2011 2nd International Conference on Networking and Information Technology
IPCSIT vol.17 (2011) © (2011) IACSIT Press, Singapore

Prototyping and Evaluating BGP-Based Solutions to Overcome Malicious IISP Blocking

Amer Al-Ghadhban, Ashraf S. Mahmoud, Marwan Abu-Amara, *Farag Azzedin and Mohammed H. Sqalli

Computer Engineering Department, *Information and Computer Science Department, King Fahd University of Petroleum & Minerals, Saudi Arabia

{g199642170, ashraf,marwan, fazzedin, sqalli}@kfupm.edu.sa

Abstract--Recent worldwide events have shown that some governments have the ability to dictate citizens right to communication by blocking communication services at will. Such blocking affects individual citizens as well as businesses that have become dependent upon unhindered access to the Internet. It is imperative to take measures in order to avoid blocking these services. We consider a scenario where a region of concern is intentionally isolated from accessing the Internet by its primary International Internet Service Provider (IISP). Under the assumption that connectivity to another IISP is available, we prototype and evaluate BGP-based solutions proposed by *Alrefai et al. [1]*. The prototyping and evaluating of these solutions were performed under conditions designed to capture the real Internet's ASes connectivity layout and traffic conditions. To design automated, consistent and repeatable testing procedures, we created four Java based programs which were able to detect the blocking action of malicious IISPs and measuring network convergence time. The resulting convergence time was in the range of 63 – 64 seconds for all of the evaluated solutions.

Keywords-malicious ISP, intentional Internet isolation, controlling outgoing and incoming Internet traffic, BGP configuration and prototyping, Internet resilience.

1. Introduction

The exchange of information over the Internet travels from source to destination through multiple interconnected networks. Some of these networks are small local networks which users are directly connected to and others are large networks that are responsible for interconnecting the smaller ones. The small networks are called Autonomous Systems (AS). An AS is a collection of connected computer networks under the control of a single entity that is usually an ISP or a larger network called an International Internet Service Provider (IISP). To gain control over ISPs or IISPs, means to possess small (local) or potentially large (international) scale determination over the ability of others to communicate and exchange information.

A survey of the literature demonstrates that some governments and private entities have the ability to control the activity of ISPs and thus the transfer of information over the Internet. In [2],[3] it is mentioned that during recent political protests certain governments were able to isolate citizens from gaining access to the Internet. *Stone-Gross et al* [4] provides examples of criminal enterprises influencing ISPs for personal gain such as the Russian Business Network (RBN). Additionally, [5],[6] mention many occasions wherein ISPs have been taken out of service due to malicious activities at the hands of hackers.

Efforts to control internet access for personal and political reasons is a real problem as demonstrated by the examples given in the aforementioned literature. Such control goes against one of the fundamental goals of the establishment of the Internet which can be summarized as open access to information and communication. This work focuses on circumventing intentional internet isolation which occurs by gaining control of IISPs. IISPs have the ability to block the incoming and outgoing internet traffic of one or more ASes. The routing protocol that interconnects different ASes with each other is Boarder Gateway Protocol

(BGP) [7]. While conducting such isolation a malicious IISP creates the appearance that BGP has successfully identified a functioning path for the transfer of information from source to destination. It furthermore disguises the fact that the traffic to and from the affected region is actually being blocked. Whether it is called intentional blocking, intentional isolation or connectivity failure, real solutions must be identified and implemented in order to protect the Internet and the right to open access thereof.

The main objectives of this work are to prototype and evaluate BGP-based solutions for intentional internet isolation. The work is characterized by the implementation of these solutions in the laboratory through a detailed set of experiments. The testing environment and test cases are as close to reality as possible in their configuration and parameters. Performance figures for the different types of traffic considered and the representative configurations were collected and compared to identify suitability and scalability of the proposed solutions. As a condition of validity the testing procedures will be shown to be consistent in time and repeatable. The evaluated schemes were tested in different scenarios and traffic loads. Finally, the performance figures, which are included during the blocking were collected for each evaluated scheme.

