,000 real nodes
Virtual Nodes	4 virtual nodes/real node
Successor List Size	$2\times (\log_2 N)$
Lookup Window Size	full routing table size
Query Message	key & live source are chosen randomly
Number of Queries	10,000 lookup queries
Query Inter-arrival Time	Exponential with mean of 50 ms
Timeout	500 ms
Stabilization Period	Uniform [15, 45] seconds

In Figures 1 to 5, after a network of 1,000 real nodes is stable, the state of 50% of the nodes is changed from ACTIVE to BLOCKED (i.e., blocking scenario), and from ACTIVE to FAILED (i.e., failure scenario). Then, uniform random queries are generated using the default parameters.

Figure 4 shows the lookup failure for the blocking and the failure scenarios. In case of the blocking scenario, there is no lookup failure when using Chord-20. That is because the blocked node will notify its neighbor. Accordingly, the neighbor updates its Successor List. In contrast, Round-10 depends on the routing table that is only updated during the stabilization process (i.e., finger table updating). Hence, it incurs a negligible lookup failure (about 0.5%). For the failure scenario, we note that the Chord protocol is stable whenever the Successor List is accurate. However, when failed nodes are 50%, the network stability decreases and the dropped messages increase. In contrast, Round-10 suffers from a minor lookup drop similar to its behavior in the blocking scenario.

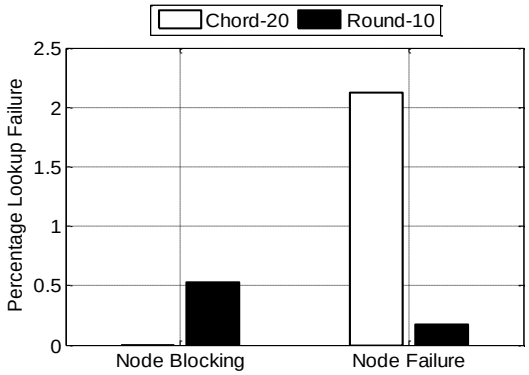

Figure 4: Lookup failure.

The load and the load fairness results are shown in Figure 5(a) and Figure 5(b), respectively. For both the blocking and the failure scenarios, each non-BLOCKED node receives an equal load (i.e., keys) for both Chord-20 and Round-10. Thus, both Chord-20 and Round-10 distribute queries equally among the remaining nodes. In contrast, it is observed that for both

cases of the blocking and the failure scenarios, Round-10 provides higher fairness between the nodes than Chord-20 due to its natural behavior.

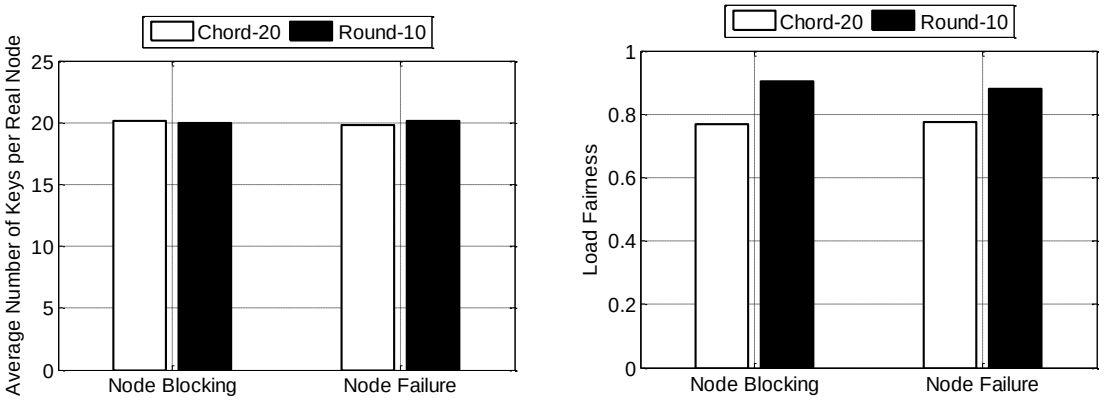

Figure 5: (a) Load, and (b) load fairness.

For the blocking scenario, we note that Round-10 does not forward the queries to intermediate nodes, so there is no extra hop to resolve the queries as noticed from Figure 6. In Chord-20, the path length follows the equation $Path\ Length = (log_2\ M)/2 - (log_2\ r)/2 + 1$ provided in [56], where M = network size, and r = virtual nodes per real node. In the node failure scenario, Chord-20 experiences an increase in the path length because of the failed nodes. This is because the node attempts to submit the query to the failing node, which results in a timeout. Then the node will utilize the closest predecessor procedure and that results in an additional submission of the same query, resulting in an increase in the path length. On the other hand, there are no hops in resolving queries in the Round-10.

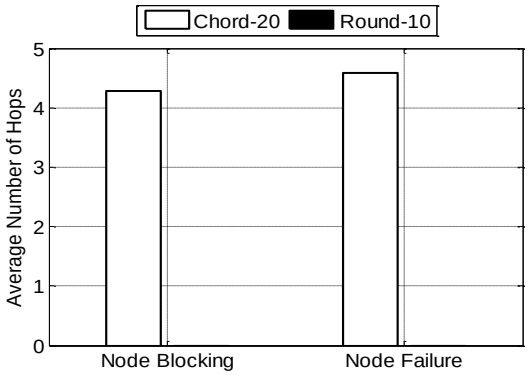

Figure 6: Message path length.

In the case of the blocking scenario, it was stated earlier that the BLOCKED nodes are still considered as live nodes, and that results in no timeout as illustrated in Figure 7. Whereas in the failure scenario, a failing node obviously does not notify others about its condition, then it is possible that some of the nodes along the path to the responsible resolver node are DEAD. Hence, Chord-20 experiences longer timeouts than Round-10 as shown in Figure 7.

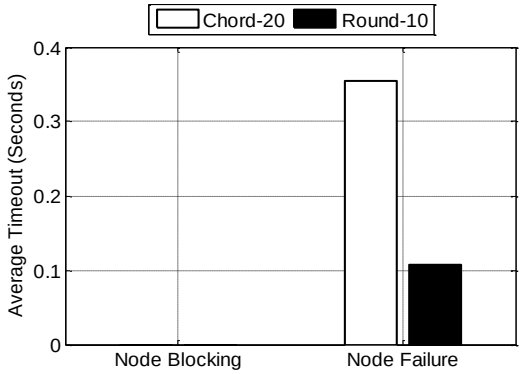

Figure 7: Message timeout.

Figure 8(a) and Figure 8(b) illustrate the traffic and traffic fairness, respectively. For the blocking scenario, the average number of traffic increases due to the additional traffic introduced by the stabilization process. The difference between Chord-20 and Round-10 is due to the fact that Round-10 does not forward messages during the lookup process. For the failure scenario, a slight increase in the average traffic between the nodes is observed when compared with the average traffic in the blocking scenario presented in Figure 8(a). This occurs in the failing scenario because the only ALIVE nodes participating in the system are not BLOCKED. In contrast, in the blocking scenario, BLOCKED nodes are part of the ALIVE nodes but do not participate in forwarding messages. Also, as shown in Figure 8(b), the drop in traffic fairness between the nodes, especially for Chord-20, is a result of updating the finger table during the stabilization process. Accordingly, when the stabilization process is completed, there will be no BLOCKED nodes in the table. That causes the BLOCKED nodes to only generate queries, and to stop helping in forwarding the messages. As a result, the traffic fairness among all ALIVE nodes is affected.

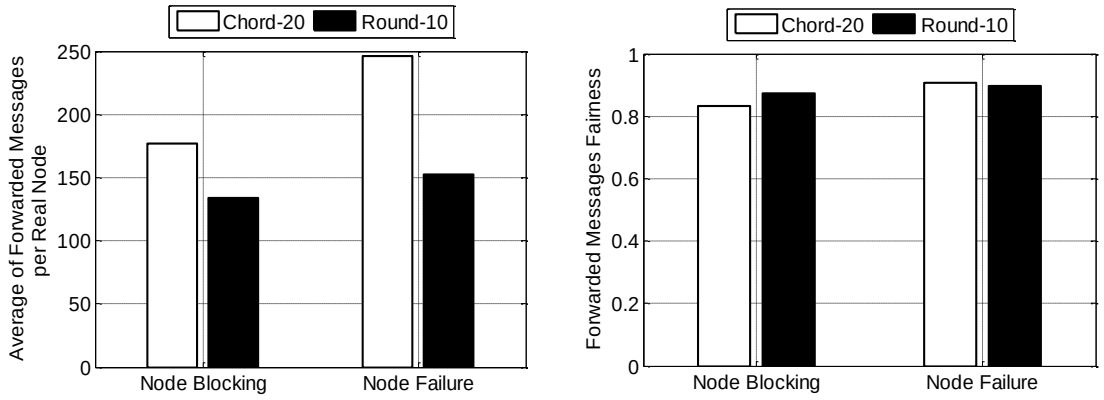

Figure 8: (a) Traffic, and (b) traffic fairness.

The details of the simulation model, the simulation code, and the simulation results can be found in [57].

In summary, the P2P solution provides an alternative path to resolve the query which increases the availability of the DNS to any specific region. The proposed solution minimizes the effects of the intentional blocking while being scalable and resilient.

BGP Tuning Solution to Routing Level Internet Access Denial Problem

BGP Tuning refers to the use of the available BGP policies in order to modify the BGP selection process to enforce the selection of the intended route as the best one. In order to make sure that the traffic passes through the non-malicious IISP (i.e., good IISP), the outgoing traffic needs to be directed through the non-malicious IISP and the incoming traffic needs to be influenced to select the non-malicious IISP as the best path in its BGP selection process. Therefore, we considered ways of controlling the outgoing and incoming traffic to direct routes through the non-malicious IISP.

Controlling the outgoing traffic is easier than controlling the incoming traffic. This is because it is easier to configure the local router to prefer a route than to affect the selection process of all other routers in the Internet that are outside the control of the local ISP. To control the outgoing traffic, the administrator of the Regional ISP (RISP) can set higher `local-pref` attributes of the routes learned from the good IISP. This assignment will ensure that all destinations reachable through the good IISP will go through it. However, if there is a destination that is only reachable through the malicious IISP, the traffic will go through it. Subsequently, it will be blackholed and will not reach the destination. Figure 9 depicts the use of the `local-pref` attribute to control the outgoing traffic.

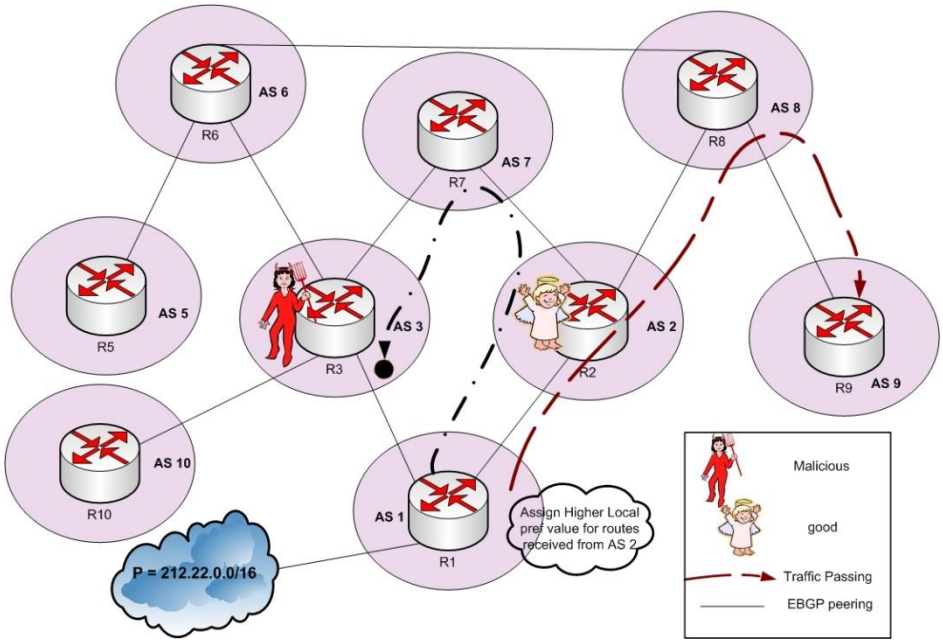

Figure 9: Control of outgoing traffic using BGP tuning.

As shown in Figure 9, the traffic will go to all reachable destinations through the good IISP. For example, the traffic destined from AS1 to AS9 will pass through AS2 and AS8 and it will reach the destination through the good IISP. However, if we assume that router R1 in AS1 wants to send traffic to router R10 in AS10. The only way to reach the destination is through router R3. In this case, the traffic will be filtered out and the destination will not be reachable. Using any identity hiding technique can help to solve this specific scenario.

On the other hand, three methods to control incoming traffic were investigated. These methods are *AS-Path shortening*, *more specific announcement*, and the *use of communities*. To achieve better control of the incoming traffic, the three methods can be combined and used together.

The first method, *AS-Path shortening*, utilizes the fact that in the selection process, when comparing two routes, if an AS Path of one route is shorter than the other, it will become more preferred. AS Path prepending has been widely used to make the path to a specific destination less preferred [58]. If the good IISP sends an announcement of the local blocked region prefixes directly without adding the AS number of the local region in the AS-path, the length of its AS-path will be reduced by one. In this way, the incoming traffic can be influenced to come through the good IISP. In Figure 10, AS1, which wants to control its traffic, will not have its AS number included in the AS-path advertised by AS2. In this way, AS7, for example, will prefer the route that comes from AS2 because it has a shorter AS Path to the prefix. However, if an AS has a local preference that leads to preferring routes to the prefix through the malicious router, then shortening the route will not influence the traffic. This is because `local-pref` is considered first in the selection process of BGP. In addition, if the only way to reach a destination is through the malicious router, then the routes will be blocked. All these scenarios are shown in Figure 10.

