Arab J Sci Eng (2012) 37:1911–1928
DOI 10.1007/s13369-012-0307-1

RESEARCH ARTICLE - COMPUTER ENGINEERING AND COMPUTER SCIENCE

Ashraf S. Mahmoud · Ahmad Salam Alrefai · Marwan Abu-Amara · Mohammed Sqalli · Farag Azzedin

Qualitative Analysis of Methods for Circumventing Malicious ISP Blocking

Received: 12 June 2010 / Accepted: 30 January 2011 / Published online: 24 May 2012

Abstract Today, the internet is crucial for everyday needs including business and governmental applications, therefore its resiliency to attacks and outage is critical. International Internet Service Providers (IISPs) usually provide connectivity to customers, but can intentionally, as in the case of enforcing an internet embargo, or unintentionally, as in the case of a security breach, block incoming and outgoing traffic while still advertising reachability information to the prefix they seem to provide connectivity for. These two scenarios result in isolating the prefix owner from the internet. Under the assumption that another cooperating IISP exists, the paper investigates three major techniques to overcome internet blockage due to internet embargo or a security breaches. First, a solution based on BGP tuning is presented where the focus is on configuring router(s) to direct the outgoing traffic and to influence the incoming traffic to pass through the cooperating IISP. The second solution utilizes virtual peering which uses a multi-hop BGP session and establishes a tunnel through the intended ISP to provide a deterministic control of incoming traffic. For the third solution, we propose a virtual transit approach in which multiple routers distributed across the internet work as a transit for the blocked local region. This solution extends virtual peering where routers advertise a shorter path to other peers on the internet. We compare the three proposed solutions in terms of traffic filtering, setup overhead, communication overhead, difficulty to offset the solution, and scalability. Finally, we also present a brief validation and proof of concept for the BGP-based solution utilizing simulations.

Keywords Network security · IP hijacking · Malicious ISP · BGP tuning · Virtual peering · Virtual transit

A. S. Mahmoud (✉) · A. S. Alrefai · M. Abu-Amara · M. Sqalli
Computer Engineering Department, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
E-mail: ashraf@kfupm.edu.sa

A. S. Alrefai
E-mail: salam@kfupm.edu.sa

M. Abu-Amara
E-mail: marwan@kfupm.edu.sa

M. Sqalli
E-mail: sqalli@kfupm.edu.sa

F. Azzedin
Information and Computer Sciences Department, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
E-mail: fazzedin@kfupm.edu.sa

الخلاصة

تعتبر الإنترنت أساسية لجميع الاستخدامات اليومية و للتطبيقات الحكومية والأعمال التجارية في هذا العصر، لذا فإن القدرة على صمود الإنترنت أو حظره في حال الهجمات أمر في غاية الأهمية. فعادة ما يقوم مزودو الإنترنت الدوليون بتوفير خدمة الاتصال للزبائن، ولكن من الممكن أن يحجب هؤلاء المزودون وصول الإنترنت، بينما يستمرون في الإعلان عن قدرتهم على إيصال الحزم إلى وجهتها المطلوبة، إما بشكل متعمد كما في حالة فرض حظر على الإنترنت عن منطقة محددة أو بشكل غير متعمد، كما في حالة وجود خرق أمني. وهذا الأمر سينتج عنه عزل للوجهة المطلوبة عن الوصول للإنترنت.

سوف تدرس هذه الورقة العلمية ثلاث تقنيات رئيسية لمقاومة انقطاع الإنترنت بسبب الحظر المفروض أو الخروقات الأمنية، وذلك بافتراض وجود مزود دولي آخر للإنترنت من الممكن الاستعانة به. يعتمد الحل الأول على ضبط بروتوكول الـ BGP لتوجيه الحزم الصادرة وللتأثير على الحزم القادمة لتمر عبر مزود الإنترنت الدولي المتعاون. أما الحل الثاني فيستخدم التناظر الافتراضي باستخدام multi-hop BGP ووضع نفق يمر من خلال مزود الإنترنت المتعاون لكي يتم التحكم بشكل قاطع بالحزم الواردة. ويتمثل الحل الثالث في استخدام العابر الافتراضي الذي يعتمد على توزيع عدد من الموجهات حول الإنترنت لكي تعمل كنقطة عبور للمنطقة المحظورة. والعابر الافتراضي هو تطوير للتناظر الافتراضي حيث يتم من خلاله الإعلان عن مسار أقصر للجهة المطلوبة في الإنترنت. لقد قمنا بمقارنة هذه الحلول من حيث تصفية حركة المرور، والوقت اللازم للإعداد، والاتصالات الإضافية، وصعوبة كسر الحل، وقابلية التوسع.

1 Introduction

The internet is one of the most important means to communicate with others, and to obtain or to provide services. The number of people using the internet in March 2009 exceeded 1.5 billion users which constitutes about one quarter of the whole population of the world [1]. Enhancing the resiliency, availability, and security of the internet is one of the major requirements for Internet Service Providers (ISPs) and for internet users who benefit from its services. If the internet is unavailable, then the user will not be able to access services during the unavailability time. If it is not resilient or not secure then, in the presence of any problem such as internet blockage or attack, the access to the internet will be denied and the user will not be able to get to the subscribed services.

There are several causes of internet outage. For example, internet outage may occur due to software or hardware failure, or misconfiguration of routers. Link cuts can also cause internet unavailability [2]. Internet outages may also occur when the International Internet Service Provider (IISP) does not, whether intentionally or unintentionally, provide the needed internet services either at the application level or at the routing level [2]. The unintentional case may arise due to malicious attacks such as distributed denial of service (DDoS) that target the IISP. Security breaches and tampering with the IISP hardware and/or software may also lead to internet outages as the needed internet services for connectivity such as Domain Name System (DNS) and routing may become unavailable. On the other hand, for the intentional case the IISP is deliberately making some of the internet services unavailable not due to attacks but as an attempt to enforce an internet embargo or blockage on the targeted region. The former intentional act may be politically motivated. Accessing the internet even when the IISP is denying the service, must be a concern for Regional Internet Service Providers (RISPs) who wish to provide resilient internet access for their subscribers.

This paper considers the scenario of a region of concern that is connected to the internet through a primary IISP and at least one secondary IISP. At some point in time, the primary IISP starts to drop traffic belonging to the region of interest, while still advertising reachability to the region. The paper focuses on providing a number of solutions to the above IISP blocking problem with little or no modification to the conventional internet functions. The prescribed solutions herein do not make a distinction whether the internet blocking is intentional or unintentional and function in the exact same manner for the two cases. Finally, the paper provides a qualitative analysis and comparison between the given solutions.

For the remainder of the paper, it is assumed that the blocking enforced by the IISP is intentional and the term "malicious" ISP is used to refer to the IISP that is dropping the traffic belonging to the region of interest. However, the developed solutions and analysis are equally applicable to the case when the blocking is caused unintentionally and is due to a security breach at the IISP. The term "good" ISP will be used to refer to the cooperating ISP.

The paper is organized as follows. The next three sections provide a brief motivation for the work, the needed background, and a literature review of related work in the field, respectively. Section 5 gives the formal description of the problem, while we present possible BGP-based solutions of the problem in Sect. 6. Section 7 provides the discussion and qualitative analysis of the proposed methods for circumventing IISP blocking, while Sect. 8 presents a brief validation and proof of concept for the BGP-based solutions. Finally, the paper concludes in Sect. 9.

2 Motivation

The impact of the intentional internet blocking by IISPs depends on many factors, such as the locations and addresses of source and destination networks, the location and size of the malicious IISP, and the routing policies of the intermediate networks. In general, the malicious IISP will block access to its own networks. In addition, other networks may be blocked because the malicious IISP is in their routing path. As a result, the victim network may become unable to reach most of the other networks on the internet, causing a complete internet isolation for the victim network.