2. Literature Review

BGP was designed to provide reliability with minimum overhead. It is not designed with security in mind, which makes it defenseless to imminent routing attacks. In Hu et al [8] they discuss the security weaknesses of BGP which are categorized into three main categories. First, BGP does not provide message integrity and message origin authentication mechanisms and it is vulnerable to a reply attack. Second, BGP does not provide a mechanism to verify the legality of the AS-Path or prefix advertisements from the AS. Third, BGP does not verify the validity of BGP attributes included in the BGP advertisements.

BGP attacks have been discussed in Nordstrom and Dovrolis [9] and they name four main purposes for these attacks as follows: 1. *Blackholing* 2. *Redirection* 3. *Instability* and 4. *Supervision*. *Blackholing* is an attack method of dropping all the traffic bypassing the attacking router. Also, the attacker may drop only traffic that belongs to a specific AS. *Redirection is* a method of redirecting all traffic or a specific user's traffic to another destination or server for content analysis. *Supervision* is similar to the previous method, but the purpose is to modify the traffic content then forward it to the right destination. *Instability* is an attack method initiated to harm the network with destablizing events such as injecting false updates, link flapping or announcing successive advertisement then withdrawals. In this work we are considering the *Blackholing* attack method where a malicious IISP isolating a region of concern from the Internet.

In Omer et al [10] a new method and network model are proposed to measure the resilience of the Internet's infrastructure by identifying the vulnerabilities of global undersea optical fibers. They evaluate the effect of the possible losses in these cables against the Internet infrastructure and the recovery from it. Soo Kim et al. [11] conducted a study which proved that modifying the network topology improves its resilience. Cohen et al. [12] have shown that scale-free networks like the Internet are susceptible to an intentional attack because there exists few ASes, e.g. IISPs, which aggregate a large number of the internet connectivity. They proved mathematically that the removal of one or more of these ASes causes momentous Internet outage.

A more realistic study with practical analysis has been conducted by Dolev et al. [13] wherein they assume the Internet ASes are connected as a directed graph (policy-based). They made their analysis and measurements of the resiliency of the Internet based on that assumption. In addition, they concluded that the Internet is highly sensitive to an intentional attack and could possibly crumble very fast. In contrast, the Internet is resilient to random failure. A major investigation into the sensitivity of the Internet to a random fault and attacks was made by Park et al. [14]. They concluded that the Internet is robust and is becoming more robust with time against random failures; and the average internet diameter is stable even though the number of internet users is increasing.

3. Problem Statement and BGP-Based Solutions to Be Tested

3.1 Problem Statement

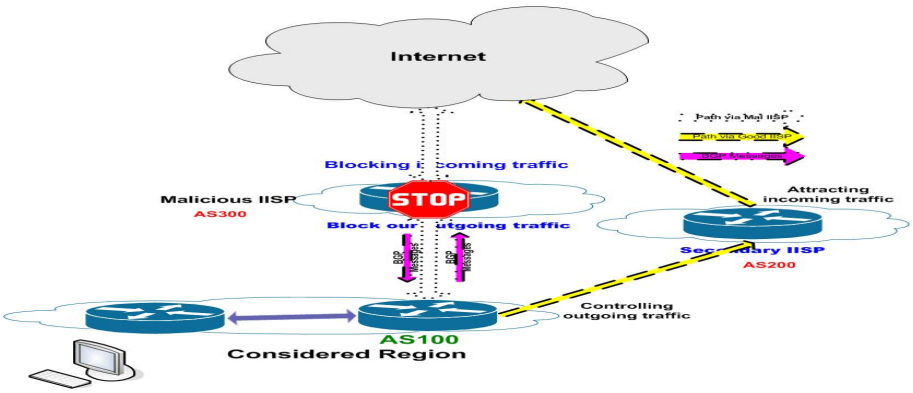

Fig.1 Malicious IISP blocking the concerned (considered) region traffic while still exchanging BGP messages.