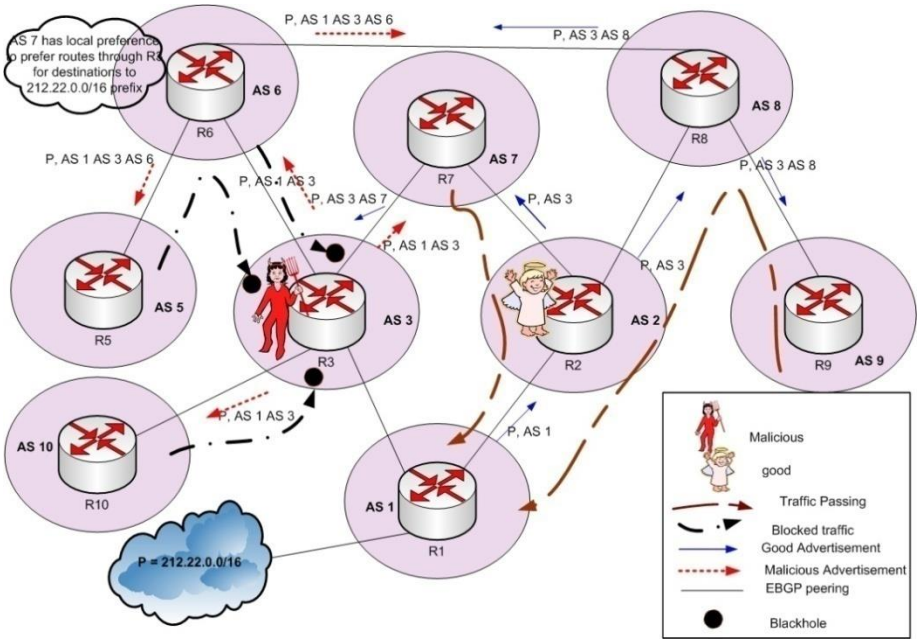

Figure 10: AS Path shortening.

As shown in Figure 10, the traffic will be influenced to go through the shorter path. For example, router R9 in AS9 will select to go through the path AS8 AS2 AS1. Therefore, it passes through the good IISP, i.e., AS2. In case the path has a local preference to go through the malicious IISP, then advertising a shorter path will not help. For example, if R6 in AS6 decides to send traffic through the malicious IISP AS3 because of local preference, then the traffic will be blackholed. Also, if the only path to a destination must pass through the malicious IISP as the case for router R10 in AS10, then the traffic will be blackholed when destined to the prefix of AS1. This is because the traffic must pass through the malicious IISP AS3.

The second method, *more specific announcement*, is based on the fact that when forwarding traffic, the destination of the traffic is matched to the longest match of the prefixes in the routing table. In this way, advertising a more specific prefix through the good IISP makes all the routers select the good IISP to reach the destination. Figure 11 depicts this technique.

As shown in Figure 11, even if a path has a higher `local-pref` for routes that are destined to a certain prefix, a router selects the path with the longest prefix match even before it performs the selection process. For example, AS6 has a higher local preference to routes received from the malicious IISP AS3. However, because more specific prefixes are advertised to AS6 from AS8, the traffic selects the path AS6 AS8 AS2 AS1. Therefore, the traffic will pass through the good IISP. Obviously, if the only provider for an AS is the malicious IISP, then the route must go through it and it will be blackholed like the case for router R10 in AS10.

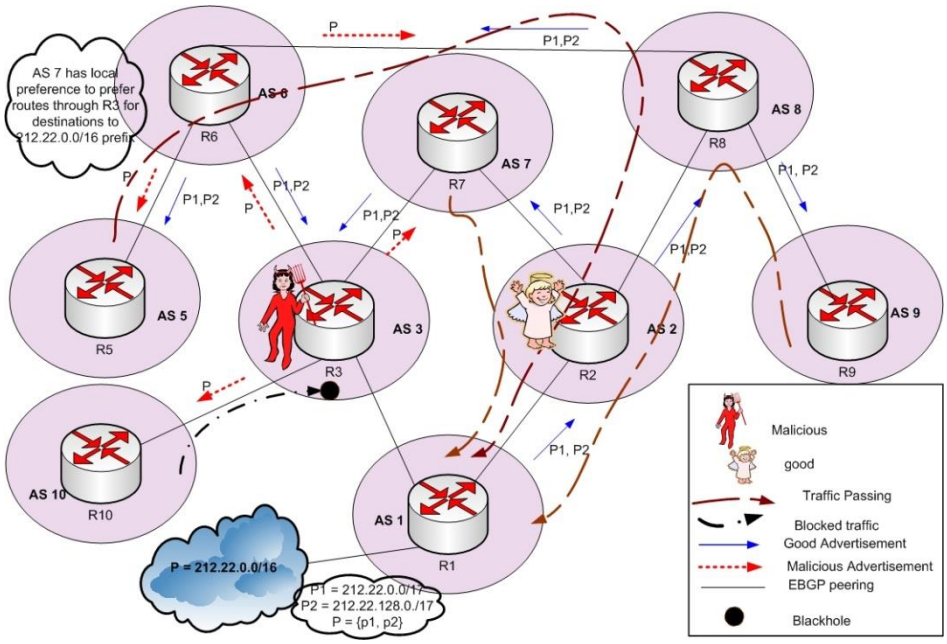

Figure 11: More specific prefixes announcement.

The third method to direct the incoming traffic through the good IISP is through the *use of communities*. An AS can advertise a path and assign a certain community number to it. A route map condition can then be set in some of the ASes in between to assign a higher local preference for routes with the community number assigned to the advertised path. As a result, these ASes will prefer the paths through the good IISP. Figure 12 shows the use of community to control the traffic.

As shown in Figure 12, cooperation with some of the ASes in the Internet, belonging to the same community, may be needed in order to prefer the routes with a specific community. Because the routes advertised from router R2 in AS2 belong to a certain community number, AS7 will set a higher `local-pref` for paths learned through AS2. So, the traffic will go through AS2, the good IISP. If some ASes select to direct the traffic through the malicious AS because it is the only way to reach the destination, then the traffic will be blocked since it will face a blackhole as in the case of router R10 in AS10.

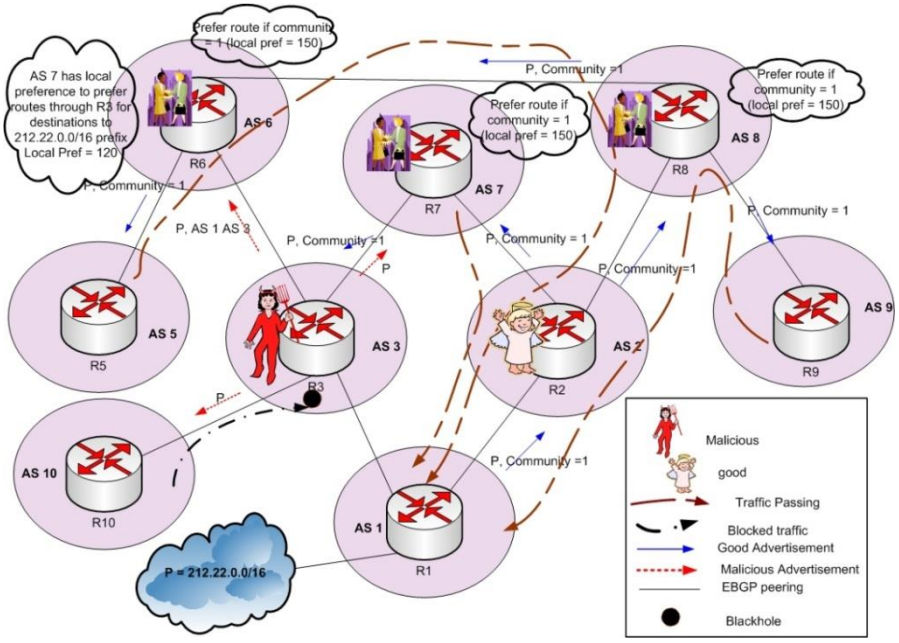

Figure 12: Control the traffic through the use of communities.

Note that with each of the three BGP tuning methods, the network performance is not affected once the method is activated and the alternative BGP path that passes through the good IISP is chosen. Hence, to compare between the three BGP tuning methods, we provide a qualitative comparison amongst the methods in terms of the following criteria: traffic filtering, setup overhead, communication overhead, difficulty to combat the method, and scalability. Each of these criteria is discussed in the next few paragraphs. Table 3 provides a summary of the qualitative comparison.

Traffic filtering refers to the amount of the traffic that is expected to be filtered out, i.e., blackholed, using the methods listed in the columns. The AS-Path shortening has a higher traffic filtering than more specific prefixes and the use of communities as traffic will be filtered for the cases when routers have a local preference that prefers the malicious route, routers that can still see the path through the malicious route as a shorter path depending on their location, and routers that have the malicious router being the only provider.

The second criteria in table 1 is the **setup overhead** which is a measure of the time needed and the difficulty level faced in order to get all the required configurations performed to execute the method. The setup overhead is higher for the use of communities method as it requires more coordination with other ASes than the other methods to make it functional.

In contrast, **communication overhead** refers to the number of messages that need to be exchanged between routers before the method is effective. There are no additional exchanged messages between routers other than those that are part of normal BGP conversations.

The difficulty to combat the method is a measure of the amount of effort required by the malicious IISP to overcome the solution. One obvious disadvantage of BGP tuning based techniques is that the malicious IISP can easily mimic the tuning implemented by the local region to neutralize all the benefits gained. For example, it can shorten the AS-Path to be more preferred in the selection process. It can also advertise more specific prefixes so it can gain advantage of the longest prefix match. Moreover, it can advertise routes with the same community number advertised by the good IISP to eliminate the advantage of the use of community for the local region.

Finally, in terms of **scalability** which refers to the easiness of extending the method or using it for the entire Internet, all BGP tuning techniques provide high scalability since the configuration will affect the decision taken by the traffic in the Internet without any additional connections.

Table 3: Comparison between BGP tuning methods

	BGP Tuning		
	AS-Path Shortening	**More Specific Prefixes**	**Communities**
Filtering the traffic	Medium	Small	Small
Setup overhead	Small	Small	Medium
Communication overhead	No	No	No
Difficulty to combat the method	Easy	Easy	Easy
Scalability	High	High	High

Another way to compare between the three BGP tuning methods is to consider the BGP convergence time required by each method for the new path to be setup through the good ISP. Such a comparison was examined by conducting simulations using the OPNET network simulator [55]. Figure 13 depicts the experiment setup implemented in OPNET. Assume the region of interest, designated by AS12, is connected to the Internet through two IISPs: AS3 representing the main IISP, and AS4, representing the secondary IISP. The IISPs AS3 and AS4 provide connectivity for the region of interest as well as to other ASes such as AS6, AS100, and AS7. The main objective is to evaluate how long it takes for the traffic to switch to the secondary IISP (i.e., good IISP) after the main IISP AS3 starts to drop or blackhole traffic coming to or going out of AS12. The traffic sessions are between AS12 and AS7 with the intermediate AS100 modeling the Internet cloud. For this setup, we considered the convergence time of three types of traffic: hypertext transfer protocol (HTTP), file transfer protocol (FTP), and Voice over IP (VOIP). Note that HTTP and FTP run over TCP, whereas VOIP runs over UDP. Finally, the designed experiment accounted for the average Internet delay and the background load through specifying input parameters for AS100 and the links connecting the various autonomous systems.

The intended Internet outage occurs at a specified time instant during the simulation where the behavior of Router 3 at the core of the main IISP network, as shown in Figure 13, is modified to drop traffic originating or destined to AS12, while it continues to advertise reachability to the region of interest AS12. Following the outage occurrence, the edge router of AS12, Router 2, starts implementing the specified BGP-based solution. The OPNET model allows the testing of the following solutions: AS-Path shortening, more specific prefixes, and the use of communities, one at a time. The internal code implementing BGP for the involved routers is modified to implement these solutions and respond to the Internet outage event as per the specified solution. The details of the simulation model and programming code are elaborated on in [50].