From a business point of view, an ISP that performs internet blocking on a network is risking its reputation. ISPs are supposed to provide the promised service of traffic routing without such filtering or blocking. Hence, when the victim networks detect and report the act of internet blocking, that ISP may lose its reputation, and eventually its customers.

However, there are many other forces and motivations that may push an ISP to perform internet blocking on an organization or a country. An attacker that targets a specific organization can perform the attack at the ISP level, by hacking into the ISP's network and reconfiguring it to block that targeted network. Internet blocking could also be driven by political motivations. Governments may force their ISPs to block some services from a specific country or region in attempt to establish an internet embargo on the targeted region.

Many large services and networks have been attacked recently for political motivations. Gmail, for example, had many recent attacks targeting email accounts of Chinese human rights activists [3]. Twitter, a popular social network, has also been attacked recently by hackers from Iran [4]. These types of attacks are driven by political forces.

These reasons, and many others, can encourage ISPs to perform internet blocking on a specific network. The potential risk and impact of internet blocking may be critical. Therefore, solutions for this problem should be studied and deployed.

3 Background

Networks are managed by ISPs that can be classified into tiers depending on the size of the networks operated by the ISP. Tier 1 ISPs often have global backbone networks while tier 2 ISPs have regional wide networks. Tier 3 ISPs normally provide internet access to end users [5]. A network that is administered by a single organization is called an *Autonomous System* (AS). The internet is composed of a large number of Autonomous Systems (ASes). Figure 1 depicts a general overview of the internet.

Each cloud in Fig. 1 represents an AS. Within each AS, routers run an Interior Gateway Protocol (IGP) such as the Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol. On the other hand, for communication between ASes, routers run an Exterior Gateway Protocol (EGP) such as the Border Gateway Protocol (BGP).

The Border Gateway Protocol (BGP) defined in RFC 4271 [6] is the de facto standard for routing in the internet. BGP is a path vector protocol which is based on a distance vector approach. In contrast to a simple distance vector protocol, such as RIP, which has a simple metric coupled to it, BGP utilizes many attributes and enables the selection of the best route based on a local policy. Attributes supported by BPG include the `Local-Pref` and the `AS-PATH`. The `Local-Pref` value enables the ISP administrator to prioritize the routes for a specific destination, while the `AS-PATH` is a sequence of AS numbers that defines the route. BGP is used by ASes to determine how to route traffic to other ASes. Thus, the main function of BGP is the exchange of *Network Reachability Information* (NRI) between BGP systems. This information includes a set of ASes which is sufficient for building a graph of ASes in order to reach a destination. In this case, a *BGP speaker*, a router that runs the BGP protocol, can eliminate routing loops and enforce routing policies. BGP uses a destination-based forwarding paradigm where the router's decision to forward a packet is only based on the destination address. Therefore, BGP supports only policies that obey the rules of this forwarding paradigm, and constructed a policies should be based only on the destination address. In order to have a reliable communication for BGP, BGP runs over the Transmission Control Protocol (TCP).

Contractual agreements or Service Level Agreements (SLAs) are defined between ISPs and their clients to determine the quality of service that must be provided. BGP allows one to filter received routes from other peers and to filter advertisements to them; also it might select routes based on specific conditions. In essence, BGP is utilized as a "policy-based inter-domain routing protocol." [7].

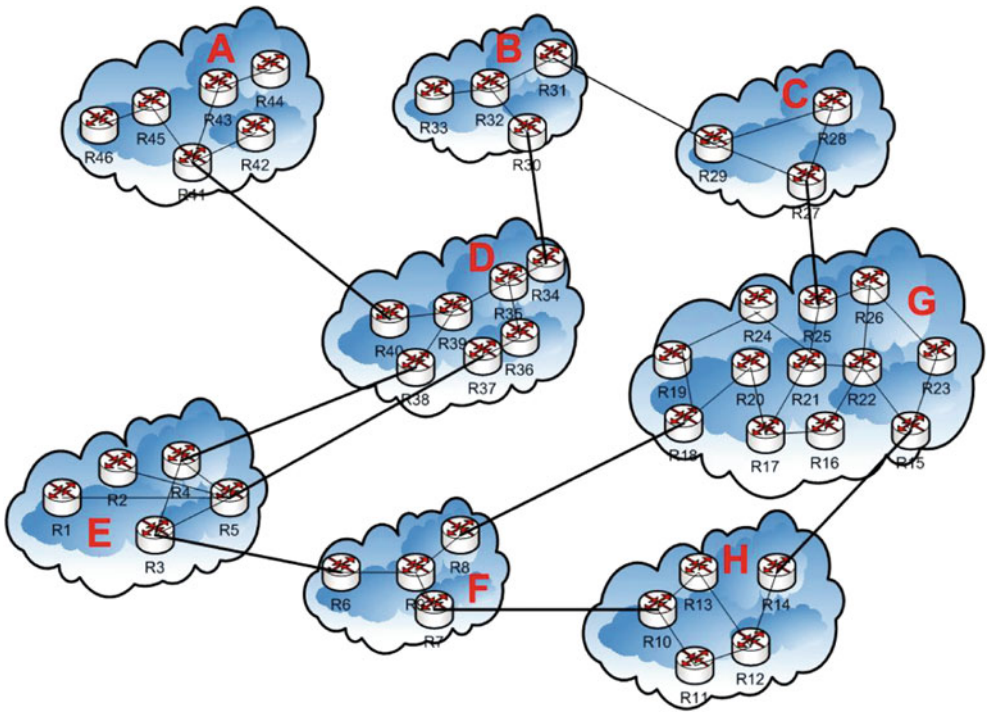

Fig. 1 Overview of the internet

BGP defines a number of rules that help in selecting the best route. The first rule is to take the route with a higher degree of preference that is usually set by the `Local-pref` attribute. If more than one route has the same local preference, BGP will look for the shortest `AS_Path`. If they are the same, BGP will look for the lowest `Multi_Exit_Disc (MED)` value if they have the same `Next_Hop` value. Then, if a decision cannot be taken, the path with the lowest cost to the `Next_Hop` of the route is selected. Finally, the route with the lowest BGP identifier is chosen if more than one route is still present [5].

A *transit* network is a network that delivers traffic from one network to another, whereas a *stub* network is a network that does not work as a transit network. The stub network only carries traffic if it is either a source or a destination for that traffic. The relationship between routers in the internet is either *peer-to-peer* where routers exchange traffic in one to one relationship without money involvement, or *customer–provider* where the customer pays money for the provider for connectivity to the internet. A *multihomed* customer is one that has more than one provider that it is connected to.

Since BGP is vector-based, one router does not have the full picture of the whole network. Moreover, an IISP has control over BGP messages and user traffic, and can maliciously deny access to its customers. In addition, BGP messages are not verified, and hence we believe that service denial at the routing level should be investigated thoroughly.

4 Related Work

BGP was not designed with protection and security in mind and that makes it vulnerable to many attacks. The study in [8] 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. For instance, BGP does not check that the message is not modified (integrity) or being replayed, i.e., an old message is sent when it is not valid anymore (freshness). Moreover, it does not verify the legitimacy of the originator of the message (origin authentication). 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 [7,8]. A more recent and comprehensive survey of BGP security issues can be found in [9]. 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: (a) uncertainty of the mapping between the IP prefixes and the AS numbers for the ASes that manage them; (b) the utilization of the Transmission Control Protocol (TCP) as the underlying transport protocol; and (c) 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 [10].

Nordstrom and Dovrolis [11] 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 *supervision* which is similar to the redirection attack, but with the intention to modify the data (i.e. eavesdroping) and then forwading 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 [12]. 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 2 h [11,13,14]. 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 [13,14]. There are many proposed counter measures against these types of attacks identified in [11], 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 [11]. The second solution is the use of a Secure Border Gateway Protocol (S-BGP) [15]. Although this method can provide high security against attacks, it adds high overhead on the internet. Wang and Wang [16] 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 standarization front is also briefly presented in [17].