This study is going to focus on the network configuration depicted in Fig.1 where the region of concern, denoted by AS100, is connected to the Internet through a primary IISP, defined here as the malicious IISP and denoted by AS300. The region of concern is also connected through a secondary IISP, called here the good IISP and specified by AS200. As indicated by its definition the primary IISP for intentional reasons blocks the incoming and outgoing internet traffic of the region of concern. Although, the malicious IISP isolates the internet traffic of the region of concern, the malicious IISP's BGP speaker is still exchanging *keepalive* and BGP messages with the concerned region's BGP speaker and advertising its prefixes on the internet. However, without these messages being exchanged, the concerned region's BGP speaker will directly route the outgoing traffic through the good IISP and acquire incoming traffic through it as well. The boarder router that carries the traffic between different ASes is called a BGP speaker.

There are two methods to get around Internet connectivity isolation: traffic identity hiding and traffic control. The former method can be achieved through implementing NATing or tunneling techniques in local ISP BGP speakers and a cooperative AS in the Internet which are implemented by [15][16]. The latter method may be accomplished through exploiting the availability of a secondary IISP and employing the BGP attributes and routing configuration commands to route the outgoing traffic and attract the incoming traffic through the secondary IISP.

3.2 BGP-Based Solutions To Be Tested

In order to attract the incoming and outgoing traffic out of the primary IISP path a modification to the BGP path selection procedure of the intermediate routers is needed. The BGP routing protocol has a unique path selection procedure explained in [7]. The BGP protocol offers configuration commands capable of controlling the BGP path selection procedure attributes, such as *AS-Path Pre-pending* and *Local-Preference*. Some of the evaluated solutions can influence the incoming traffic to go through the good IISP and others can control the outgoing traffic. Table.1 shows the classification of the BGP functions based on their ability in controlling the outgoing traffic or attracting the incoming one.

First of all, the solutions that can influence incoming traffic are *AS-Path Shortening, More Specific Prefixes,* and

TABLE.1 THE CLASSIFICATION OF THE BGP METHODS.

Ability / BGP function	Incoming *Attracter*	Outgoing *Outforwarder*
AS-Path Shortening	Yes	No
More specific Prefixes	Yes	No
BGP Community	Yes	No
Local Preference	No	Yes

BGP Community. In this work these solutions are called *Attracters*. *AS-Path Shortening [1]* is implemented by the region of concern advertising its prefixes using *distribute-list* configuration commands and the good IISP originating the concerned region's prefixes using *network* commands. As a result, the prefixes appear in the Internet as belonging to the good IISP and the first AS number in the AS-Path associated with these prefixes is the good IISP's AS number (i.e. here is AS200). Hence, the concerned region's prefixes that are advertised via the good IISP appear in the Internet with a shorter AS-Path than the ones that advertised via

the malicious IISP. In *More Specific Prefixes*, the routing table algorithm selects the longest prefix match as a network destination to the forwarding traffic. Based on this, attracting the traffic through the good IISP can be achieved by advertising long prefixes. However, the accepted length of the prefix on internet routers is limited to a fixed length [17]. As for *BGP Community*, the BGP protocol has a community attribute which is used in the evaluated solutions to influence the incoming traffic going through the good IISP. This attribute enables any AS to advertise its prefixes associated with a community value to its neighbor ASes. When the neighboring AS gets the *community* advertisement it's going to look at the community value then perform an action based on it. The action performed here is assigning a higher *Local Preference* value to the path where the neighbor gets the *community* advertisement from. This solution requires cooperation from most of the ASes between a source and the region of concern. Table 2 shows part of the community values that are used by Sprint [18], one of the largest IISP in the world. Any subscriber ISP can influence the BGP path selection procedure of Sprint by associating the appropriate community value with its advertisements. To forward the region of concern outgoing Internet traffic through the good IISP the *Local Preference* attribute is used.

The BGP-based solutions are a combination of *Attracter* with *Outforwarder.* In [1] they propose only one solution to control the outgoing traffic because the outgoing traffic is under the control of the region of concern but the incoming traffic is under control of the source AS and the ASes in between. Table 3 illustrates the BGP-based solutions that were evaluated in this work.