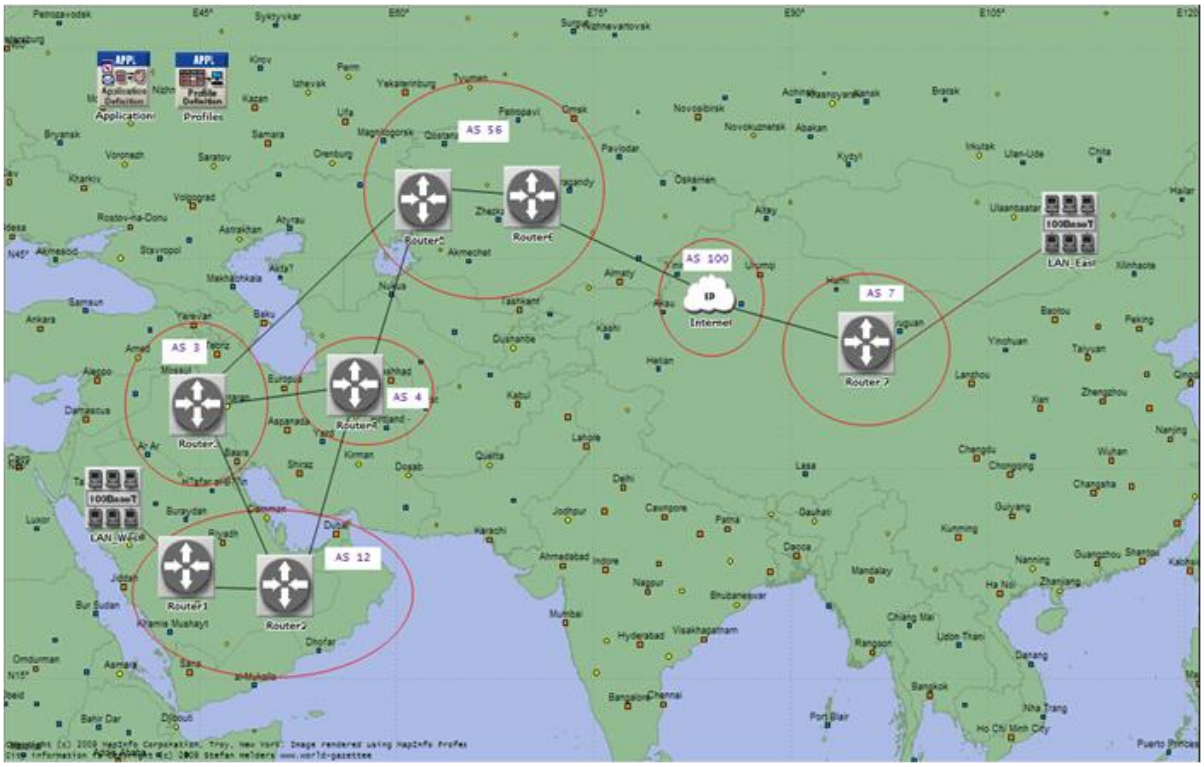

Figure 13: OPNET simulation model.

Note that the convergence time includes the duration it takes the BGP messages and updates to propagate through the Internet and the routers of concern to update their tables with the required entries. Figure 14 shows the convergence time for the case where the Internet average delay is equal to 100 milliseconds for three cases of background loads: 20%, 50%, and 80%. The graph also shows the 95% confidence interval for each convergence figure.

It can be seen that the bulk of the convergence times are in the order of 1 second with the BGP solution based on more specific prefixes having the least convergence times. This is true for background loads and for the three types of traffic. This is because the solution based on more specific prefixes requires the update of nominally one prefix, the one belonging to AS 7 (or LAN_East in Figure 13). However, for the case of path shortening and the use of communities, BGP has to search through and update a longer list of paths for the new path (reachability through AS2) to take place. The figure also shows that the convergence time increases slightly with the increase in the Internet background load. Furthermore, as the VOIP sessions used in the simulation have an average throughput that is greater than that for the corresponding HTTP and FTP sessions, they present more load to the network and consequently lead to longer BGP convergence times. The same network setup is also tested for an average Internet delay of 5 seconds, as opposed to the 100 milliseconds assumed earlier, and the convergence times were ranging from 15 to 50 seconds. The convergence times patterns observed for the 5 second Internet delay are consistent with those for the 100 millisecond scenario. Similar to the previous case of 100 millisecond average Internet delay, the more specific prefix based solutions have lower convergence times compared to those for the other two solutions, and the VOIP traffic leads to the longest convergence time relative to HTTP and FTP traffic. Finally, again the convergence time increases slightly with the increase of the Internet background load.

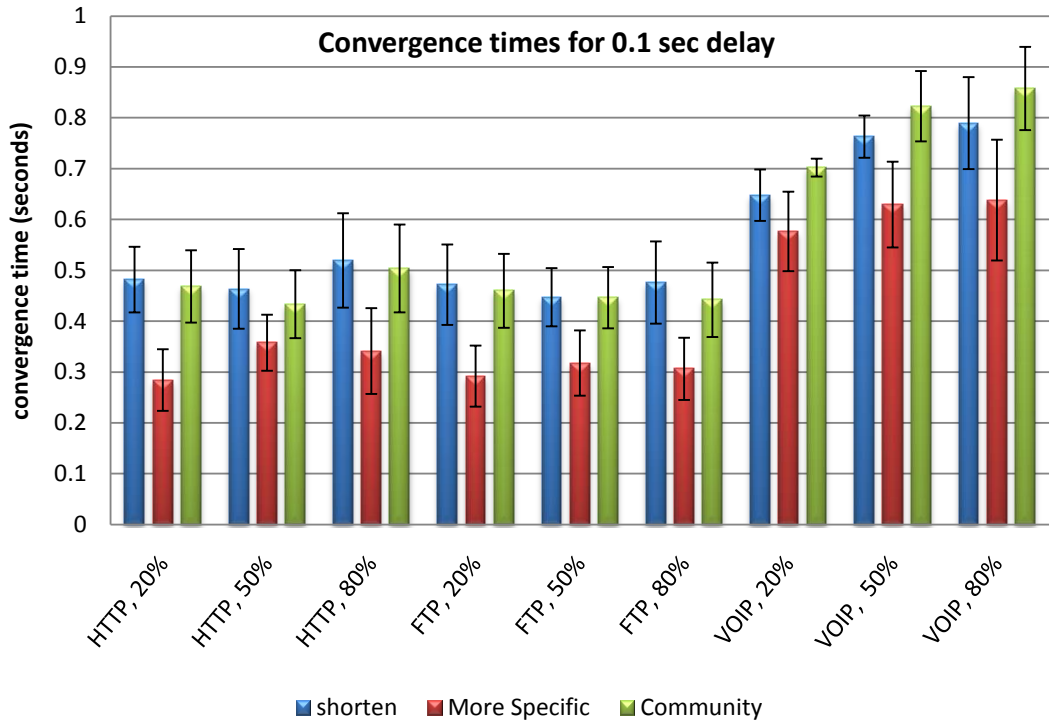

Figure 14: Convergence time for BGP-based solutions with 0.1 s average Internet delay.

Moreover, the results of the simulations were confirmed by conducting lab experiments using the testbed shown in Figure 15. The laboratory set up contained seven Cisco 2811 routers, four Catalyst 2950 switches, one workstation, and three servers. The three servers were set up as they would be on the Internet side (AS600), and the workstation as it would be on the local side (AS100). Also, each server was assigned to one of the Internet applications; FTP, HTTP, and VoIP. Furthermore, each server and the workstation were equipped with WireShark [51] network analyzer to collect the statistics of each test. The routers of the testbed were interconnected using 1.544 Mbps DS-1 links. The testing procedure was used with three different traffic loads; 384 Kbps (about 25% of DS-1 link capacity), 768 Kbps (about 50% of DS-1 link capacity), and 1.28 Mbps (about 80% of DS-1 link capacity). Using the different traffic loads, the **convergence time** of each BGP-based solution was collected.

Four Java network programs were also developed to automate the testing environment. The first and main program, referred to as *detector*, is capable of checking the Internet connectivity and is installed on the workstation that is connected to the local side (AS100). When the *detector* program detects a sequence of timeout messages, it can immediately and remotely login to the local side (AS100) BGP speaker and configure it with the BGP-based solution to be tested. At a pre-determined time for turning on the malicious activity a second program, referred to as *malicious*, connects to the malicious ISP (AS300) BGP speaker to configure it to act maliciously by setting its Access Control List (ACL) to block the outgoing and incoming traffic of the local side (AS100). The third and the fourth programs, referred to as *reset*, were designed to reset the configurations of the testbed to the initial state so a new testing can be conducted.

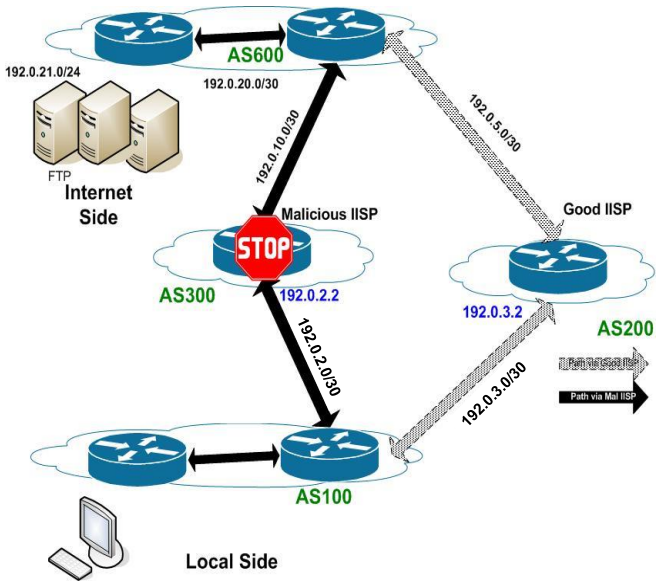

Figure 15: BGP-based solution laboratory testbed setup

To achieve the behavior of a malicious ISP, the *malicious* program was used to configure the AS300 BGP speaker router's ACL to permit the exchange of only BGP messages and advertisements to and from the local side (AS100) while blocking the rest of the traffic from and to the local side (AS100). Two ACL lists were implemented in the malicious ISP BGP speaker to block traffic belonging to the local side (AS100), one blocking the outgoing traffic and another blocking the incoming traffic. The first ACL was implemented on the closest interface to the local side (192.0.2.2). On the other hand, the second ACL was implemented on the interface that is closest to the Internet side (192.0.10.1).

To verify that the malicious behavior is taking place, the *PING* command was used. Figure 16 shows how the local side BGP speaker cannot reach the server at the Internet side (AS600) after running the *malicious* program. Note that although an alternative physical path that did not pass through the malicious ISP existed, the local side BGP speaker continued to forward the outgoing traffic via the malicious ISP which in turn dropped it. The *traceroute* results demonstrated that the local BGP speaker still preferred the path that passed through the malicious ISP (192.0.2.2). After implementing one of the evaluated BGP-based solutions, the local side can ping the Internet side and the *traceroute* results show that the packets were going through the good ISP (192.0.3.2).

The testing procedure commences by having the client residing at the local side (AS100) initiating communication with a server that resides at the Internet side (AS600), and running the *detector* program. At a pre-determined time for turning on the malicious activity the *malicious* program connects to the malicious ISP router (AS300), and configures its ACL with the blocking configuration. When the *detector* program detects an Internet connectivity loss, it immediately connects to the local side BGP speaker (AS100), and configures it with one of the BGP solutions. For the solution to take effect, the *detector* program resets the BGP and the routing tables of the local side BGP speaker. The time between detecting the malicious action and recovering from it is measured by the *detector* program. Once the testing is concluded the *reset* programs are executed to put the testbed back to its initial state. This procedure was repeated 10 times and the average results were considered.

```
Over malicious ISP path
C:\Users\lab>PING 192.0.21.6
Pinging 192.0.21.6 with 32 bytes of data:
Reply from 192.0.2.2: Destination net unreachable.
Reply from 192.0.2.2: Destination net unreachable.

C:\Users\lab>tracert 192.0.21.6
Tracing route to 192.0.21.6 over a maximum of 30 hops
  1    <1 ms    <1 ms    <1 ms  192.0.1.1
  2    192.0.2.2  reports: Destination net unreachable. ← malicious ISP
Trace complete.

Over alternate path after implementing one of the solutions
C:\Users\lab>ping 192.0.21.6 ← Server in AS600
Pinging 192.0.21.6 with 32 bytes of data:
Reply from 192.0.21.6: bytes=32 time=23ms TTL=124
Reply from 192.0.21.6: bytes=32 time=23ms TTL=124

C:\Users\lab>TRACERT 192.0.21.6
Tracing route to 192.0.21.6 over a maximum of 30 hops:
  1    <1 ms    <1 ms    <1 ms  192.0.1.1
  2     6 ms     5 ms     5 ms  192.0.7.1
  3    11 ms    11 ms    11 ms  192.0.3.2 ← good ISP
  4    19 ms    19 ms    18 ms  192.0.4.2
  5    27 ms    27 ms    27 ms  192.0.5.1
  6    33 ms    32 ms    32 ms  192.0.21.6
Trace complete.
```

Figure 16: Verification of the malicious activity and the BGP-based solution

The convergence time testing procedure is summarized in the following steps:

1. Choose one of the BGP-based solutions and configure the *detector* program accordingly.
2. Run the *detector* program so as to maintain a sequence of PING commands with the application server at the Internet side (AS600).
3. At a pre-determined time for turning on the malicious activity run the *malicious* program to configure the malicious ISP router's ACL so as to block traffic from and to the local side (AS100).
4. When the *detector* program receives a sequence of failed replies, it configures the local side (AS100) router with the chosen BGP-based solution and records the time afterwards. Subsequently, the *detector* program resumes the sequence of PING commands to the same application server, and records the time when it gets a successful reply from the application server.
5. Run the *reset* programs to put back the testbed into its initial state including clearing the BGP and the routing tables from all routers.
6. Wait 60 seconds that are needed to ensure that the BGP and the routing tables of all routers are cleared out.
7. Repeat steps 1 through 6 for 10 times.

The procedure was repeated for each of the BGP-based solutions and the corresponding convergence time was collected. Accordingly, the results of the lab experiments confirmed the results of the simulations provided in Figure 14.

Tunnel-based Solution to Routing Level Internet Access Denial Problem

The proposed tunnel-based solution to the Internet access denial utilizes available tunneling protocols such as IP-in-IP and GRE. A tunnel is created from the local AS to a destination AS only if the normal path to the destination AS passes through the malicious ISP. For the proper establishment of the tunnel, the solution assumes the presence of at least one cooperating AS that is placed before the malicious ISP on the tunnel path, and at least another cooperating AS that is placed after the malicious ISP on the tunnel path. As a result of creating the tunnel, the malicious ISP can be bypassed.

To illustrate the concept further, we note that Quoitin in [37] proposed the use of *virtual peering* to deterministically control the incoming traffic through one of the providers. Virtual peering was used to achieve load balancing of incoming traffic among providers and to reduce the latency by choosing paths that have the lowest delay. The virtual peering method can be used to automate the establishment of tunnels as it automates the setting up and the removal of the virtual peering (i.e., tunnels) by using Virtual Peering Controller (VPC). Figure 17 depicts the use of virtual peering to solve the malicious ISP blocking.