The prefix hijack attack can be addressed using a Prefix Hijack Alert System (PHAS) [18]. PHAS is an email notification system that alerts a prefix owner whenever there is a change in the origin AS that owns the prefix. Everyday 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 [18]. The system examines the data collected in RouteViews [19] and notifies the prefix owner about any possible hijack.

Similarly, Zheng et al. [20] 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.

Hu and Mao [21] 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 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 like 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 [22] 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 [23]. The study considers approaches of controlling outgoing and incoming traffic when the Regional ISP is multi-homed. Usually, this control is either

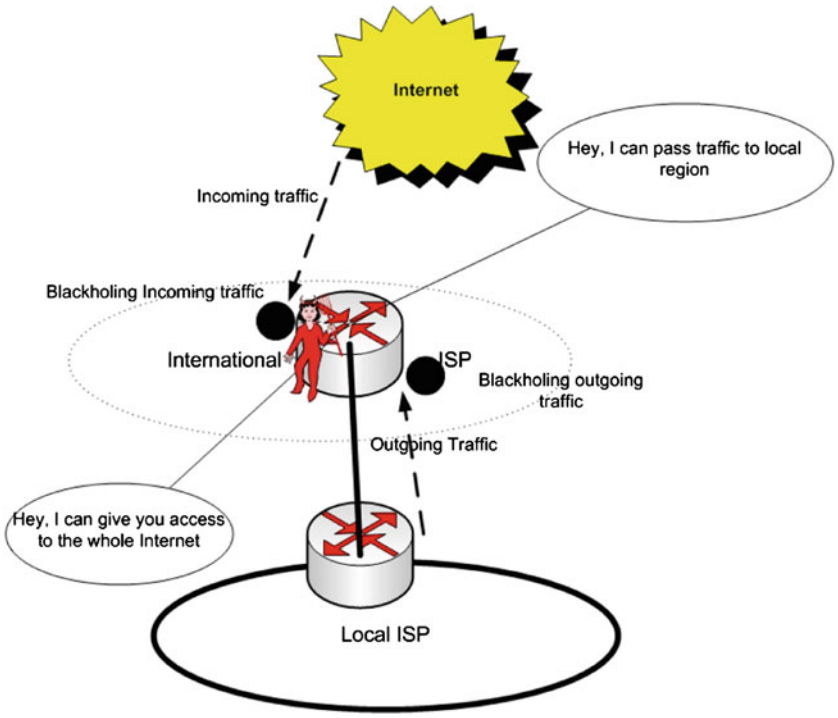

Fig. 2 Malicious IISP blocking

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 [24]. 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. In this paper, a virtual peering solution is adapted to solve the problem of malicious internet blocking. More details about virtual peering are described later in this paper.

5 Problem Description

Local ISPs which provide internet reachability for their customers are getting internet access through IISPs. These IISPs are paid in customer–provider relationship to provide internet access. The IISPs can make use of the routing protocol they run to maliciously block access to some services. These IISPs can also claim that they have a route to destinations inside the local region while blackholing the traffic destined to the prefix owned by the local region. Therefore, a malicious IISP can be defined as a seemingly legitimate provider that advertises a prefix of a local region and advertises internet prefixes to the local region with the ill intentions of enforcing internet isolation. This isolation is achieved by blackholing traffic originating or destined to the local region. Figure 2 depicts the described scenario. This paper focuses on proposing and evaluating solutions to bypass the malicious IISP despite its false advertisements. Blockage detection is not part of this work, however it may be simple to design and implement blockage detection mechanism as per [21].

Two major approaches exist to solve our problem. One approach is to hide one's identity using techniques such as NATing, tunneling, etc. Another approach is to have more than one IISP, i.e. multihoming, and to control the traffic either using an overlay protocol, fine tuning, or modifying the BGP protocol to perform traffic engineering. The latter is used to direct the outgoing and incoming traffic so that they do not pass through the malicious IISP. We will concentrate in this paper on the BGP-based solutions and not on NATing or pure tunneling. Figures 3 and 4 illustrate all of these solutions. It should be noted that if the IISP is only denying access service without false advertisements with the intention to blackhole legitimate traffic, then the existence of another IISP that is not denying the access will solve the problem. This is because all incoming and

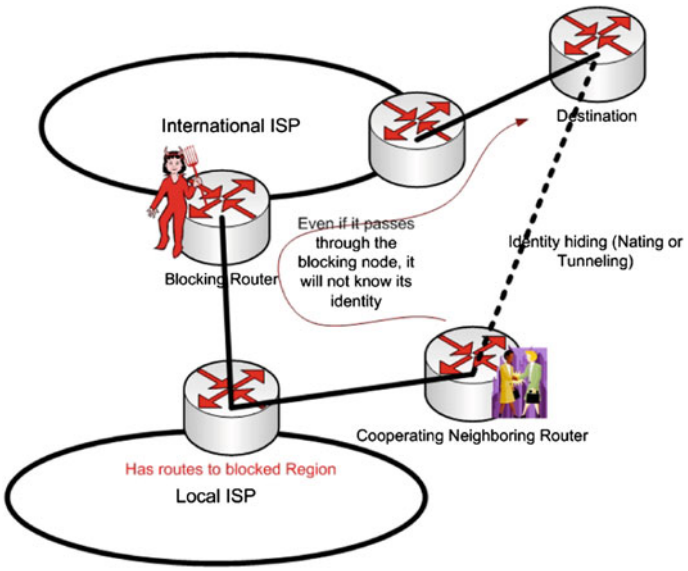

Fig. 3 Identity hiding techniques

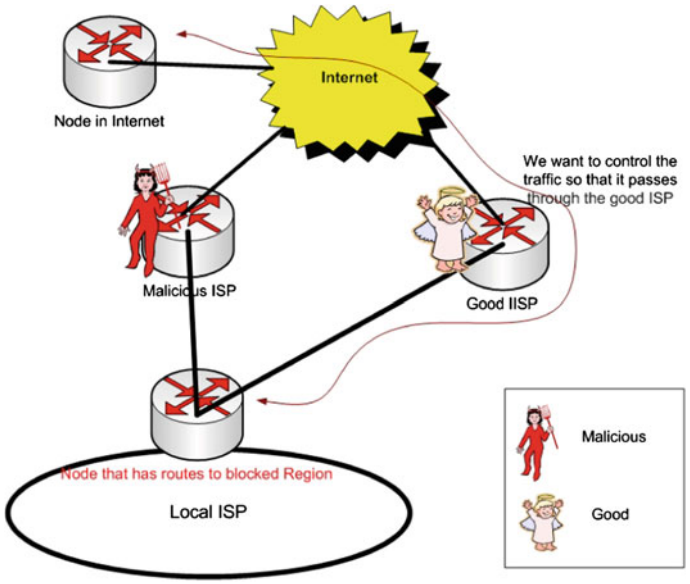

Fig. 4 Traffic engineering-based techniques

outgoing traffic will pass through the other IISP since it is the only provider that advertises the prefixes of the local region and a BGP session can be established with it alone. However, a solution needs to be investigated if the IISP is advertising the prefix of a specific region to the internet and/or advertising internet prefixes to this region while blackholing traffic originating or destined to the local region.

In Fig. 3, the local ISP may be aided by a cooperating neighboring router to set up a tunnel, or to use the path through a neighbor while changing the source address (NATing) in order to send traffic to destination. In the same way, the destination should direct its traffic through the cooperating neighboring router.

Figure 4 depicts traffic engineering approaches that focus on directing the traffic, both incoming and outgoing, through the good or non-malicious IISP. This may be performed through configurations of the BGP protocol, and may also be done through some cooperative algorithms [22].