4. Prototype Design and Implementation

The BGP-Based solutions are evaluated in a real laboratory. The laboratory set up contains seven Cisco 2811 routers, four Catalyst 2950 switches, one workstation and three servers. The three servers are set up as they would be on the Internet side and the workstation as it would be on

TABLE.2 Sprint Local Preference BGP Community Value

BGP Community Value	Resulting Local Pref
1239:70	70
1239:80	80
1239:90	90
1239:100	100
1239:110	110

the local side. Also, each server is assigned to a specific Internet application: FTP, HTTP or VoIP. Furthermore, one of these servers and the workstation are equipped with WireShark [19] network analyzer to collect the statistics of each test.

Four Java network programs are also programmed to automate the testing environment. The first and main software program, called here *checker*, is capable of checking the Internet connectivity installed in one of the machines on the local side (AS100). When it faces a sequence of timeout messages, it can immediately and remotely login to the local side (AS100) BGP speaker and configure it with one of the recommended BGP-Based solutions. The second software configures the malicious IISP (AS300) BGP speaker with the Access Control List (ACL) commands to block the outgoing and incoming traffic of the local side (AS100). The third and fourth software are designed to erase the previous configurations to conduct new testing attempts.

These four programs are precisely modeled to enhance the manual procedure followed by a network administrator that relies on human intervention. Whereas the manual procedure provides inaccurate results and inconsistent movements between steps, the implemented procedures are automated and provide consistency in time and movement between steps.

4.1 Laboratory Scenarios

The AS-Path length from a local AS to a remote AS through two different IISPs is not always identical. Based on this fact, the evaluated solutions are examined in two dissimilar laboratory scenarios. The first scenario, called here identical scenario, is shown in Fig.2 which demonstrates where the AS-Path length from the local side (AS100) to the Internet side (AS600) over the two IISPs are the same. The second scenario, called the non identical scenario, is shown in Fig.3 which demonstrates where the AS-Path from the local side to the Internet side through the good IISP (AS200) is longer than the AS-Path to the same side

when it goes through the malicious IISP (AS300). In Fig.2 and 3, AS100 represents the region of concern and AS600 represents the Internet side where three servers are installed with different Internet applications: FTP, HTTP and VoIP. Also, the two figures show AS300 as the malicious IISP that is blocking the outgoing and incoming traffic of the AS100.

TABLE.3 THE BGP-BASED SOLUTIONS

	Local Preference
AS-Path Shortening	✓
More specific Prefixes	✓
BGP Community	✓

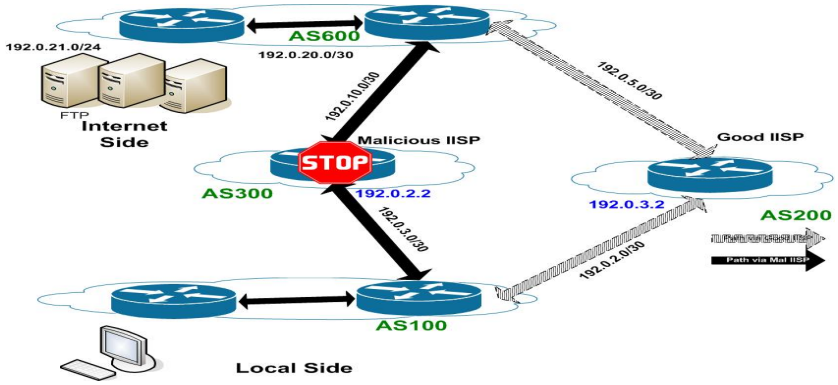

Fig.2 Identical Scenario

However, AS100 is *multihomed* to a secondary IISP (good IISP) where the evaluated BGP-solutions route the outgoing traffic and attract the incoming traffic through it. Routers in the laboratory are configured to provide the desired connectivity.

Every solution is tested with the same testing procedure. The testing procedure consists of three dissimilar traffic configurations; the load, being one Internet application, within each configuration is tested with three different link capacities. Additionally, performance figures for the implemented Internet applications are measured and analyzed. Performance figures include convergence time, number of lost packets and end-to-end delay. During the switching, from the malicious IISP to the good IISP, the performance figures of a current HTTP, FTP or VoIP session are affected. Also, a check is made on whether or not the BGP-solutions face the same effect. This allows for a comparison of these solutions based on the effect of switching upon the performance figures. To measure this effect, a network analyzer capable of measuring the number of lost packets and end-to-end delay of the current session is used in addition to another program capable of measuring network convergence time.