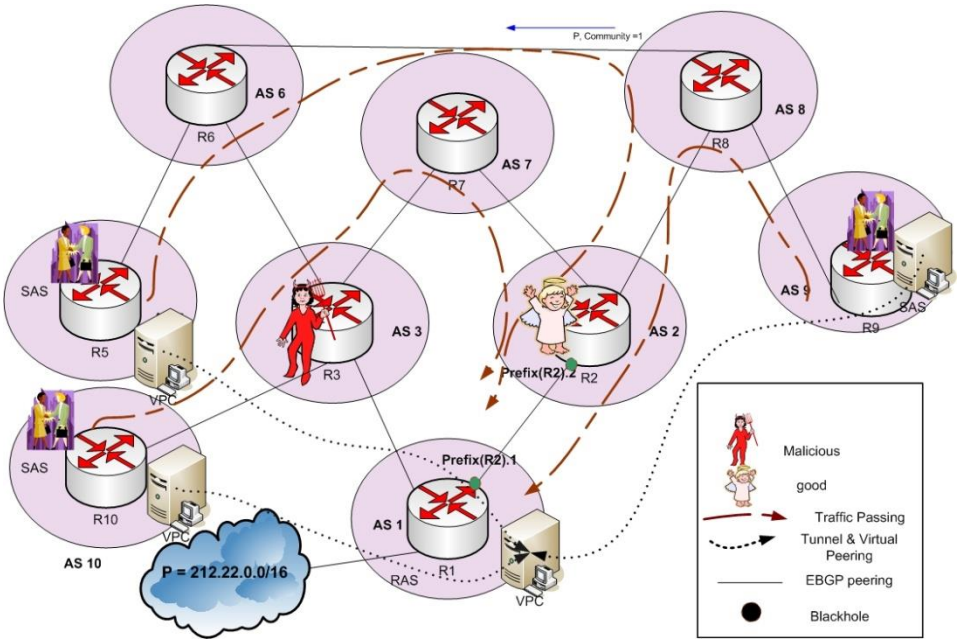

Figure 17: Virtual peering-based solution.

As shown in Figure 17, the Requestor Autonomous System (RAS), which is the destination ISP, requests the establishment and removal of virtual peering connections. The Source Autonomous System (SAS) originates the traffic destined to RAS. The VPC in RAS needs to know the IP address of the VPC in SAS. This can be done either manually or automatically by including the IP address of the VPC in the BGP update message as an extended community attribute. When the RAS knows the IP address, it can establish a multi-hop BGP session with the VPC in the SAS. It can then send a Virtual Peering Establishment (VPE) or a Virtual Peering Removal (VPR) to the SAS. After that, one of the border routers in the SAS establishes a tunnel with one of the border routers of the RAS so that the IP address used as destination address belongs to a prefix owned by the provider AS, which is the good IISP for this case. Knowing that the ISP assigns one of its IP addresses to the connection between the border routers, this IP address is used as a destination to the tunnel. Hence, the malicious router will not be able to figure out that this IP address belongs to the blocked prefix. The

router can then decapsulate the traffic and send it to the destination of interest. In order to force routes to go through the tunnel, a higher `local-pref` is assigned to them. One thing to note is that even if the traffic passes through the malicious ISP, the traffic will not be blocked because it will be destined to an IP not belonging to the prefix being blackholed.

Setting up the multi-hop BGP connection should be established such that the traffic of BGP messages does not pass through the malicious ISP. If they pass through the malicious ISP, a BGP connection will not be established, since the BGP messages will be blackholed. Therefore, the choice of the IP address of the VPC in a virtual peering must be the same as the one used to set up the tunnel. Therefore, virtual peering shall be adapted such that the VPC is the same router that is used for establishing the tunnel to force traffic to pass through the good ISP.

In virtual peering, every source of traffic must do virtual peering in order to control all traffic coming from or going to the Internet. This is not always very practical, and makes the solution not scalable. Therefore, the project introduced a number of modifications to the virtual peering method to make the solution scalable. We refer to the modified solution by a *virtual transit* solution. The main difference is that the transit router advertises the prefix to other routers in the Internet. Another difference is that the AS-Path can be set as a path of length two: the transit AS and the originating (or provider) AS. A number of virtual transit routers may be distributed in the Internet attracting the traffic destined to the prefix of interest and then passing it to the destination in a tunnel established in the same way it is done with virtual peering. The concept of virtual transit is illustrated in Figure 18.

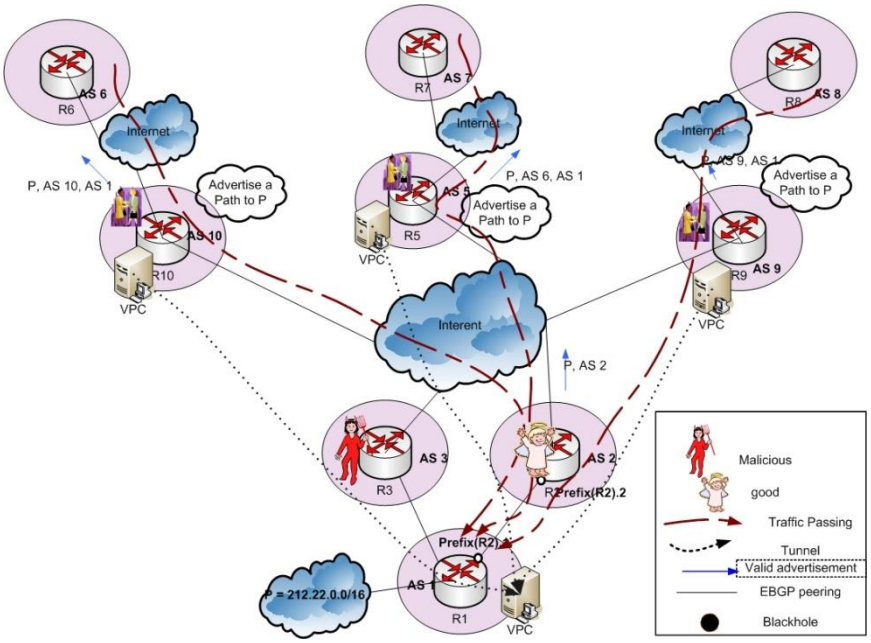

Figure 18: Virtual transit scheme.

As shown in Figure 18, the VPC in AS10 is a cooperative router. This router establishes a virtual transit connection with the VPC in AS1 in the same way virtual peering is established. However, it advertises to other ASes in the Internet that it has a path to R1, and in order to influence more traffic, it shortens the AS-Path to two. The philosophy behind this is that since a tunnel is established between AS10 and AS1, then the path will consider only these ASes in its advertisement. It may also advertise more specific prefixes in order to attract even more traffic.

The basic proposed tunnel-based solution (i.e., without virtual transit) was validated through simulations using the OPNET network simulator [55]. In the simulation setup, the typical network shown in Figure 2 was used. In addition, a tunnel between R2 (from the blocked AS) and R5 (from the distant AS) was created. The created tunnel passes through R3 and R4. The non-blocked IP address provided by the neighboring AS (i.e., AS200) was used to create the tunnel. Thus, with the help of a neighboring AS, a tunnel that passes through the malicious ISP (i.e., AS300) was created. The use of a non-blocked IP address will prevent the malicious router (i.e., router R4) from dropping incoming and outgoing traffic to and from the affected AS.

To create a tunnel, we need a prefix to be used for the tunnel interface. In the simulation, the chosen prefix belongs to subnet 200.0.0.0/24 (i.e., AS200). The tunnel starting point IP address is 200.0.0.1, and the tunnel ending point IP address is 200.0.0.2. The routing protocol used for the tunnel interface is OSPF.

To validate that the proposed solution is setup to forward the traffic properly through the tunnel, we first examined the IP forwarding table on both routers R2 and R5. From the tables we can determine that the incoming and the outgoing traffic on router R2 and router R5 use the created tunnel. Furthermore, Figure 19(a) and Figure 19(b) show the IP tunnel traffic received, in bits per second (bps), at router R2, and the IP tunnel traffic sent, in bps, by router R2, respectively. In contrast, Figure 19(c) and Figure 19(d) show the IP tunnel traffic received, in bps, at router R5, and the IP tunnel traffic sent, in bps, by router R5, respectively. This validates the proper setup and operation of the tunnel, and the proposed solution.

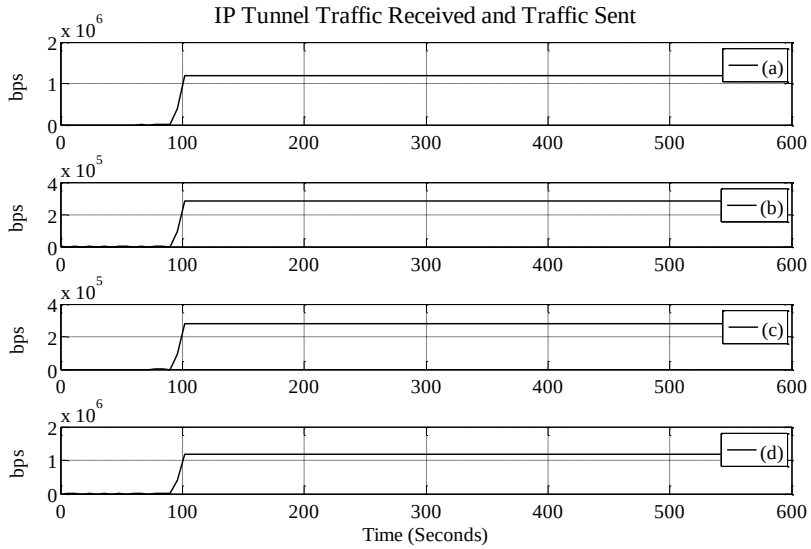

Figure 19: IP tunnel traffic received and sent by routers R2 and R5.

When tunneling protocols are used, extra overhead bytes are added to the packets entering the tunnel as compared to normal packets. The amount of extra overhead bytes that will be added depends on the type of the tunneling protocol used. The tunneling protocols IP-in-IP, GRE, and GRE with check sum add 20, 24, and 28 of extra overhead bytes, respectively, to each packet entering the tunnel. As such, an impact on the network performance is expected as a result of using tunneling protocols in the proposed solution. Thus, the effect of the proposed tunnel-based solution on the network performance was evaluated. Hence, we compared the network performance when the tunnel-based solution is used, to the normal operation (i.e., no malicious activity by the ISP) in terms of end-to-end delay, and traffic throughput overhead

under different types of traffic and with 75% network load. The performance evaluation is conducted by means of OPNET simulations. The network model shown in Figure 2 was used for the performance evaluation. The local and remote local area networks (i.e., LAN_R1 and LAN_R6) were set to 100 Mbps Fast Ethernet networks. The gateway routers were based on the generic router model in OPNET that supports BGP, OSPF, and tunneling. All routers were interconnected to each other as shown in Figure 2 using DS-1 links, providing a data rate of 1.544 Mbps. The total network simulation duration was set to 600 seconds. Moreover, several simulation scenarios are considered by varying the type of tunneling protocol, the type of traffic, and with 75% network load. The tunneling protocols considered in the simulations were IP-in-IP, GRE, and GRE with checksum. Furthermore, each simulation scenario was repeated for 5 different seeds/runs, and the average of the 5 results was reported. The solution was evaluated for different types of traffic generated by well known applications such as file transfer protocol (FTP) and video conferencing. FTP represents a network application that runs over TCP, whereas video conferencing represents a network application that runs over UDP. Each simulation was run with 75% of the available link bandwidth (i.e., 1,158 kbps). The performance statistics collected were the end-to-end delay and the traffic throughput overhead. The end-to-end delay, measured in seconds, was collected at LAN_R1, whereas the throughput, measured in bits per second, was collected at the link between R5 to R4. The R5 to R4 link was selected because the tunnel overhead can be examined at this link with respect to each tunneling protocol.

The end-to-end delay refers to the duration of time that a packet takes to travel from the client to the server. The end-to-end delay includes the transmission time, the propagation time, and the queuing delay.

FTP was simulated as requests to download a file from the server in LAN_R6. The results for the relative increase in the end-to-end delay, which is computed as $(Delay_{Tunnel} - Delay_{NoTunnel})/Delay_{NoTunnel}$, are shown in Figure 20.

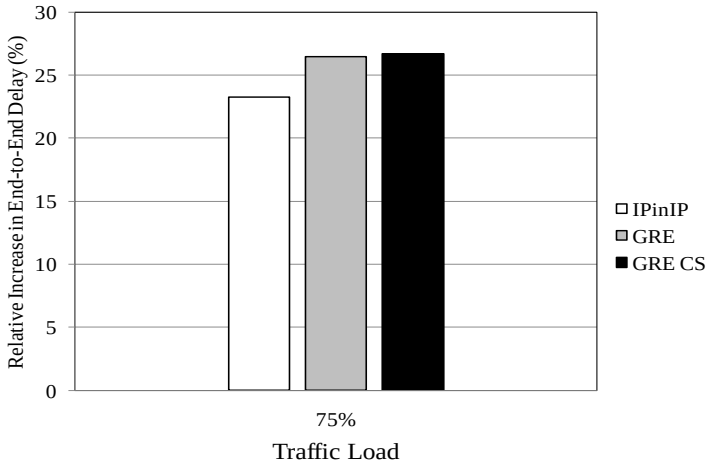

Figure 20: FTP relative increase in end-to-end delay.