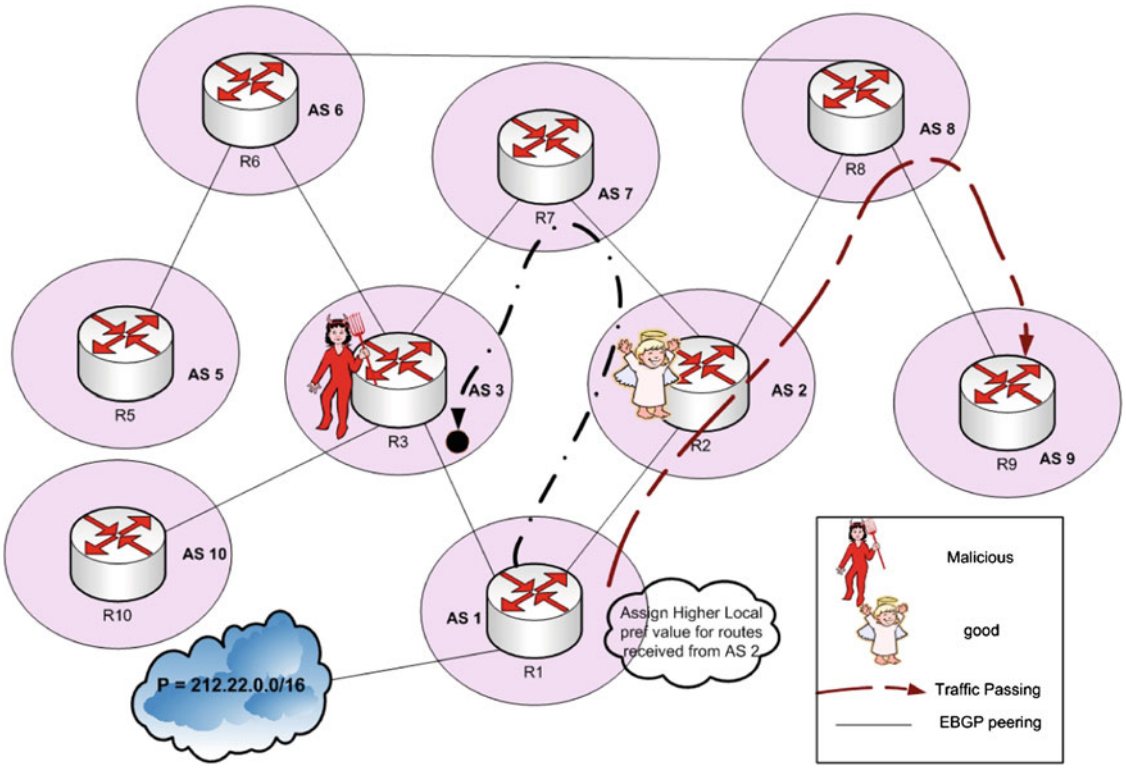

Fig. 5 Control of outgoing traffic using BGP tuning

6 Methods to Combat Internet Service Denial

6.1 BGP Tuning

BGP Tuning means to use the available BGP policies to modify the BGP selection process to enforce the selection of the intended route as the best one. BGP tuning has been used as way for traffic engineering [25]. In order to make sure that the traffic passes through the good IISP, the outgoing traffic needs to be directed through the good IISP and the incoming traffic needs to be influenced to select the good IISP as the best path in its BGP selection process. Therefore, we will be looking at ways of controlling the outgoing and incoming traffic to direct routes through the non-malicious IISP.

6.1.1 Control of Outgoing Traffic

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 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 5 depicts the use of the `local-pref` attribute to control the outgoing traffic.

As shown in Fig. 5, 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. Setting up a tunnel or using any identity hiding technique can help to solve this specific scenario.

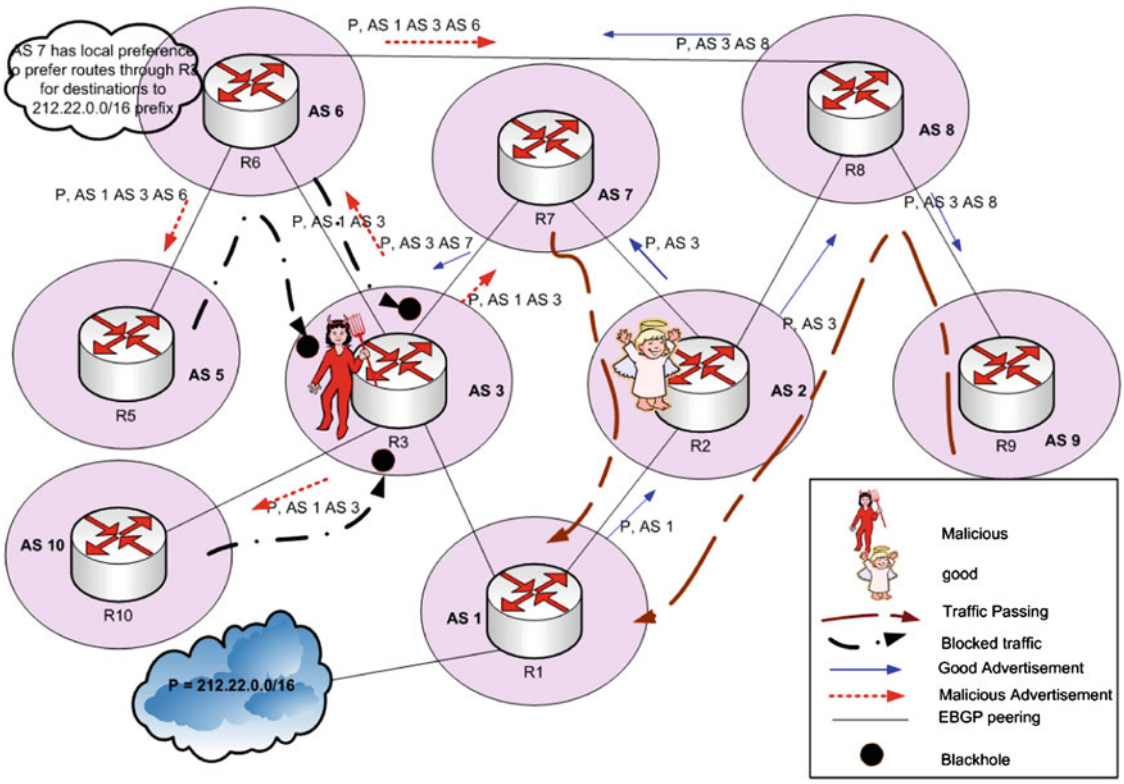

Fig. 6 AS Path shortening

6.1.2 Control of Incoming Traffic

Three ways to control traffic using BGP tuning are 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 techniques can be combined and used together rather than individually.

AS Path Shortening: 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 [26]. If the good IISP sends an announcement of our 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 Fig. 6, 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 `ASPath` 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 Fig. 6.

As shown in Fig. 6, 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.

More Specific Announcement: 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 7 depicts this technique.