4.2 Malicious IISP Blocking Configuration

We use ACL to permit the exchange of BGP messages and advertisements between the malicious IISP and the region of concern BGP speakers and to deny the rest of the traffic. Two ACLs are implemented in the malicious IISP BGP speaker, one blocking the outgoing traffic and another blocking the incoming traffic. The first ACL is implemented in the closest interface to the local side (192.0.2.2). The second is implemented in the interface that is closest to the Internet side (192.0.10.1). Fig.4 shows how the local side BGP speaker can't *ping* to the FTP server after implementing the ACL. Even though the alternative path is available, the local side BGP speaker still sends the outgoing traffic via the malicious IISP that is dropping it. The *traceroute* result also demonstrates that the local BGP speaker still prefers the path via the malicious IISP (192.0.2.2). After implementing one of the evaluated

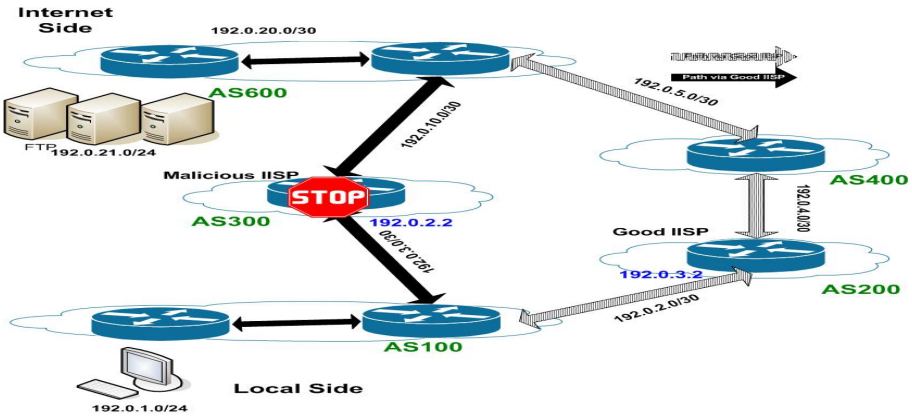

Fig.3 Non-Identical Scenario

solutions, the local side can ping the Internet side and the *traceroute* result shows the packets have gone over the good IISP (192.0.3.2).

4.3 Internet Application Testing Procedure

The Internet applications testing procedure is divided into multiple configurations. In each configuration a combination of the recommended BGP solutions is tested with one of the following Internet applications: FTP, HTTP or VoIP. During testing of the configurations, a network analyzer, Wireshark, is used to collect the required performance figures.

1) Procedure 1

In configuration 1, one of the Internet application's clients, such as FTP client, residing in the Internet side (AS100) starts communicating with the compatible server that resides in another Internet side (AS600). At the same time Wireshark and *checker* programs are running. At a specified instance in time a program connects to the malicious IISP router (AS300) and configures it with the blocking configurations. When the *checker* faces Internet connectivity loss, it immediately connects to the local side BGP speaker (AS100) and configures it with one of the BGP solutions, such as *Local Preference* with *AS-Path Shortening*.

Configurations 2 and 3 are the same as configuration 1, except using HTTP and VoIP respectively.

2) Procedure 2:

Procedure 2 is identical to procedure 1, but using a different BGP solution.

4.4 Convergence Time Procedure

The time between detecting the malicious action and recovering from it is measured by the software *checker*. This time includes, the required time for detecting the action, telnet to the BGP speaker and the required waiting time for getting the echo-replies from the Internet side's server. In this work this time is called *convergence time*. This procedure is repeated 10 times and the average results are considered.