The relative increase in the end-to-end delay is caused by the introduction of the tunnel. Furthermore, we note that the IP-in-IP tunneling protocol has the least relative increase in the end-to-end delay. The observation is justified by noting that IP-in-IP tunneling protocol adds the least amount of overhead among the tunneling protocols considered, and therefore, produces the least amount of fragmentation. Although the increase in the relative end-to-end delay is around 25%, the absolute increase in the end-to-end delay shown in Figure 21, and

computed as ($Delay_{Tunnel}$) − ($Delay_{NoTunnel}$), is less than 1.5 ms, which is considered to be negligible for FTP.

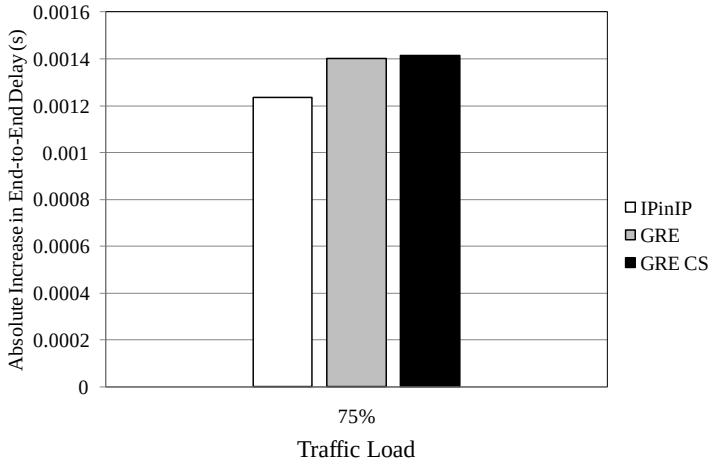

Figure 21: FTP absolute increase in end-to-end delay.

For the video conferencing scenario, the end-to-end delay results show similar behavior to the results obtained for FTP. The relative increase in the end-to-end delay is shown in Figure 22, and as evident it shows a lower relative increase in the end-to-end delay than the results for FTP. This is mainly due to the fact that video conferencing runs over UDP which has a considerably smaller header than the TCP header, over which FTP runs. Also, UDP is a connectionless protocol that does not wait for acknowledgements as in the case of TCP.

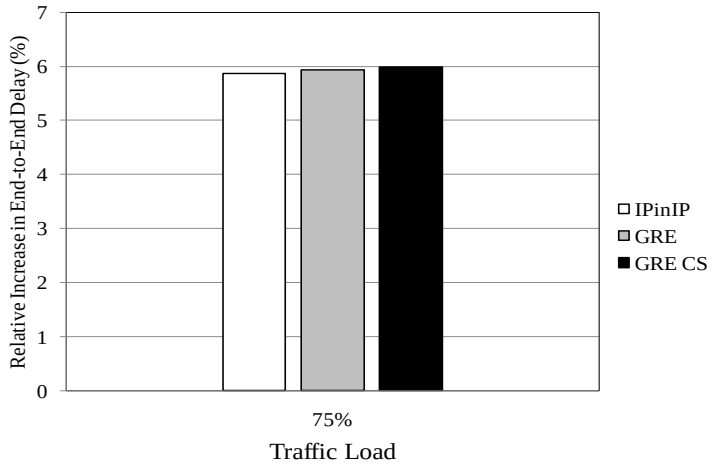

Figure 22: Video conferencing relative increase in end-to-end delay.

Figure 23 shows the absolute amount of increase in the end-to-end delay caused by the introduction of the tunnel. Similar to FTP, we see that the IP-in-IP tunneling protocol has the least amount of increase in the end-to-end delay. Although the increase in the relative end-to-end delay is around 6%, the absolute increase in the end-to-end delay shown in Figure 23 is less than 7 ms, which is considered to be insignificant for video conferencing.

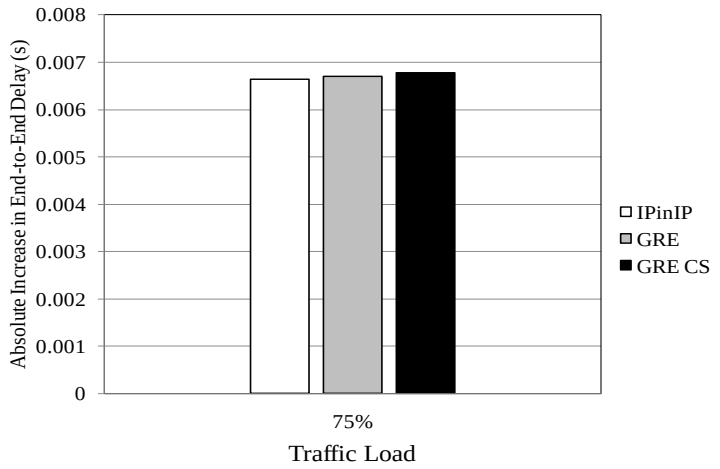

Figure 23: Video conferencing absolute increase in end-to-end delay.

The other performance measure that was investigated was the traffic throughput overhead, in bits per seconds, which measures the amount of overhead bits added to each packet entering the tunnel. The simulation was set to measure the traffic throughput overhead at the link R5 to R4 where the tunnel starts.

The results for both the FTP and the video conferencing applications are shown in Figure 24 and Figure 25. The relative increase in the throughput overhead is shown in Figure 24, whereas the absolute overhead of the tunnel is shown in Figure 25.

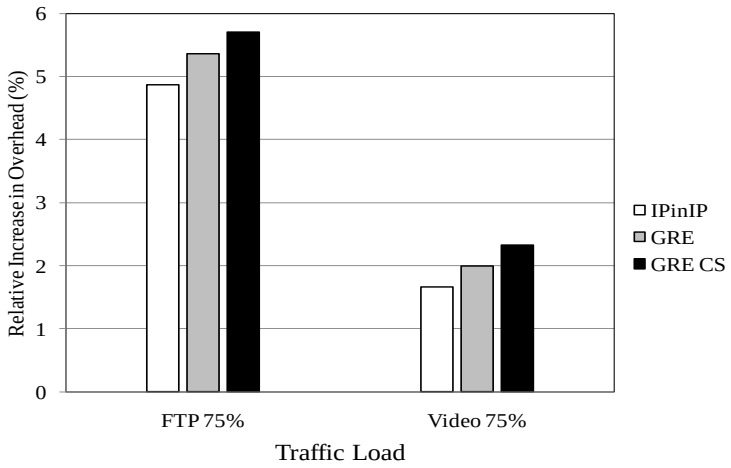

Figure 24: Relative increase in throughput overhead.

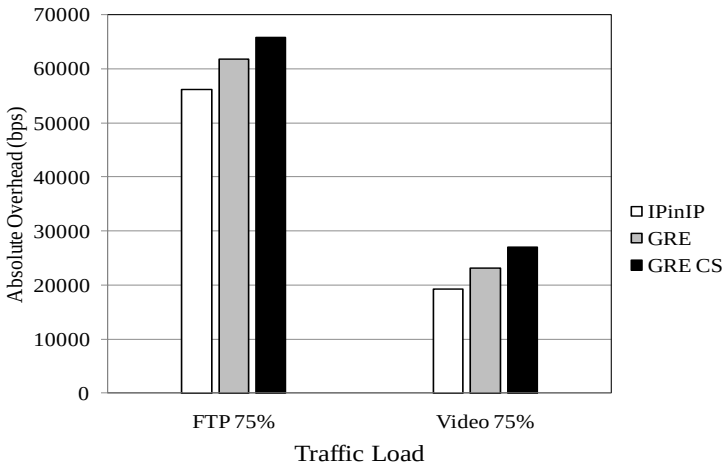

Figure 25: Absolute throughput overhead.

It is obvious from Figure 24 that the relative increase in the traffic throughput overhead for video conferencing is lower than the relative increase in the throughput for FTP for the same reason stated earlier for the end-to-end delay. Likewise, It can be observed from Figure 24 that the IP-in-IP tunneling protocol has the least amount of relative increase in traffic throughput overhead as it adds the smallest header size among the other tunneling protocols. Furthermore, the amount of absolute traffic throughput overhead that is caused by the introduction of the tunnel is shown in Figure 25. In both the FTP and the video conferencing scenarios, the amount of absolute traffic throughput overhead is considered to be negligible relative to the overall throughput of about 1,158 kbps.

In conclusion, UDP-based traffic has the least amount of relative increase in the tunnel overhead when compared to the TCP-based traffic. To further make the comparison fair, an experiment was conducted with a small file size for the FTP application. The results of the experiment confirmed that the UDP-based traffic incurs less amount of tunnel overhead than the TCP-based traffic. As stated earlier, this is mostly attributed to the differences between UDP and TCP in header size and connection type.

The details of the simulation model, and the simulation results can be found in [60]. Moreover, the results of the simulations were confirmed by conducting lab experiments using the testbed shown in Figure 26. The laboratory set up contained six Cisco 2811 routers, four Catalyst 2950 switches, one workstation, and three servers. The blocked local side was represented by AS100, while AS300 represented the malicious ISP. On the other hand, AS400 represented the remote cooperative AS, and AS200 represented the neighboring cooperative AS. The three servers were set up in LAN_R4 as they would be on the Internet side (AS400), and the workstation as it would be on the local side (AS100). Also, each server was assigned to one of the Internet applications; FTP, HTTP, and VoIP. Furthermore, each server and the workstation were equipped with WireShark [51] network analyzer to collect the statistics of each test. In addition, Iperf [59] was used to generate real-time traffic. The routers were interconnected using 1.544 Mbps DS-1 links. The testing procedure was used with three different traffic loads; 384 Kbps (about 25% of DS-1 link capacity), 768 Kbps (about 50% of DS-1 link capacity), and 1.158 Mbps (about 75% of DS-1 link capacity). Using the different traffic loads, the **end-to-end delay**, **traffic throughput**, and **packet drop rate** of each tunnel-based solution were collected.

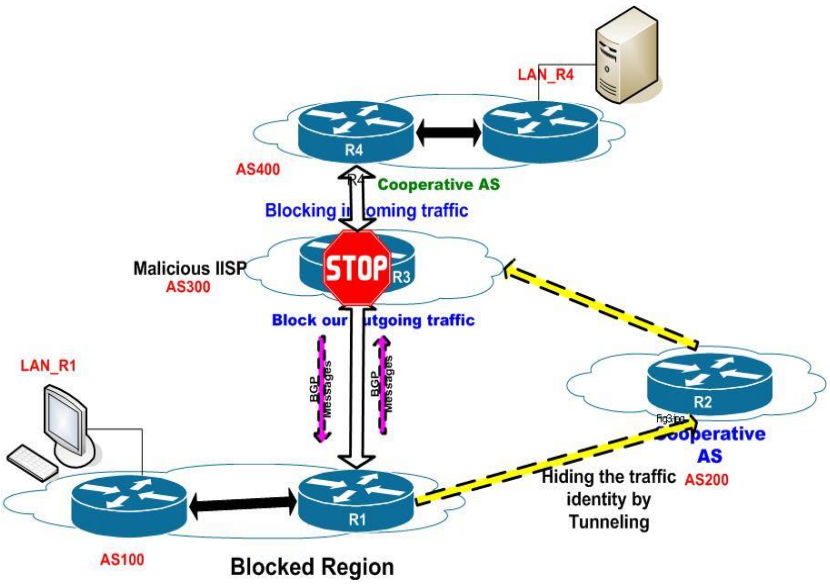

Figure 26: Tunnel-based solution laboratory testbed setup

The four Java network programs (*detector*, *malicious*, and *reset*) that were developed to automate the testing of the BGP-based solution were modified and used to automate the testing of the tunnel-based solution. After turning on the malicious activity, a tunnel would be created between the blocked region gateway router (i.e., R1 in AS100) and the remote cooperative AS router (i.e., R4 in AS400). Subsequently, the incoming and outgoing traffic would be forwarded through the cooperative AS router (i.e., R2 in AS200).

Similar to the testing for the BGP-based solution, the testing procedure of the tunnel-based solution commences by having the client residing at the local side (AS100) initiating communication with a server that resides at the Internet side (AS400) using Iperf. Accordingly, the **baseline** end-to-end delay, traffic throughput, and packet drop rate are collected using WireShark. At a pre-determined time for turning on the malicious activity the *malicious* program connects to the malicious ISP router (AS300), and configures its ACL with the blocking configuration. The *detector* program is then run to connect to the local side BGP speaker (AS100), and to configure it with one of the three tunneling protocols used in the solution. Note that prior to the beginning of the testing of the solution the remote cooperative AS is configured with the pre-selected tunneling protocol. Subsequently, the **solution** end-to-end delay, traffic throughput, and packet drop rate are collected again using WireShark. Once the testing is concluded the *reset* programs are executed to put the testbed back in its initial state. This procedure was repeated 10 times and the average results were considered. The testing procedure is summarized in the following steps:

1. Choose one of the tunnel-based solutions and configure the *detector* program accordingly.
2. Run the Ipref to start generating traffic from the local side (AS100) to the application server at the Internet side (AS400).
3. Collect the **baseline** end-to-end delay, traffic throughput, and packet drop rate using WireShark.
4. Run the *malicious* program to configure the malicious ISP router's ACL so as to block traffic from and to the local side (AS100).
5. Run the *detector* program to connect to the local side BGP speaker (AS100), and configure it with the selected tunneling protocol used in the solution.

6. Collect the **solution** end-to-end delay, traffic throughput, and packet drop rate again using WireShark.
7. Run the *reset* programs to put back the testbed into its initial state including removing the tunnel, and clearing the BGP and the routing tables from all routers.
8. Repeat steps 1 through 7 for 10 times.