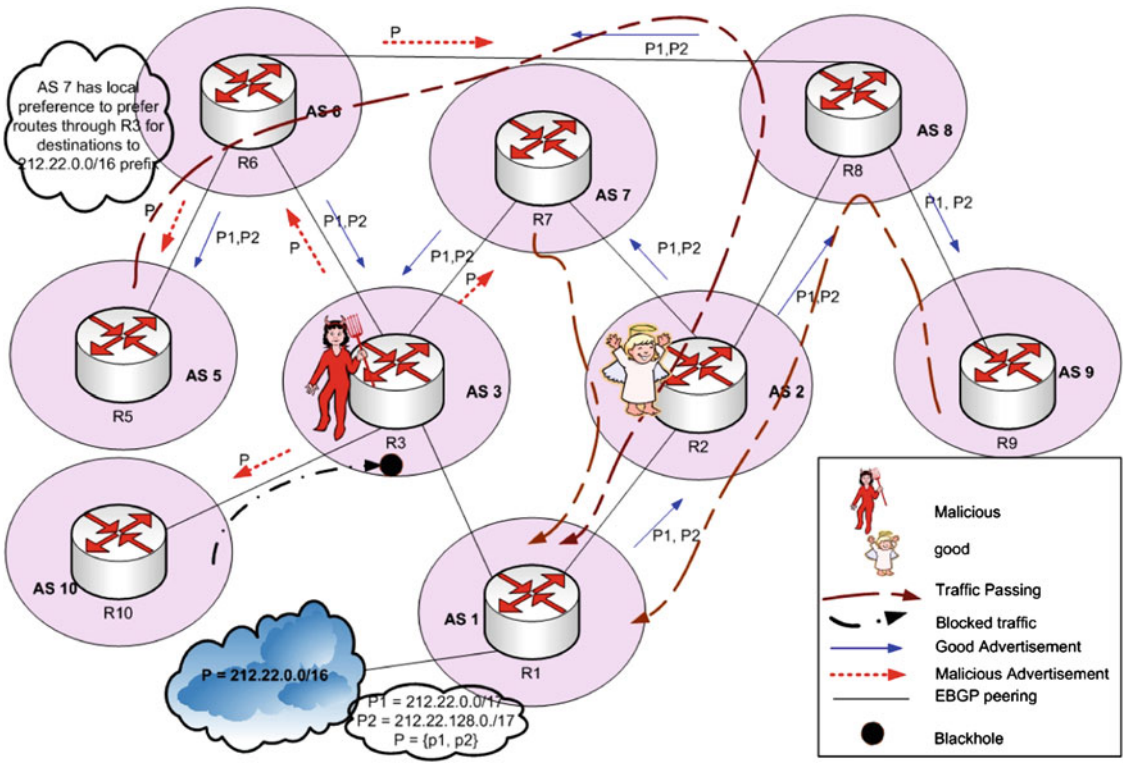

Fig. 7 More specific prefixes announcement

As shown in Fig. 7, 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. Of course, 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.

The Use of Communities: The third scheme 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 8 shows the use of community to control the traffic.

As shown in Fig. 8, 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. Actually, there are certain community numbers that tell the AS to set a higher local preference for certain paths. Therefore, the router has to make sure that the routers will do this in case it advertises the paths with this community number. 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.

6.2 Virtual Peering

Quoitin [22] proposes the use of virtual peering to deterministically control the incoming traffic through one of the providers. He used virtual peering to achieve load balancing of incoming traffic among providers and to reduce the latency by choosing paths that have the lowest delay. Since this method controls incoming traffic, we adapt it in order to direct the incoming traffic through the non-malicious IISP. The main contribution of [22] is the automation of setting up the virtual peering using a Virtual Peering Controller (VPC) that manages

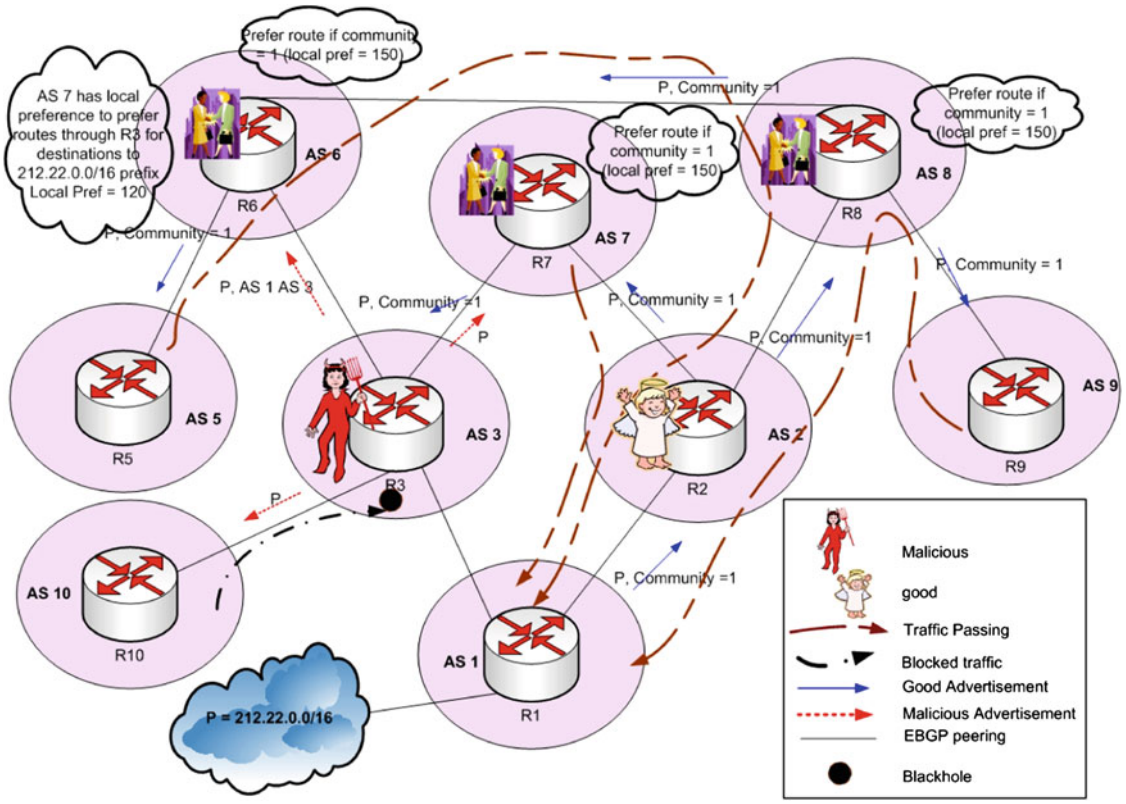

Fig. 8 Control the traffic through the use of communities

virtual peering establishment and removal. Figure 9 depicts the use of virtual peering to solve the malicious IISP blocking.

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 our 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 IISP, 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 IISP. 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 IISP.

6.3 Virtual Transit (Hijack the Hijacker)

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, therefore a number of modifications of virtual peering

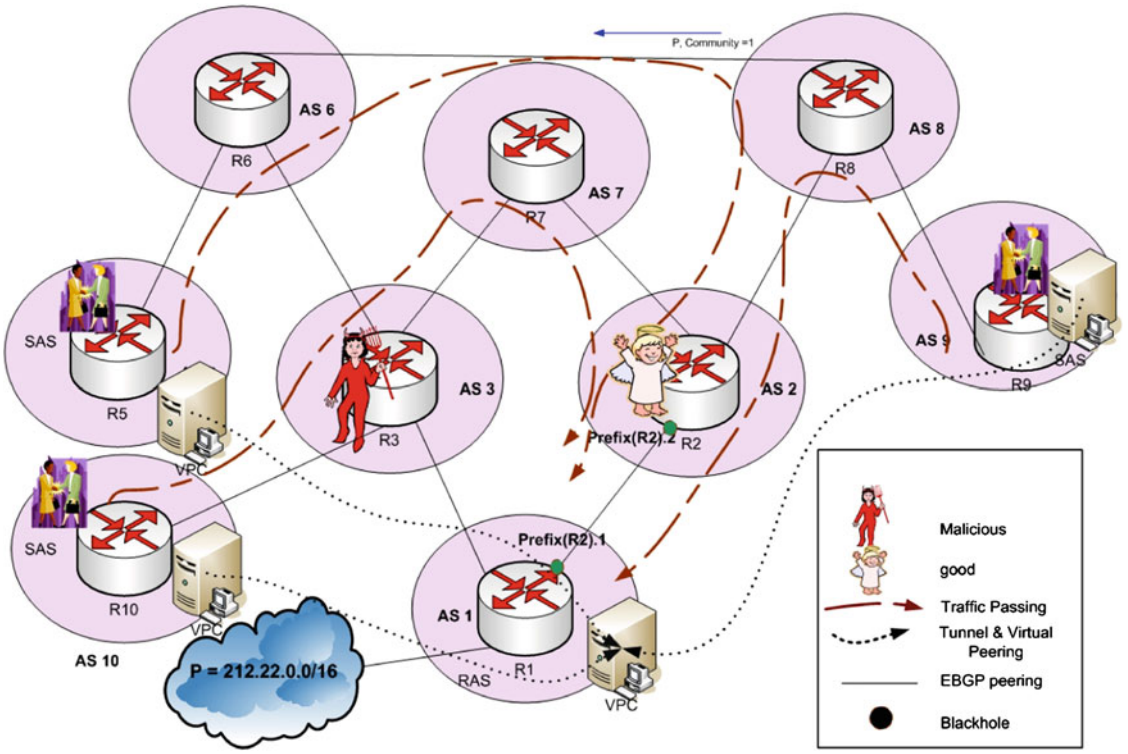

Fig. 9 Virtual peering-based solution

can be done to make the solution scalable. We refer to the modified solution by the 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 Fig. 10.