Over malicious IISP path
C:\Users\marwan>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\marwan>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 IISP
Trace complete.
Over alternate path after implementing one of the solutions

C:\Users\marwan>ping 192.0.21.6 → Server in AS 600
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\marwan>TRACERT 192.0.21.6
Tracing route to ALIEN-PC [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 IISP
  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  ALIEN-PC [192.0.21.6]
Trace complete.
```

Fig.4 Ping and trace route results from local to Internet side before and after performing one of the evaluated solutions.The convergence time testing procedure is as follows.

1. Merging one solution from the *Outforwarder* with one from the *attracter* from the recommended BGP solutions and programming the *checker* with this combination.
2. Running the *checker*, the *checker* maintains a sequence of pings to the FTP server, the second software configures the malicious router with the ACL. When the *checker* gets a sequence of failed replies, it records the time and configures the router with the combination, then maintains a sequence of pings to the same server, and records the time when it gets a reply from the server.
3. The third and fourth programs erase the blocking and the solution configurations and clear all the BGP tables.
4. Waiting 60 seconds, the estimated *convergence time,* to make sure the BGP and routing table in all the routers are cleared from the previous configurations
5. Go to step 2, if the number of tries is less than 10
6. Go to step 1 to test another solution, if the number of tries equals 10.

5. Results And Analysis

Due to the length of the AS-Path from an Internet source to a destination through any two or more different upper ISPs not always being identical and some of the proposed solutions not being able to work with non-identical scenarios, we have measured the BGP-Based solutions with different background traffic loads: 80%, 50% and 25%. The laboratory is configured with data rates of 1.544 Mbps. The 80%, 50% and 25% background traffic loads means the link capacity is 368 kbps, 786 kbps and 1.28 Mbps, respectively. Additionally, the FTP file size is 10MB, the HTTP page is 6MB (10 images) and the VoIP call lasts for 100 seconds.

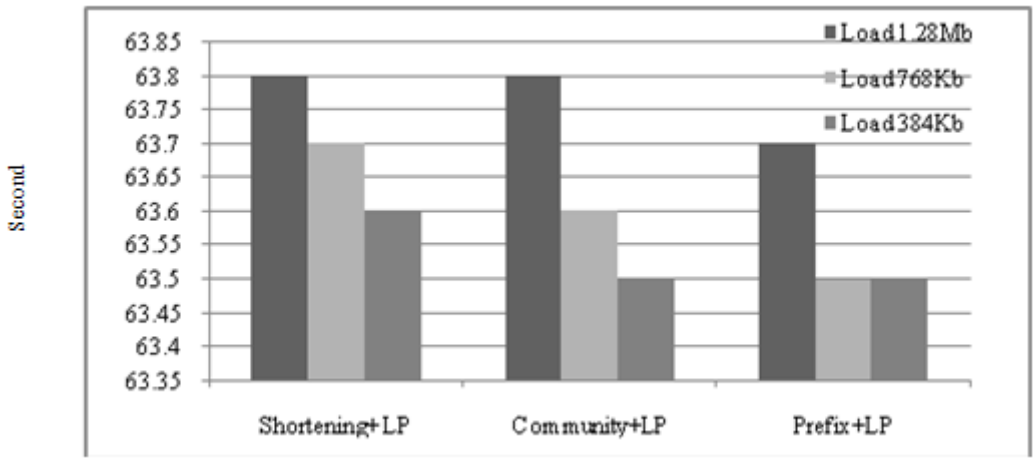

Fig.5 Convergence Time Results of the BGP-Based Solutions in Identical Scenario. Note: LP = Local Preference

We have evaluated the proposed solutions using two different laboratory scenarios, identical and non-identical. We found that the *AS-Path Shortening + Local Preference* solutions can work only with the identical scenario. In contrast, the *More Specific Prefix + Local Preference* and *BGP Community + Local*

Preference can work with the identical and non-identical scenarios. We noticed the HTTP starts slowly in begging of opening the webpage and after recovering from the blocking incident. In contrast the FTP starts fast in both situations.

5.1 Convergence Time Results