The procedure was repeated for each of the tunnel-based solutions and the corresponding end-to-end delay, traffic throughput, and packet drop rate statistics were collected. Accordingly, the results of the lab experiments confirmed the results of the simulations provided in Figure 20 through Figure 25. Hence, through simulations and lab experiments, it was shown that the tunnel-based solution does not have any significant impact on the performance of the network. Thus, deploying the tunnel-based solution as a scalable Internet access denial solution would operate without any performance impacts.

NAT-based Solution to Routing Level Internet Access Denial Problem

The proposed solution is based on NAT as it provides a level of security for the private network by hiding its internal addressing structure and topology. Hence, NAT is 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. 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.

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. For a small network, a single NAT router with an external IP address is used. 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. To resolve these issues, load-balancing can be used by adding more NAT routers at the gateway level as shown in Figure 27. Each NAT router handles a portion of the private network, and has its own pool of IP addresses.

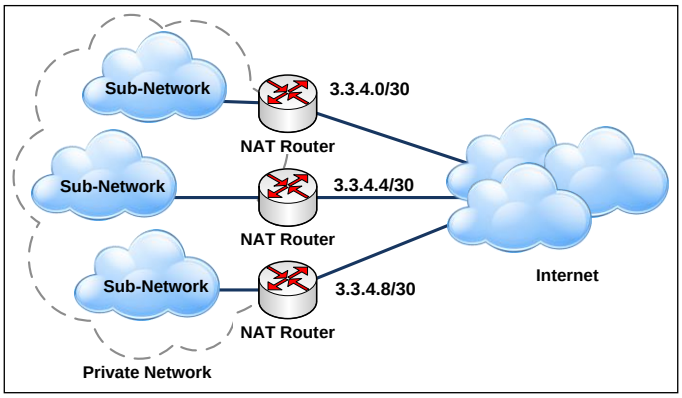

Figure 27: Extended NAT design using load-balancing over a number of NAT routers.

In order to evaluate the impact of implementing the proposed NAT solution on the network, simulations are performed using the OPNET network simulator [55]. The objective of the simulations is to compare the network performance before and after implementing the NAT solution. The range of the simulated NAT delay values is between 10μs and 250μs. In reality, the range for real routers is between 10μs and 50μs. The remaining range 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.

The simulated network shown in Figure 28 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.

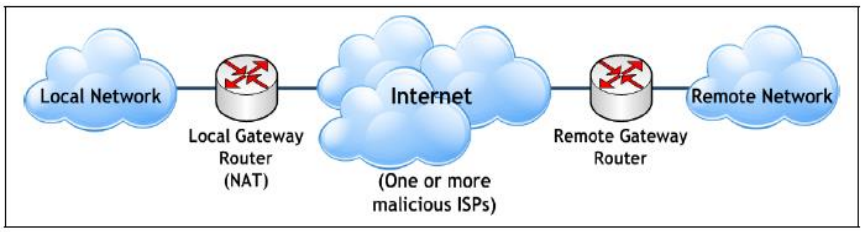

Figure 28: Simulated scenario to measure the effect of NAT delay on network performance.

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.

Each simulation measures the end-to-end delay which 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 29. 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 packets suffer to reach the destination increases linearly.

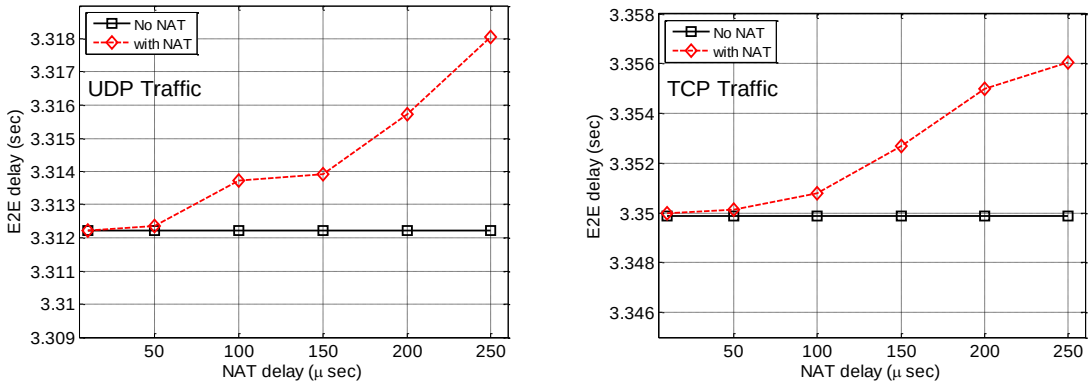

Figure 29: End-to-end delay for high UDP and TCP traffic.

The relative increase of the end-to-end delay for UDP and TCP traffic is shown in Figure 30. The relative increase is computed as $(Delay_{NAT} - Delay_{NoNAT})/Delay_{NoNAT}$. It can be seen 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).

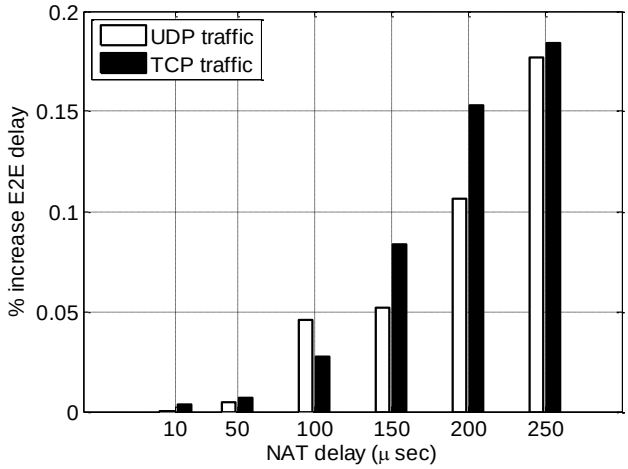

Figure 30: Relative increase of end-to-end delay for high UDP and TCP traffic.

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 was 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 31 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.

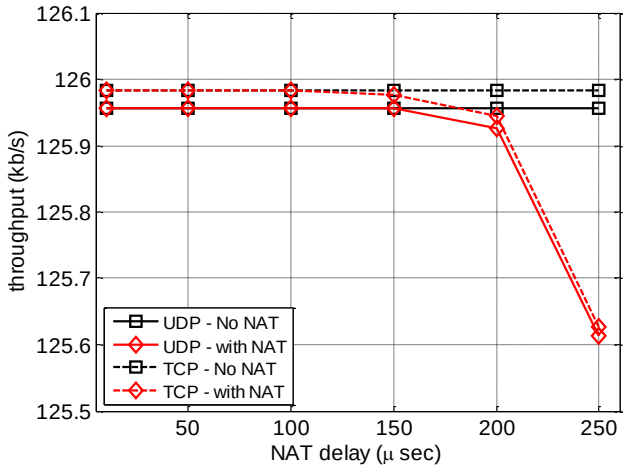

Figure 31: Throughput of high UDP and TCP traffic.

The relative decrease of throughput, which is computed as ($Throughput_{NoNAT} - Throughput_{NAT}$)/$Throughput_{NoNAT}$, is shown in Figure 32. 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.

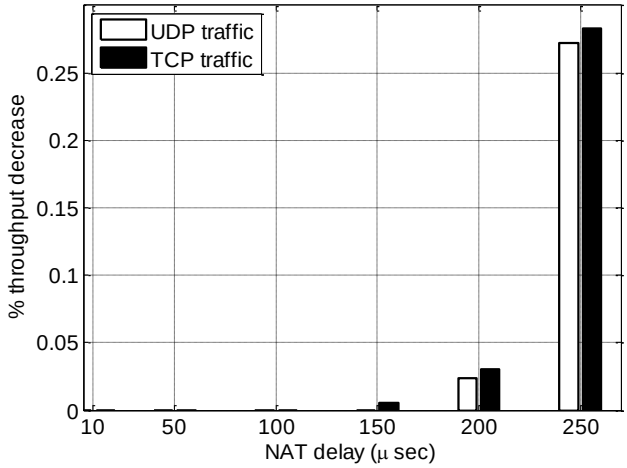

Figure 32: Relative decrease of throughput for high TCP and UDP traffic.

As for the drop rate, the simulation results show insignificant increase in the drop rate and, accordingly, are not shown for brevity.

The details of the simulation model, and the simulation results can be found in [61].

Furthermore, the results of the simulations were confirmed by conducting lab experiments using the testbed shown in Figure 33. The laboratory set up contained six Cisco 2811 routers, four Catalyst 2950 switches, one workstation, and three servers. The blocked local side was represented by AS100, while AS300 represented the malicious ISP. On the other hand, AS400 represented the remote cooperative AS, and AS200 represented the neighboring cooperative AS. The three servers were set up in LAN_R4 as they would be on the Internet side (AS400), and the workstation as it would be on the local side (AS100). Also, each server was assigned to one of the Internet applications; FTP, HTTP, and VoIP. Furthermore, each server and the workstation were equipped with WireShark [51] network analyzer to collect the statistics of each test. In addition, Iperf [59] was used to generate real-time traffic. The routers were interconnected using 1.544 Mbps DS-1 links. The testing procedure was used with three different traffic loads; 384 Kbps (about 25% of DS-1 link capacity), 768 Kbps (about 50% of DS-1 link capacity), and 1.158 Mbps (about 75% of DS-1 link capacity). Using the different traffic loads, the **end-to-end delay**, **traffic throughput**, and **packet drop rate** of the NAT-based solution were collected.

The four Java network programs (*detector*, *malicious*, and *reset*) that were developed to automate the testing of the BGP-based solution were modified and used to automate the testing of the NAT-based solution. After turning on the malicious activity, the NAT feature would be turned on the gateway router of the blocked region (i.e., R1 in AS100). Subsequently, the incoming and outgoing traffic would be forwarded over the cooperative AS router (i.e., R2 in AS200).

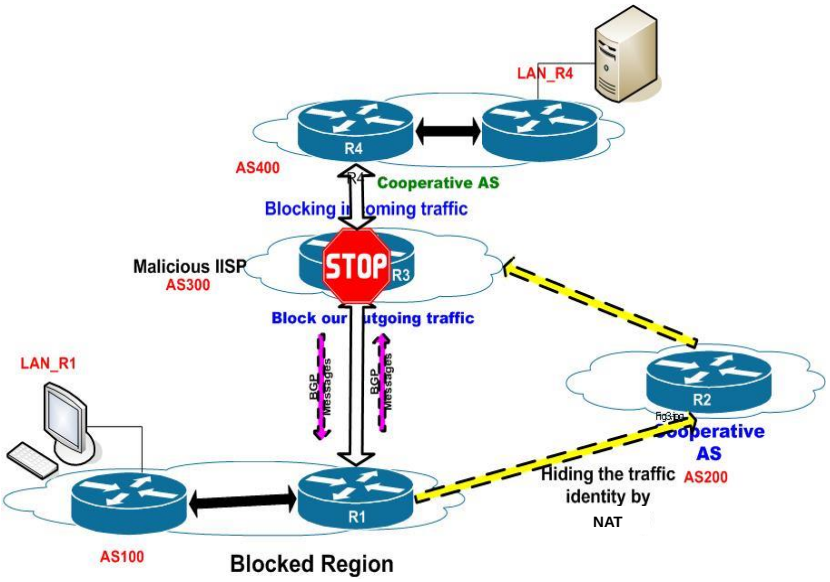

Figure 33: NAT-based solution laboratory testbed setup

Similar to the testing for the tunnel-based solution, the testing procedure of the NAT-based solution commences by having the client residing at the local side (AS100) initiating communication with a server that resides at the Internet side (AS400) using Iperf. Accordingly, the **baseline** end-to-end delay, traffic throughput, and packet drop rate are collected using WireShark. At a pre-determined time for turning on the malicious activity the *malicious* program connects to the malicious ISP router (AS300), and configures its ACL with the blocking configuration. The *detector* program is then run to connect to the local side BGP speaker (AS100), and to configure it with the NAT feature. Subsequently, the **solution** end-to-end delay, traffic throughput, and packet drop rate are collected again using WireShark. Once the testing is concluded the *reset* programs are executed to put the testbed back in its initial state. This procedure was repeated 10 times and the average results were considered. The testing procedure is summarized in the following steps:

1. Run the Ipref to start generating traffic from the local side (AS100) to the application server at the Internet side (AS400).
2. Collect the **baseline** end-to-end delay, traffic throughput, and packet drop rate using WireShark.
3. Run the *malicious* program to configure the malicious ISP router's ACL so as to block traffic from and to the local side (AS100).
4. Run the *detector* program to connect to the local side BGP speaker (AS100), and configure the NAT feature on it.
5. Collect the **solution** end-to-end delay, traffic throughput, and packet drop rate again using WireShark.
6. Run the *reset* programs to put back the testbed into its initial state including removing the NAT feature, and clearing the BGP and the routing tables from all routers.
7. Repeat steps 1 through 6 for 10 times.

The results of the lab experiments confirmed the results of the simulations provided in Figure 29 through Figure 32. Hence, through simulations and lab experiments, it was shown that NAT does not have any significant impact on the performance of the network. Thus, deploying NAT as a scalable Internet access denial solution would operate without any performance impacts.

6.0. CONCLUSIONS

The project has uncovered a new possible malicious Internet activity and proposed a number of solutions. Specifically, the project considered the problem of *Internet access denial* by malicious ISPs at both the application level and the routing level. More precisely, when a higher-level domain name system (DNS) server denies a specific country or region access to DNS services, then that specific country or region will lose access to many of the Internet applications that are highly dependable on DNS services such as HTTP, and an *application level* Internet access denial takes place. On the other hand, when a malicious ISP filters transit traffic for the purpose of dropping packets that belong to a specific country or region, then a *routing level* Internet access denial occurs.