As shown in Fig. 10, 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 prefixes 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 idea of virtual transit can also be called 'hijack the hijacker'. To clarify this, we can see that R3, a seemingly legitimate provider, is advertising a path to the local region to attract traffic with the ill intention of blackholing the traffic. The advertised path is superficially valid because AS3 is a provider for the local region and is geographically close. The Internet Routing Registry (IIR), a distributed database maintaining routing information and priorities, may even list AS3 as the legitimate ISP for AS1. On the other hand, the routers R10, R5, and R9, though are not obliged to advertise paths leading to the local region, they are doing so even with their distant geographical location and lack of physical connection to the local region. These routers (i.e. R10, R5, and R9) are cooperating with the local region to provide connectivity despite the malicious activity of R3. It is worth mentioning that with some implementation of secure BGP, this incongruity may identify the cooperative routers as hijackers, while it accepts the advertisements by the malicious router. Therefore one may consider this proposed mechanism as an ethical hijacking of the hijacking provider in order to counteract its malicious actions and gain access to the internet.

7 Analysis and Discussion

The use of local preference to control the outgoing traffic is sufficient for most cases except for the case when the route must pass through the malicious provider. A tunnel might be used in order to hide the identity in this

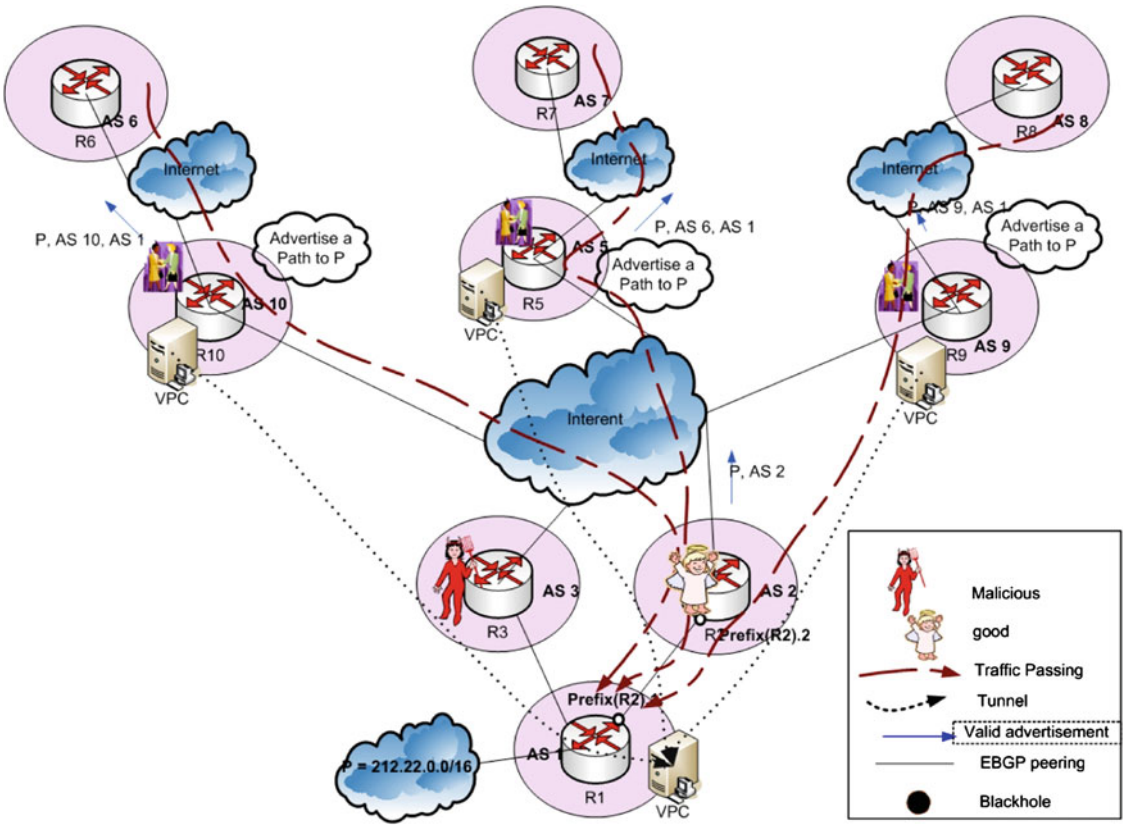

Fig. 10 Virtual transit scheme

case and complete the solution. The use of `AS-Path` shortening can attract traffic to the good IISP. However, major flows of the traffic will probably not be attracted. Examples include those routers that have a local preference to prefer the malicious route, those routers that can still see the path through the malicious route as a shorter path depending on their location, and those routers where the malicious router is the only provider. The use of more specific prefixes can attract more traffic through the good IISP. There is a limit, however, on the length of the prefix, and those that exceed this prefix can be filtered out by routers [22]. In addition, a disadvantage is that traffic will be blackholed when the only path to the destination must pass through the malicious router.

Virtual peering combines the benefits of traffic engineering techniques and hiding one's identity. Moreover, it provides a deterministic technique for the traffic to be forwarded to its destination through a specific provider. However, it requires cooperation between the sources of traffic and the destination. This means that every source of traffic must do virtual peering with the destination in order to reach the destination through the good IISP. The rationale behind having a virtual transit is to make a limited number of virtual transit routers that advertise a prefix in a short `AS-Path`. These virtual transit routers will attract almost all of the internet traffic if they are distributed around the globe and a sufficient number of them are installed. However, using virtual transit does not absolutely guarantee that all the traffic will pass through the good IISP and will not be filtered. For example, if a router has the malicious IISP as its only direct provider, then the router's traffic will be filtered out.

In Table 1, we provide a comparison amongst all the methods considered in this study 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.

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 method that achieves no filtering of traffic is virtual peering because a tunnel is established to an IP address that does not belong to the blocked prefix between every source of traffic and the destination; this tunnel hides the identity in case the traffic flows pass through the malicious IISP. For

Table 1 Comparison between methods

	BGP tuning			Virtual peering	Virtual transit
	AS-Path shortening	More specific prefixes	Communities		
Filtering the traffic	Medium	Small	Small	No	Very small
Setup overhead	Small	Small	Medium	High	High
Communication overhead	No	No	No	Medium	Medium
Difficulty to combat the method	Easy	Easy	Easy	Difficult	Difficult
Scalability	High	High	High	Small	High

the case of virtual transit, almost all of the traffic will not be filtered because most likely the traffic will be attracted by one of the distributed cooperative routers that establish a tunnel to the affected routers. However, a very small amount may be filtered if one source of traffic is attracted by the malicious ISP because it may be the only provider. For BGP tuning, if it happens that the traffic needs to pass through the malicious router, then the traffic is filtered out.