The average results of 10 runs of the *convergence time* procedure of each of the BGP-based solutions with different background traffic loads are illustrated in Fig.5. The y-axis represents the time in seconds and the x-axis represents the evaluated solutions with different background traffic loads. The resulting *convergence time* of the evaluated solutions is between 63 to 64 seconds. The *convergence time* exchanged messages are few in number and small in size. Thus, the affect of the background traffic load on the convergence time is very small. The *More Specific Prefix + Local Preference solution* always gives the fastest convergence time even with the different background traffic loads.

5.2 End-to-end Delay

The end-to-end delay of the examined Internet applications is shown in Fig.6. The y-axis in the figure displays the time in seconds and the x-axis displays the examined Internet applications with different background traffic load. We examined the FTP end-to-end delay by downloading a file stored on the FTP server residing in the Internet side from FTP client installed in the workstation resides in the local side. During the downloading, at the instant in time the blocking action is performed the solution is activated. The same procedure is performed to examine the HTTP but with 6 MB webpage. This means the *convergence time* is included in the posted end-to-end delay results in Fig.6. The *More Specific Prefix + Local Preference* solution provided the lowest end-to-end delay among the evaluated solutions.

5.3 Percentage of Traffic Drop

The percentage of the lost packets for the evaluated BGP solutions is displayed in Fig.7. The y-axis represents the

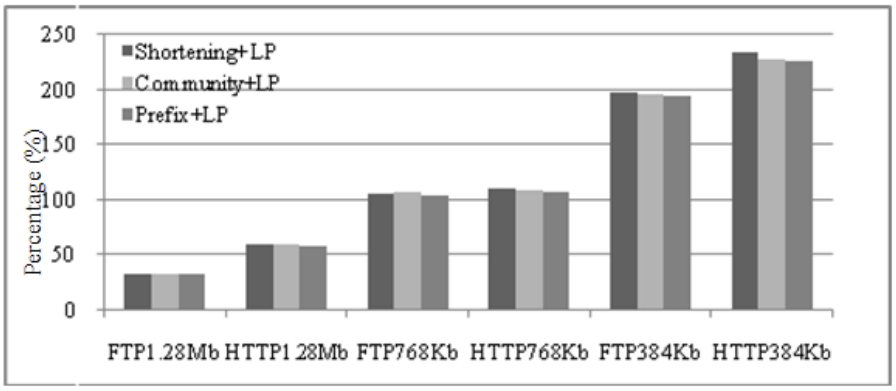

Fig.6 End-to-end Delay of the Examined Internet Applications.

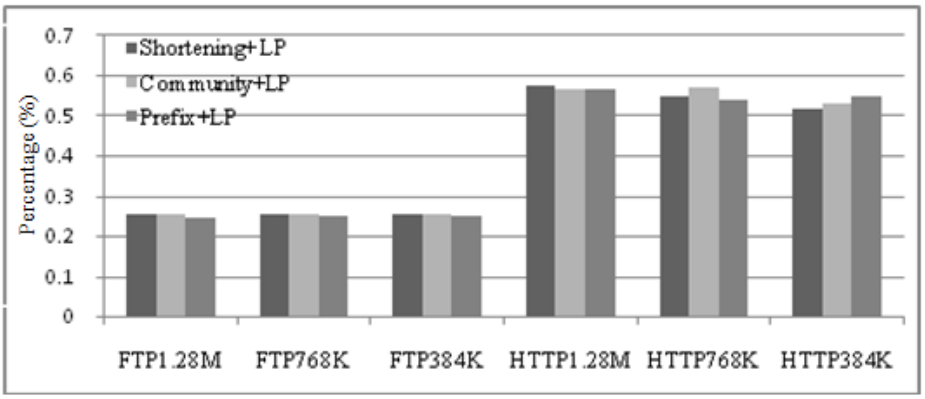

Fig.7 Percentage of the lost packets during the blocking action

percentage of lost packets in relation to sent packets and the x-axis represents the examined Internet applications with dissimilar background traffic load. By traffic drop, we mean the number of lost packets that

were dropped during the blocking incident. The percentage of the lost packets with the HTTP application is double the value of the FTP application. This small percentage proves the sensitivity of the TCP protocol to the carrier. The percentage of the lost packets for the VoIP is in the range of 40% to 41% for all the evaluated solutions and can't be set in the same graph together with the TCP applications.

6. Conclusion

The BGP-based solutions that are proposed by Ahmed et al [1] were prototyped and evaluated in a real laboratory. The laboratory was configured with the configurations that are usually applied in ISP routers. Furthermore, the solutions were evaluated in two different lab scenarios: identical and non-identical. The effects of these solutions were measured by implementing them during different Internet application streams. The evaluating procedures were also conducted with different background traffic loads. Four Java programs were programmed to automate the testing environment and make repeatable evaluation procedures. The resulting convergence time is in the range of 63 – 64 seconds for all of the evaluated solutions. The maximum percentage of the end-to-end delay is about 230% found with HTTP25% and about 190% with FTP25%. The minimum percentage is about 30% found with FTP80% and about 55% with HTTP80%. All the evaluated solutions work fine with identical and non-identical scenarios with the exception of the *AS-Path Shortening* solution.