As a result, the project devised solutions for both types of Internet access denial. Specifically, one solution based on the concept of peer-to-peer (P2P) networks was developed to solve the *application level* Internet access denial problem. Through simulations, the developed solution was proven to be highly scalable and robust as a direct result of using a P2P approach for the solution. Hence, the solution is suitable to be deployed in a country or a region. Likewise, three different solutions based on border gateway protocol (BGP) tuning, tunneling protocols, and network address translation (NAT) routers were proposed to bypass the *routing level* Internet access denial problem. The BGP tuning-based solution directs the traffic belonging to a specific country or region around the malicious ISP, and thus protects the traffic from being dropped by that malicious ISP. Alternatively, the tunneling protocol-based solution and the NAT-based solution aim to hide the identity of the traffic belonging to a specific country or region so that the malicious ISP will be misled into routing that traffic normally since its identity is hidden. Through simulations and laboratory experiments, the three developed solutions were proven to be efficient, and to have negligible effect on the existing network performance. Hence, the solutions are suitable to be deployed in a country or a region. In particular, the more specific prefixes method of the BGP-based solutions is suitable when an additional ISP other than the malicious ISP is available. Otherwise, if an additional ISP is not available, the NAT-based solution is more appropriate to be deployed for traffic originating from within the victim region. On the other hand, due to the server reachability problem associated with the NAT routers, the IP-in-IP tunnel-based solution is more suitable to be deployed for traffic originating from outside the victim region. All proposed solutions follow the standards, and hence do not require major modifications to the existing network infrastructure.

Hence, by ensuring better Internet resiliency to KSA, the project significantly contributed to the ***computer systems and networks*** priority technology area of the Information Technology program of KSA's National Science, Technology and Innovation plan. Moreover, the project addressed a new type of malicious Internet activities that has not been addressed previously by other researchers, and opened new research directions. Furthermore, the results of the project allow the Internet network operators and regulators in KSA, in the region, or any other region to be less reliant on international service providers. Likewise, the solutions provided by the project are beneficial to Internet providers and regulators such as King Abdulaziz City for Science and Technology (KACST) and the Communications and Information Technology Commission (CITC), carriers such as Saudi Telecom Company (STC), and local businesses and institutions. Finally, the project produced a specialized team of researchers in this particular field.

7.0 PROJECT OUTCOMES

Outputs	Status	Date
Publications (Journal Papers)		
1. Ashraf Mahmoud, Ahmad Alrefai, Marwan Abu-Amara, Mohammed Sqalli, Farag Azzedin, "Qualitative Analysis of Methods for Circumventing Malicious ISP Blocking," Arabian Journal for Science and Engineering (AJSE).	**Accepted**	**30/1/2011**
2. Marwan Abu-Amara, Abdulaziz Al-Baiz, Ashraf Mahmoud, Mohammed Sqalli, and Farag Azzedin, "A Scalable NAT-Based Solution to Internet Access Denial by Higher-tier ISPs," Journal of Security and Communication Networks, John Wiley.	**Accepted**	**19/9/2011**
3. Marwan Abu-Amara, Mohammed Asif, Mohammed Sqalli, Ashraf Mahmoud, and Farag Azzedin, "A Tunnel-Based Solution to Internet Access Denial by Higher-tier ISPs."	**Under Preparation**	
4. Marwan Abu-Amara, Fahd Abdulhameed, Farag Azzedin, Mohammed Sqalli, and Ashraf Mahmoud, "Dynamic Round-Robin P2P DNS."	**Under Preparation**	
Publications (Conference Papers)		
1. Marwan Abu-Amara, Farag Azzedin, Fahd Abdulhameed, Ashraf Mahmoud, and Mohammed Sqalli, "Dynamic Peer-to-Peer (P2P) Solution to Counter Malicious Higher Domain Name System (DNS) Nameservers," The 24th Canadian Conference on Electrical and Computer Engineering, Niagara Falls, Ontario, Canada, May 8-11, 2011.	**Published**	**8-11/5/2011**
2. Abdulaziz Al-Baiz, Marwan Abu-Amara, Ashraf Mahmoud, Mohammed H. Sqalli, and Farag Azzedin, "Internet Access Denial by Higher-tier ISPs: A NAT-Based Solution," The 24th Canadian Conference on Electrical and Computer Engineering, Niagara Falls, Ontario, Canada, May 8-11, 2011.	**Published**	**8-11/5/2011**
3. Marwan Abu-Amara, Mohammed Asif, Mohammed Sqalli, Ashraf Mahmoud, and Farag Azzedin, "Resilient Internet Access Using Tunnel-Based Solution for Malicious ISP Blocking," The 3rd IEEE International Conference on Communication Software and Networks, Xi'an, China, May 27-29, 2011.	**Published**	**27-29/5/2011**
4. Fahd Abdulhameed, "Dynamic Round-Robin P2P DNS to Improve Internet Access Resiliency at KSA," First Scientific Conference for Graduate and Undergraduate Students, Riyadh, March 2010.	**Published**	**March, 2010**
5. Mohammed Asif, "Resilient Internet Access for KSA using Tunnel-Based Solution," Second Scientific	**Published**	**28-31/3/2011**

Conference for Graduate and Undergraduate Students, Jeddah, Kingdom of Saudi Arabia, 28-31 March, 2011. 6. Amer Al-Ghadhban, Ashraf S. Mahmoud, Marwan Abu-Amara, Farag Azzedin, and Mohammed H. Sqalli, "Prototyping and Evaluating BGP-Based Solutions to Overcome Malicious IISP Blocking," The 2nd International Conference on Networking and Information Technology, Hong Kong, November 25-27, 2011.	**Accepted**	**25-27/11/2011**
Presentations 1. Marwan Abu-Amara, "Dynamic Peer-to-Peer (P2P) Solution to Counter Malicious Higher Domain Name System (DNS) Nameservers," The 24th Canadian Conference on Electrical and Computer Engineering, Niagara Falls, Ontario, Canada, May 8-11, 2011. 2. Marwan Abu-Amara, "Internet Access Denial by Higher-tier ISPs: A NAT-Based Solution," The 24th Canadian Conference on Electrical and Computer Engineering, Niagara Falls, Ontario, Canada, May 8-11, 2011. 3. Fahd Abdulhameed, "Dynamic Round-Robin P2P DNS to Improve Internet Access Resiliency at KSA," First Scientific Conference for Graduate and Undergraduate Students, Riyadh, March 2010. 4. Mohammed Asif, "Resilient Internet Access for KSA using Tunnel-Based Solution". Second Scientific Conference for Graduate and Undergraduate Students, Jeddah, Kingdom of Saudi Arabia, 28-31 March, 2011.	**International** **International** **National** **National**	**10/5/2011** **11/5/2011** **March 2010** **March 2011**
Technical Outputs		
Patents, licenses or other research commercialization activities 1. Marwan Abu-Amara, Fahd Abdulhameed, Farag Azzedin, Ashraf Mahmoud, and Mohammed Sqalli, "Dynamic Round-Robin Peer-to-Peer (P2P) Domain Name System (DNS)," Docket No. 32185.89. 2. Marwan Abu-Amara, Mohammed Asif, Mohammed Sqalli, Ashraf Mahmoud, and Farag Azzedin, "Tunneling-Based Solution to Bypass Internet Access Denial by Higher-tier Internet Service Providers," Docket No. 32185.83. 3. Marwan Abu-Amara, Abdulaziz Al-Baiz, Ashraf Mahmoud, Mohammed Sqalli, and Farag Azzedin, "NAT-based Solution to Internet Access Denial By Higher-Tier ISPs." 4. Ashraf Mahmoud, Marwan Abu-Amara, Ahmad Alrefai, Mohammed Sqalli, and Farag Azzedin, "BGP-	**Submitted** **Submitted** **Under Preparation** **Under Preparation**	**15/6/2011** **29/6/2011**

based Solution to Intentional Blocking by Higher-tier ISPs.”		
Other		

RELATIONSHIP OF THE PROJECT OUTCOMES TO NSTIP STRATEGIC FRAMEWORK

PROJECT OUTCOMES	STRATEGIC TECHNOLOGY PROGRAM GOALS			PROJECT OBJECTIVE ACHIEVED
	GOAL 1 (Computer prototypes, systems, executable product(s), process(es), or procedure(s) useful to the local industry and relevant to the strategic technologies roadmap of IT program tracks)	**GOAL 2** (Experimental setups and equipments contributing to building computer prototypes/systems, executable product(s), process(es), or procedure(s) in IT program tracks)	**GOAL 3** (Experienced teams in the technologies related to the projects)	
1. Contribute to the improvement of Internet access resilience and robustness by providing measures and solutions to the Internet unavailability caused by service providers' denial of Internet access.	**X**			**2, 3, 4, 5, 7**
2. Demonstrate scenarios where such activity, i.e., denial of Internet access, may take place, how disastrous it can be, and how the new proposed solutions may overcome it.		**X**		**1**
3. Show the effectiveness of the proposed measures and solutions.		**X**		**6**
4. Provide practical solutions and prototypes that can be used by organizations to prevent the Internet isolation from occurring, or at least minimize its impact. Such solutions and prototypes can be commercially productized.	**X**			**6**
5. Discover other possible Internet isolation techniques that can be used by service providers to deny access to the Internet. Such discovery can be used to initiate other possible research and product ideas and investigations.			**X**	**1**
6. Develop a team of subject matter experts in the area of Internet access and its impact on the Internet applications and routing.			**X**	**1, 2, 3, 4, 5, 6, 7**

8.0 ADDITIONAL ACHIEVEMENTS

The project resulted in four completed master theses [50][57][60][61], and at least 4 additional ongoing master theses. The completed theses resulted in a number of journal and conference publications as well as patent applications as outlined in section 7.

9.0 VALUE TO THE KINGDOM

The project's proposed solutions assist in ensuring KSA's Internet resiliency against malicious activities, namely the malicious act by ISPs to deny access to the Internet. Hence, the results of the project allow the Internet network operators in KSA, in the region, or any other region to be less reliant on international service providers. Consequently, an Internet embargo against KSA can be thwarted by utilizing the results of the project.

The potential positive impacts on the economy and the society at large are obvious as many businesses and government agencies are highly dependent on the availability of the Internet for conducting their daily affairs. Furthermore, since most of the means of communication such as telephone, mobile and email are merging with the Internet, virtually all people are directly or indirectly affected by the unavailability of the Internet. Therefore, by placing the results of the project in action, the economic and the social impacts on KSA can be severely reduced in case of an Internet embargo against KSA.

The direct utilization of the project's expected outcomes can be summarized in the following:

1. Make Internet network operators in the Kingdom, in the region, or any other region less reliant on external service providers by:
   a. Providing an Internet network operator with a set of solutions to the service provider denial of Internet access problem by addressing the problem at the applications and routing levels. Thus, the network operator can chose the part of the solution to apply depending on the level at which the problem is occurring.
   b. Allowing an Internet network operator to choose a suitable solution based on the operator's need for quick deployment.
2. At the Kingdom level, one of the major beneficiaries of the outcome of this project can potentially include many Saudi companies, as well as governmental and non-governmental institutions, that require access to the Internet. The implementation of the project's proposed solutions is expected to increase the network resilience, minimize the disruption to network services, and therefore increase the operation and production.
3. In general, major operators, providers, and regulators, including local ones such as KACST, Saudi Telecom Company (STC), and Communication and Information Technology Commission (CITC) may use the results of this project to build a robust infrastructure that allows them to provide Internet access even when one or more of the international service providers decide to deny them such access.
4. Local companies can make use of the proposed solutions to be able to get Internet access even when their main service provider decides to deny it.
5. The P2P DNS solution can help in ensuring the DNS availability, and hence the Internet availability, under physical infrastructure failures that may be due to natural disasters, terrorist attacks, or human errors.
6. Develop in-Kingdom and in-house expertise and subject matter experts in the area of Internet access and its impact on the Internet applications and routing. The experience

was gained through researching the area in the literature, devising different solutions, and experimenting with prototypes of some of the devised solutions.

7. In addition, such a project lays the ground for further research on Internet resilience and robustness. For example, the following are some of the areas identified for further research and development, perhaps in future related projects:
   a. Identification of Severity Levels: In addition to the analysis carried out in the project which focused on both the application and routing functions of the network, more analysis can be performed to cover the last segment related to the robustness of the Internet, namely, the physical layer. In addition, one could define and formulate the severity level for each of these three areas. The formulation would allow classification of these severity levels of each area in comparison with the other areas based on the level of impact it has on the Internet isolation. Further, a study and formulation could be conducted on the severity levels induced by a combination of Internet isolation causes from one or more of the three areas. A study of the impact of the different severity levels on the different network topologies can be performed together with risk assessment on different services affected by the Internet isolation caused by any of the three areas. Finally, the developed solution can be made to respond to the severity level of a particular incident of the isolation scenario by varying specific parameters in the solution.
   b. Prototyping and Experimental Setup: The project prototyped the most promising solution of the two major types of the service provider denial of Internet access problem. However, given the resources and time, the project can be extended later to produce one prototype for each of the promising solutions found for each of the three areas: application, routing, and physical layer. Experimental results pertaining to the robustness of these solutions and the impact of isolation can be collected and compared to provide a matrix characterizing the isolation type and network robustness.
   c. Network Security: The project focused on network resilience in face of Internet service provider intentional denial of service. Many of the already identified isolation types in the areas of application and routing are related to Internet security threats such as denial of service attacks. A future study could attempt to examine and utilize the solutions developed for counteracting these threats but from a network robustness point of view.
8. Moreover, results of general interest to the research community are published at key international computer networking conferences, e.g., IEEE and ACM conferences. In addition, this research work has already led to, and will lead to more publications in refereed reputable journals.
9. Further, the devised solution and, subsequently, the developed prototype can be used to develop a commercial product. The commercial product may consist of a set of software and/or hardware components as well as a set of guidelines applicable to the routing level of any interested institution’s network. Such a commercial product can

help enhance the Internet access, robustness, and/or resilience of the interested institution's network.