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 high in virtual peering and virtual transit, relative to BGP tuning, because of the establishment and removal mechanisms of a virtual peer and virtual transit connections. This overhead is relatively medium for communities because configurations on cooperating routers are still needed. The setup overhead for `AS-Path` shortening and more specific prefixes is very small since the configuration is needed only on one (the RISP) or two (the RISP and the good IISP) routers.

In contrast, communication overhead refers to the number of messages that need to be exchanged between routers before the method is effective. Only virtual peering and virtual transit have some communication overhead because of the establishment of the multi-hop BGP connection. For BGP tuning-based methods, 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. Although the malicious IISP can also make virtual peering and virtual transit, it is however very difficult to know that the traffic is destined to the blocked prefix and that these techniques are used, since that traffic is apparently not sent to the prefix that is hijacked.

Finally, in terms of scalability which refers to the easiness of extending the method or using it for the entire internet, BGP tuning techniques provide the most scalability since the configuration will affect the decision taken by the traffic in the internet without any additional connections. In the case of virtual peering, if the blocked local region wants to direct all the traffic of the internet to the destination through the good IISP, then a virtual peering connection needs to be established between all the sources of traffic and the destination. Virtual transit dramatically reduces the number of needed established tunnels while having almost all of the traffic of the internet directed to the destination through the good IISP.

8 Validation and Proof of Concept

While the purpose of this paper is to introduce the problem of internet outage due to the intentional or unintentional actions of the IISP and provide a *qualitative* analysis of the proposed solutions, we present in this section a brief validation and proof of concept for the BGP-based solutions outlined earlier. A full validation of all proposed solutions requires the consideration of a variety of network configurations/topologies and the experimentation of all solutions for each one of these network configurations. In addition, these experimentations may be required for different types of services and traffic. This full-factorial suite of experimental setups and evaluations is outside of the scope of this paper.

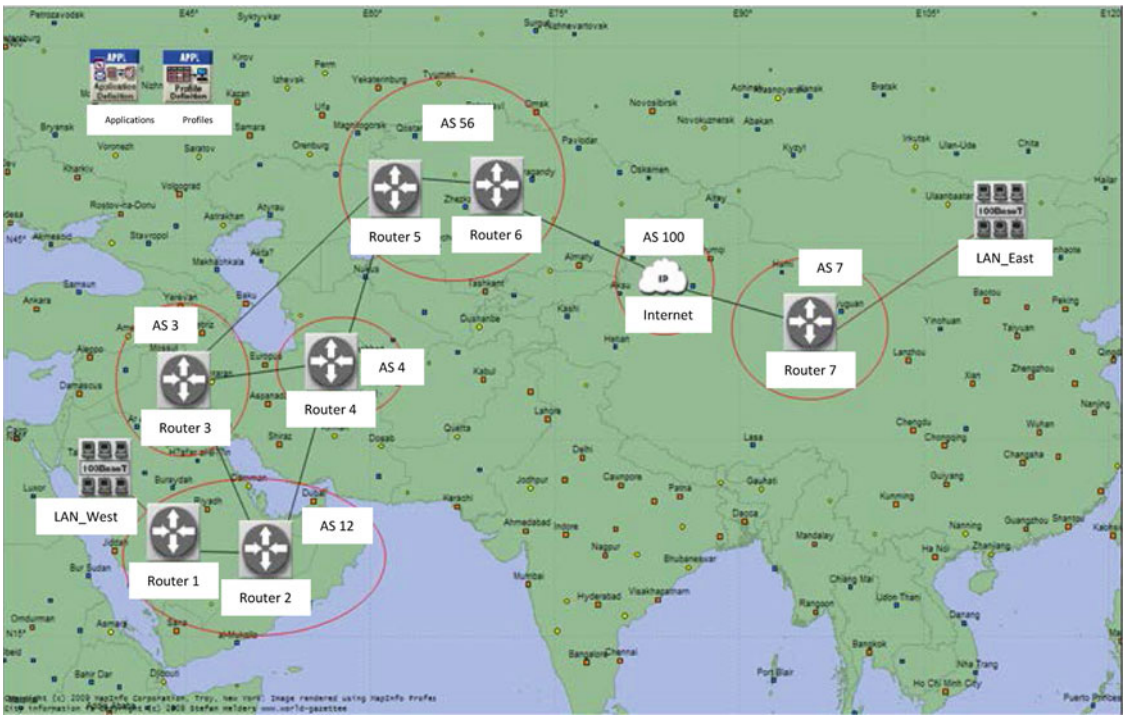

Fig. 11 OPNET simulation model

In the following two subsections, we present an experimental setup that is implemented using OPNET for testing the BGP-based solution and provide a sample of the obtained simulation results.

8.1 OPNET Simulation Model

OPNET is a network simulation and evaluation tool that is widely used for commercial as well as academic and research purposes (www.opnet.com). The tool has full and detailed implementation of most, if not all, TCP/IP protocols including the BGP with the capacity to modify or reprogram the behavior of these protocols. The latter reason along with OPNET's reputation and support in the research community made OPNET the natural candidate platform for simulating and evaluating the BGP-based solutions for circumventing the internet outage scenarios described earlier.

Figure 11 depicts the experiment setup implemented in OPNET. Assume the region of interest, designated by AS12, is connected to the internet through two international ISPs: AS3 representing the main IISP, and AS4, representing the secondary IISP. The IISPs AS3 and AS4 provide connectivity for the region of interest to other ASs such as AS6, AS100, and AS7. Our main objective is to evaluate the performance of traffic before and 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 consider the performance of three types of traffic: Hyper Text Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Voice Over IP (VOIP). Finally, the designed experiment accounts 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 Fig. 11, 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: Path Shortening, More Specific Prefixes, and the Use of Communities, one at a time. The internal code implementing BGP for 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 [27].

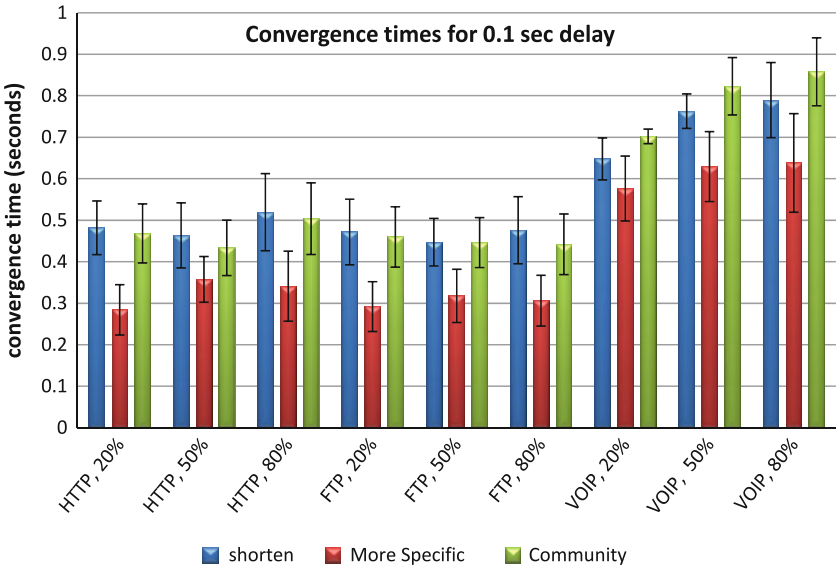

Fig. 12 Convergence time for BGP-based solutions assuming 100 msec average internet delay

8.2 Sample Results

Of particular interest amongst the possible performance figures is the convergence time for the BGP as the new path is setup through the cooperating IISP, which is AS4 in our experiment, replacing the old path through the malicious IISP AS 3. This 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 12 shows the convergence time for the case where the internet average delay is equal to 100 ms 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 s with the BGP solution based on more specific prefixes having the least convergence times. This is true for the considered 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 Fig. 11). 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 s, as opposed to the 100 ms assumed earlier, and the convergence times were ranging from 15 to 50 s. The convergence times patterns observed for the 5 s internet delay are consistent with those for the 100 ms scenario. Similar to the previous case of 100 ms 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.