7. Acknowledgment

The authors acknowledge the support provided by King Fahd University of Petroleum and Minerals (KFUPM). This project is funded by King Abdulaziz City for Science and Technology (KACST) under the National Science, Technology, and Innovation Plan (project number 08-INF97-4).

8. References

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

[2] Martyn Williams, "Blocking Internet cost Egypt at least $90M, says OECD," COMPUTERWORLD, http://www.computerworld.com/s/article/9207938/Blocking_Internet_cost_Egypt_at_least_90M_says_OECD, February, 2011.

[3] Ashish, "Indian Government blocking access to Blogger sites ?," Indian Political Blog, February, 2011.

[4] B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda, "FIRE: FInding Rogue nEtworks," In Proc. of 25th ACSAC, 2009.

[5] Catalin Cimpanu, "Australian ISPs Battle DDOS Attack," Softpedia, http://news.softpedia.com/news/Australian-ISPs-Battle-DDOS-Attack-119603.shtml, August, 2009.

[6] "Denial-of-service attack,"Wikipedia, http://en.wikipedia.org/wiki/Denial-of-service_attack

[7] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006.

[8] Y. Hu, A. Perrig, and D. Johnson, "Efficient security mechanisms for routing protocols," in Proc. ISOC Network and Distributed Systems Security Symp. (NDSS), San Diego, CA, Feb. 2003.

[9] O. Nordstrom, C. Dovrolis," Beware of BGP attacks," in *ACM SIGCOMM CCR*, April 2004.

[10] M. Omer, R. Nilchiani, and A. Mostashari,"Measuring the Resilience of the Global Internet Infrastructure System," *Accepted for publication: IEEE International Systems Journal,* 2009.

[11] S. Kim, H. Lee, and Y. W. Lee, "Improving Resiliency of Network Topology with Enhanced Evolving Strategies," in *Proceedings of the Sixth IEEE International Conference on Computer and Information Technology* 2006, p. 149.

[12] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, "Breakdown of the internet under intentional attack," *Phy Rev Lett,* vol. 86, pp. 3682-3685, Apr. 2001.

[13] D. Dolev, S. Jamin, O. Mokryn, and Y. Shavitt, "Internet resiliency to attacks and failures under BGP policy routing," *Computer Networks: The International Journal of Computer and Telecommunications Networking* vol. 50, pp. 3183 - 3196 Nov. 2006.

[14] S.-T. Park, A. Khrabrov, D. M. Pennock, S. Lawrence, C. L. Giles, and L. H. Ungar, "Static and dynamic analysis of the Internet's susceptibility to faults and attacks," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE* 2003.

[15] M.Abu-Amara, M.Asif, M. Sqalli, A. Mahmoud, F.Azzedin, "Resilient Internet Access Using Tunnel-Based Solution for Malicious ISP Blocking," in *Proceedings of the 3rd IEEE International Conference on Communication Software and Networks, May, 2011.*

[16] A.Albaiz, "INTERNET DENIAL BY HIGHER-TIER ISPs:A NAT-BASED SOLUTION," Master Thesis, King Fahd University of Petroleum and Minerals, March 2010.

[17] B. Quoitin, "BGP-based Interdomain Traffic Engineering," Ph.D. Dissertation, Universite catholique de Louvain, Louvain-la-Neuve, Belgium, Aug. 2006.

[18] "BGP Community String for Sprint AS 1239," http://www.ipbalance.com/routing/bgp/bgp-community-attributes-list/369-bgp-community-string-for-sprint-as-1239.html

[19] "WireShark," http://www.wireshark.org