10. Likewise, the acquired tools, the devised solutions, and the established prototype will greatly benefit in educating and training many graduate and undergraduate Saudi students attending networking courses at KFUPM. Examples of such courses include, but not limited to, Computer Security, Network Security, Network Design and Management, Systems and Network Administration, and Computer Networks at the undergraduate level, as well as Computer and Network Security, Computer Network Design, Network Management, and Computer Networks at the graduate level.
11. Moreover, the project addressed a new type of malicious Internet activities that has not been addressed previously by other researchers, and opened new research directions.
12. Finally, the experience gained by the researchers can be readily made available to other Saudi universities, private companies, and any other interested institution through consultations, lecture series, short courses, workshops, and/or prototype demonstrations.

The list of targeted end users includes, but not limited to:

1. KACST.
2. CITC.
3. Major Internet service providers (for example, STC, Mobily, etc.).
4. IT departments at governmental institutions and local private companies.
5. Any regional or global major Internet service operator, provider, and/or regulator.

10.0 BROADER IMPACTS OF THE STUDY

Teaching and Training Taught *Computer Network Design* (2nd semester of academic year 2010-2011), *Computer Networks* (2nd semester of academic year 2010-2011, 1st semester of academic year 2010-2011, 2nd semester of academic year 2008-2009), *Computer System Performance* (2nd semester of academic year 2010-2011), *Client Server Programming* (1st semester of academic year 2010-2011, 1st semester of academic year 2008-2009), and *Independent Research* (2nd semester of academic year 2010-2011) courses at the graduate level where the course project was directly related to the Internet access denial problem. The course project involved literature review and/or simulations, prototyping, and lab experimentations of some of the proposed solutions.
Infrastructure Aside from the direct use in the project, the purchased equipment will help in conducting additional research and lab experimentations in different fields of networking. Similar equipment is available in the university but only in teaching labs that are dedicated for regular teaching activities and that cannot be shared with research projects.
Collaborations The main collaboration was with KACST through Dr. Hesham Bin-Abbas who is the director of the Internet Services Unit (ISU) at KACST, and a consultant for the project. The collaboration took place between September 2009 and August 2010, and was focused on exchanging information about KSA's Internet configuration and setup, and sharing and discussing the proposed solutions. The collaboration can be extended to Internet providers, regulators, and operators (e.g., CITC, STC, Mobily) to get useful feedback on the proposed solutions.
Funding The results of the project can lead to further research, and additional funding will be sought from NSTIP and/or internal KFUPM funding in the near future. The focus of the additional research is as outlined in section 9.
Contributions to the Strategic Technologies goals of NSTIP Please see the appropriate table provided in section 7.
Others

11.0 OTHER CONCERNS

There are no additional concerns or comments related to the final reporting of the research program.

12.0 REFERENCES

[1] King Abdulaziz City for Science and Technology, "Strategic Priorities for Information Technology Program," Doc. No. 17P0001-PLN-0001-ER01, http://www.kacst.edu.sa/en/research/Documents/Information Technology.pdf.

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

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

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

[5] R. Ford, M. Bush, and A. Boulatov, "Internet instability and disturbance: goal or menace?," *Proceedings of the 2005 workshop on New security paradigms*, Lake Arrowhead, USA, pp. 3-8, September 2005.

[6] J. Zheng, M. Hu, and L. Zhao, "Enhancing Internet robustness against malicious flows using active queue management," *Proceedings of Second International Conference on Embedded Software and Systems*, Xi'an, China, pp. 501-506, December 2005.

[7] F. Guo, J. Chen, and T.-c. Chiueh, "Spoof detection for preventing DoS attacks against DNS servers," *Proceedings of the 26th IEEE International Conference on Distributed Computing Systems*, Washington, DC, USA, p. 37, IEEE Computer Society, 2006.

[8] M. Lad, R. Oliveira, B. Zhang, and L. Zhang, "Understanding resiliency of Internet topology against prefix hijack attacks," *Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, Edinburgh, UK, pp. 368-377, June 2007.

[9] M. Haungs, R. Pandey, and E. Barr, "Handling catastrophic failures in scalable internet applications," *Proceedings of 2004 International Symposium on Applications and the Internet*, Tokyo, Japan, pp. 188-194, January 2004.

[10] D. Dolev, S. Jamin, O. Mokryn, and Y. Shavitt, "Internet resiliency to attacks and failures under BGP policy routing," *Computer Net*works, vol. 50, no. 16, pp. 3183-3196, 2006.

[11] M. Abu-Amara, A. Mahmoud, F. Azzedin, and M. Sqalli, "Internet Access Denial by International Internet Service Providers: Analysis and Counter Measures," *NSTIP Research Proposal*, April 2008.

[12] S. Cheung, "Denial of Service against the Domain Name System: Threats and Countermeasures," *IEEE Security and Privacy*, vol. 4, no. 1, p. 40, January 2006.

[13] N. Poolsappasit and I. Ray, "Enhancing Internet Domain Name System Availability by Building Rings of Cooperation Among Cache Resolvers," *Proceedings of 2007 Information Assurance and Security Workshop*, West Point, NY, pp.317-324, 20-22 June 2007.

[14] S. Ariyapperuma and C.J. Mitchell, "Security Vulnerabilities in DNS and DNSSEC," *Proceedings of the Second International Conference on Availability, Reliability and Security*, Vienna, pp. 335-342, 10-13 April 2007.

[15] H. Ballani and P. Francis, "Mitigating DNS DoS Attacks," *Proceedings of the 15th Conference on Computer and Communications Security*, Alexandria, Virginia, pp. 189-198, 27-31 Oct. 2008.

[16] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS security introduction and requirements," RFC 4033, Mar. 2005.

[17] V. Pappas, D. Massey, and L. Zhang, "Enhancing DNS Resilience against Denial of Service Attacks," *Proceedings of the 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, Edinburgh, pp. 450-459, 25-28 June 2007.

[18] E. Cohen and H. Kaplan, "Proactive Caching of DNS Records: Addressing a Performance Bottleneck," *Proceedings of the 2001 Symposium on Applications and the Internet*, San Diego, CA, USA, pp.85-94, 8-12 Jan. 2001.
[19] V. Ramasubramanian and E.G. Sirer, "The Design and Implementation of a Next Generation Name Service for the Internet," *Proceedings of applications, technologies, architectures, and protocols for computer communications*, Portland, Oregon, USA, pp. 331 - 342, 30 Aug.-3 Sep. 2004.
[20] R. Cox, A. Muthitacharoen, and R. Morris, "Serving DNS Using a Peer-to-Peer Lookup Service," *Proceedings of first International Workshop on Peer-to-Peer Systems*, London, UK, pp. 155 – 165, 2002.
[21] 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.
[22] S. Murphy, "BGP Security Vulnerabilities Analysis," RFC 4272, Internet Engineering Task Force, January 2006.
[23] K. Butler, T. Farley, P. McDanial, and J. Rexford, "A Survey of BGP Security," AT&T Labs, Research, Florham Park, NJ, Technical Report TD-5UGJ33, 2005.
[24] K. Butler, T. Farley, P. McDaniel, J. Rexford, "A Survey of BGP Security Issues and Solutions," Proceedings of the IEEE, vol. 98, issue 1, pp. 100-122, 2010.
[25] M. O. Nicholes and B. Mukerjee, "A Survey of Security Techniques for the Border Gateway Protocol (BGP)," IEEE Communications Surveys & Tutorials, vol. 11, no. 1, pp. 52-65, Mar, 2009.
[26] O. Nordstrom and C. Dovrolis, "Beware of BGP Attacks," ACM SIGCOMM Computer Communications Review, vol. 34, no. 2, pp. 1-8, Apr. 2004.
[27] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439, Internet Engineering Task Force, November 1998.
[28] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP Misconfiguration," in Proceedings of ACM Sigcomm, Aug. 2002, pp. 3-16.
[29] J. A. Farrar. (2001, Apr.) Merit Network Email List Archives. [Online]. http://www.merit.edu/mail.archives/nanog/2001-04/msg00209.html
[30] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, pp. 582-592, Apr. 2000.
[31] N. Wang, B. Wang, "AT: an Origin Verification Mechanism based on Assignment Track for Securing BGP," IEEE International Conference on Communications, 2008 (ICC '08), pp. 5739 – 5745, 2008.
[32] S. Ortiz, "Securing the Internet's Routing Infrastructure," IEEE Computer, vol. 42, issue 4, pp. 21-23, 2009.
[33] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, L. Zhang, "PHAS: A Prefix Hijack Alert System," Proceedings of the 15th conference on USENIX Security, Vancouver, B.C., Canada, 2006.
[34] U. o. Oregon, "The Route Views Project," http://www.routeviews.org/
[35] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-Weight Distributed Scheme for Detecting IP prefix Hijacks in Real-Time," in SIGCOMM'07, Kyoto, Aug. 2007.
[36] X. Hu and Z. M. Mao, "Accurate Real-time Identification of IP Hijacking," in IEEE Symposium on Security and Privacy, May 2007, pp. 20-23.
[37] B. Quoitin, "BGP-based Interdomain Traffic Engineering," Ph.D. Dissertation, Universite catholique de Louvain, Louvain-la-Neuve, Belgium, Aug. 2006.

[38] B. Quoitin and S. Uhlig, "Modeling the routing of an Autonomous System with C-BGP," IEEE Network, vol. 19, no. 6, pp. 12-19, Nov. 2005.
[39] B. Quoitin and O. Bonaventure, "A cooperative approach to interdomain traffic," in Proceedings of the 1st Conference on Next Generation Internet Networks Traffic Engineering, Rome, Italy, 2005.
[40] J. Postel, "Internet protocol," RFC 791, Internet Engineering Task Force, September 1981.
[41] S. Kini, S. Ramasubramanian, A. Kvalbein, and A. F. Hansen. "Fast recovery from dual link failures in IP networks," Proceedings of 2009 IEEE INFOCOM, Rio de Janeiro, pp.1368-1376, 19-25 April 2009.
[42] C. Perkins, "IP encapsulation within IP," RFC 2003, Internet Engineering Task Force, October 1996.
[43] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin. "Internet routing resilience to failures: analysis and implications," Proceedings of the 2007 ACM CoNext, New York, US, 2007.
[44] R. Atkinson, "Security architecture for the internet protocol," RFC 1825, Internet Engineering Task Force, August 1995.
[45] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina, "Generic routing encapsulation (GRE)," RFC 2784, Internet Engineering Task Force, March 2000.
[46] P. F. Syverson, M. G. Reed, and D. M. Goldschlag, "Anonymous connections and onion routing," *IEEE Journal on Selected Areas in Communications*, vol. 16, pp. 482-494, 1998.
[47] L. Zhuang, F. Zhou, B. Y. Zhao, and A. Rowstron, "Cashmere: resilient anonymous routing," *Proceedings of the 2nd conference on Symposium on Networked Systems Design and Implementation - Volume 2*, Boston, USA, pp. 301-314, May 2005.
[48] M. K. Reiter and A. D. Rubin, "Anonymous Web transactions with Crowds," *Communications of the ACM*, vol. 42, no. 2, pp. 32-38, February 1999.
[49] C. Shields and B. N. Levine, "A protocol for anonymous communication over the Internet," *Proceedings of the 7th ACM conference on Computer and communications security*, Athens, Greece, pp. 33-42, November 2000.
[50] A. AlRefai, "BGP-Based Solutions for International ISP Blocking," Master Thesis submitted to Deanship of Graduate Studies, King Fahd University of Petroleum and Minerals, May 2010.
[51] "WireShark," http://www.wireshark.org.
[52] J. Liu, J. Kong, X. Hong, and M. Gerla, "Performance evaluation of anonymous routing protocols in MANETs," *Proceedings of the 2006 IEEE Wireless Communications and Networking Conference*, Las Vegas, USA, pp. 646-651, April 2006.
[53] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. Groot, and E. Lear, "Address allocation for private internets," RFC 1918, Internet Engineering Task Force, February 1996.
[54] K. Egevang and P. Francis, "The IP network address translator (NAT)," RFC 1631, Internet Engineering Task Force, May 1994.
[55] "OPNET Modeler," http://www.opnet.com/.
[56] I. Stoica, R. Morris, D. Liben-Nowell, D.R. Karger, M.F. Kaashoek, F. Dabek, H. Balakrishnan, "Chord: a scalable peer-to-peer lookup protocol for Internet applications," IEEE/ACM Transactions on Networking, vol. 11, no. 1, pp. 17- 32, Feb 2003.

[57] F. Abdulhameed, "Dynamic Round-Robin Peer-to-Peer (P2P) Domain Name System (DNS)," Master Thesis submitted to Deanship of Graduate Studies, King Fahd University of Petroleum and Minerals, November 2010.
[58] R. K. C. Chang and M. Lo, "Inbound traffic engineering for multi-homed ASes using AS path prepending," in Network Operations and Management Symposium, vol. 1, 2004, pp. 98-102.
[59] "Iperf ," http://sourceforge.net/projects/iperf/.
[60] M. Asif, "Tunneling Based Solution to Bypass Internet Access Denial by International Internet Service Providers," Master Thesis submitted to Deanship of Graduate Studies, King Fahd University of Petroleum and Minerals, December 2010.
[61] A. Al-Baiz, "Internet Denial by Higher-tier ISPs: A NAT-Based Solution," Master Thesis submitted to Deanship of Graduate Studies, King Fahd University of Petroleum and Minerals, January 2011.