9 Conclusions

In this paper, three techniques to overcome the problem of internet blocking are discussed. This blocking may be intentional as the case of internet embargo imposed by the international internet service provider (IISP) or unintentional blocking as a result of a security breach at the IISP. These techniques are BGP tuning, virtual peering, and virtual transit techniques. The setting of `Local-Pref` attribute to higher value can easily provide control of outgoing traffic. For the control of incoming traffic, three BGP tuning-based techniques are investigated; namely, `AS-Path` shortening, advertising more specific prefixes, and thirdly the use of community.

If the good IISP advertises the prefix after removing the AS number of the local region, the `AS-Path` is shortened by one which influences the decision process of other routers to favor the route through the good IISP. Since an IP router, before entering the selection of the best route of BGP, searches for the longest prefix match, the advertisement of more specific prefixes will attract more traffic than the other two kinds of BGP tuning techniques. However, a malicious IISP can easily advertise more specific prefixes to attract the traffic of interest. Moreover, routers on the internet are encouraged to remove from their routing tables prefixes subnet values longer than a certain predefined value [22]. The use of communities requires cooperation with other ASes to configure their routers to prefer routes with certain community number. BGP tuning is the simplest approach amongst those discussed in this paper, and can be easily achieved through configuration of the routers of interest.

The second method is the use of virtual peering which provides the most conservative method for controlling the incoming traffic. One can also use virtual peering in the opposite direction to hide the identity of the traffic originator, or have a bidirectional tunnel. Virtual peering provides a deterministic way to control the incoming traffic through the good IISP; however, it requires every source to establish a virtual peering connection.

Virtual transit, a hybrid technique, is a compromise between the first solution and the second solution. Virtual transit utilizes a number of cooperative nodes that advertise a prefix with a shorter `AS-Path` and can advertise more specific prefixes to attract even more traffic while establishing a tunnel with the local region through the good IISP. This provides a deterministic control of traffic through good IISP once traffic arrived to a cooperative node. The virtual transit technique does not require that all sources of traffic make a tunnel to the prefix of interest, and therefore it increases the scalability of the solution. Moreover, it has the potential to attract almost all of the traffic of the global internet.

This paper also presents a brief validation and proof of concept for the BGP-based solution utilizing OPNET simulations. A sample result of BGP convergence time is shown for various network configuration and traffic types.

Acknowledgments The authors would like to acknowledge the support provided by King Abdulaziz City for Science and Technology (KACST) through the Science & Technology Unit at King Fahd University of Petroleum & Minerals (KFUPM) for funding this work through project No. 08-INF97-4 (Project name: Internet Access Denial by International Internet Service Providers: Analysis and Counter Measures) as part of the National Science, Technology and Innovation Plan. The authors also acknowledge Bruno Quoitin for the useful discussions on the subject matter.

References

1. Internet World Stats (2009). http://www.internetworldstats.com/stats.htm
2. Abu-Amara, M.; Mahmoud, A.S.; Azzedin, F.; Sqalli, M.H.: Internet access denial by international internet service providers: analysis and counter measures, Research proposal for National Science, Technology and Innovation Plan (NSTIP), Project # 08-INF97-4, submitted to King Abdulaziz City for Science and Technology (KACST), Riyadh, KSA, April 2008
3. Drummond, D.: A new approach to china. The Official Google Blog (2010). http://googleblog.blogspot.com/2010/01/new-approach-to-china.html
4. Finkle, J.; Bartz, D.: Twitter hacked, attacker claims Iran link (2009). http://www.reuters.com/article/idUSTRE5BH2A620091218
5. Wan, T.; Oorschot, P.C.V.; Kranakis, E.: A Selective Introduction to Border Gateway Protocol (BGP) Security Issues, Carleton University (2005)
6. Rekhter, Y.; Li, T.; Hares, S.: IETF-A Border Gateway Protocol 4 (BGP-4) (2006). http://www.ietf.org/rfc/rfc4271.txt
7. Butler, K.; Farley, T.; McDanial, P.; Rexford, J.: A Survey of BGP Security, AT&T Labs, Research, Florham Park, NJ, Technical Report TD-5UGJ33 (2005)
8. Murphy, S.: IETF -BGP Security Vulnerabilities Analysis (2006). http://www.ietf.org/rfc/rfc4272.txt
9. Butler, K.; Farley, T.; McDaniel, P.; Rexford, J.: A survey of BGP security issues and solutions. In: Proceedings of the IEEE, vol. 98 (1), pp. 100–122 (2010)
10. Nicholes, M.O.; Mukerjee, B.: A Survey of Security Techniques for the Border Gateway Protocol (BGP). IEEE Commun. Surv. Tutorials **11**(1), 52–65 (2009)
11. Nordstrom, O.; Dovrolis, C.: Beware of BGP attacks. ACM SIGCOMM Comput. Commun. Rev. **34**(2), 1–8 (2004)
12. Villamizar, C.; Chandra, R.; G.R.: IETF (1998). http://www.ietf.org/rfc/rfc2439.txt
13. Mahajan, R.; Wetherall, D.; Anderson, T.: Understanding BGP misconfiguration. In: Proceedings of ACM Sigcomm, pp. 3–16 (2002)
14. Farrar, J.A.: Merit Network Email List Archives (2001). http://www.merit.edu/mail.archives/nanog/2001-04/msg00209.html
15. Kent, S.; Lynn, C.; Mikkelson, J.; Seo, K.: Secure Border Gateway Protocol (S-BGP). IEEE J. Sel. Areas Commun. **18**(4), 582–592 (2000)

16. Wang, N.; Wang, B.: AT: an origin verification mechanism based on assignment track for securing BGP. IEEE Int. Confer. Commun. (ICC'08) 5739–5745 (2008)
17. Ortiz, S.: Securing the Internet's routing infrastructure. IEEE Comput. **42**(4), 21–23 (2009)
18. Lad, M.; Massey, D.; Pei, D.; Wu, Y.; Zhang, B.; Zhang, L.: PHAS: A Prefix Hijack Alert System. In: Proceedings of the 15th conference on USENIX Security, Vancouver (2006)
19. Oregon, U.o.: The Route Views Project. http://www.routeviews.org/
20. Zheng, C.; Ji, L.; Pei, D.; Wang, J.; Francis, P.: A light-weight distributed scheme for detecting IP prefix Hijacks in real-time. In: SIGCOMM'07, Kyoto (2007)
21. Hu, X.; Mao, Z.M.: Accurate real-time identification of IP hijacking. In: IEEE symposium on security and privacy (2007)
22. Quoitin, B.: BGP-based interdomain traffic engineering. Ph.d. Dissertation, Universite catholique de Louvain, Louvain-la-Neuve, Belgium (2006)
23. Quoitin, B.; Uhlig, S.: Modeling the routing of an Autonomous System with C-BGP. IEEE Netw. **19**(6), 12–19 (2005)
24. Quoitin, B.; Bonaventure, O.: A cooperative approach to interdomain traffic. In: Proceedings of the 1st conference on next generation internet networks traffic engineering, Rome, Italy (2005)
25. Quoitin, B.; Pelsser, C.; Swinnen, L.; Bonaventure, O.; Uhlig, S.: Interdomain traffic engineering with BGP. **41**(5), 122–128 (2003)
26. Chang, R.K.C.; Lo, M.: Inbound traffic engineering for multi-homed ASes using AS path prepending. In: Network Operations and Management Symposium, vol. 1, pp. 98–102 (2004)
27. AlRefai, A.: BGP-based solutions for international ISP blocking. Master Thesis submitted to Deanship of Graduate Studies, King Fahd University of Petroleum and Minerals (2010)