Anurag Porwal, Rohit Maheshwari, B.L.Pal, Gaurav Kakhani / International Journal of Engineering Research and Applications (IJERA) ISSN: 2248-9622 www.ijera.com
Vol. 2, Issue 2,Mar-Apr 2012, pp.178-182

The Secure interdomain Routing and Forwarding

Anurag Porwal[1], Rohit Maheshwari[2], B.L.Pal[3], Gaurav Kakhani[4]
[1]Lecturer at Mewar University, Gangrar
[2]Asst.Professor at Mewar University,Gangrar
[3]Asst.Professor at Mewar University,Gangrar
[4]Lecturer at Mewar University,Gangrar

Abstract

The Internet consists of thousands of independent domains with different, and sometimes competing, business interests. However, the current interdomain routing protocol (BGP) limits each router to using a single route for each destination prefix, which may not satisfy the diverse requirements of end users. The Border Gateway Protocol (BGP) is the protocol for the interconnectivity of the Internet. the security issues that surround interdomain routing. Thus, it is required to provide a highly secure protocol to keep the normal operation of the Internet. Although the Internet's routing system has serious security vulnerabilities, none of the existing proposals for a secure variant of BGP has been successfully deployed in practice. This paper explains data forwarding process in interdomain routing.

1. Introduction

The Internet consists of thousands of independently administered domains that rely on the Border Gateway
Protocol (BGP) to learn how to reach remote destinations. the protocol requires each router to select a single "best" route for each destination prefix from the routes advertised by its neighbors. This leaves many ASes with little control over the paths their traffics takes [2].
Recent research has considered several alternatives to single path routing, including source routing and overlay networks. In source routing, an end user or AS picks the entire path the packets traverse. In overlay networks, packets can travel through intermediate hosts to avoid performance or reliability problems on the direct path [1].

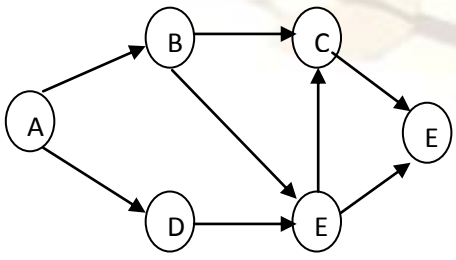

Fig 1. Single path routing to AS F

The current interdomain routing protocol Border Gateway Protocol (BGP4), have two requirements for interdomain routing are at the root of BGP's bewildering complexity [1]:

1. Policy- ASes have business relationships with one another must cooperate to achieve global reachability. Operators use routing policies to control the flow of outbound traffic and specify which routes are advertised to neighboring networks under what conditions.

2. Scalability- Routing protocols must scale with increasing network size. The main mechanism to achieve scalability is aggressive aggregation of routing information, including destination prefixes. Each router in the AS receives a summary of routing information from its route reflector, rather than relying on a "full mesh" of routers communicating with each other.

Interdomain Routing Model

We now define a model of interdomain routing that scopes our discussion. ASes exchange routing information via exterior routers at one or more locations. Each AS has interior routers that obtain information about external routes from the exterior routers. Given any set of available routes to a destination d, Sd, each router selects a best route, $rd = \lambda (S d)$. Every router must have a preference relation for all a, b $\in$ Sd: either a $\prec$ b, or b $\prec$ a. Each router applies an export policy to determine the neighboring routers to which it should readvertise its current best route in Sd. This model captures many features of BGP. The preference function, λ, incorporates the BGP decision process and the effects of routing policies on route selection. In BGP, each router propagates only the best route (or nothing) for a destination to a neighboring router. The notion of "exterior" and "interior" routers reflects the general property that some routers in a network will exchange routes in other administrative domains and others will not; it also allows for distinctive behavior in the two realms: eBGP for exchanging routes between ASes, and iBGP for exchanging routes between interior routers. The model also recognizes that each router in an AS may make different decisions, as in BGP. This model also reflects several limitations of the current interdomain routing policy. BGP does not permit policies that dictate which ASes must be and must not be traversed en route to a destination [1].

Routing Architectures

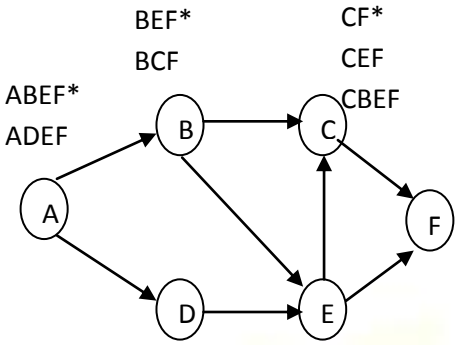

Fig 2.(a) Current BGP protocol

DABEF ECF

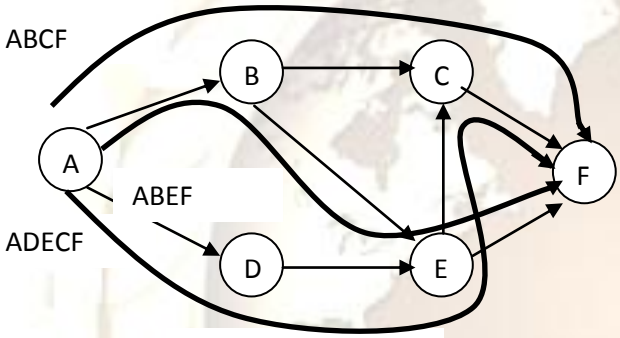

Fig 2 (b) Source routing

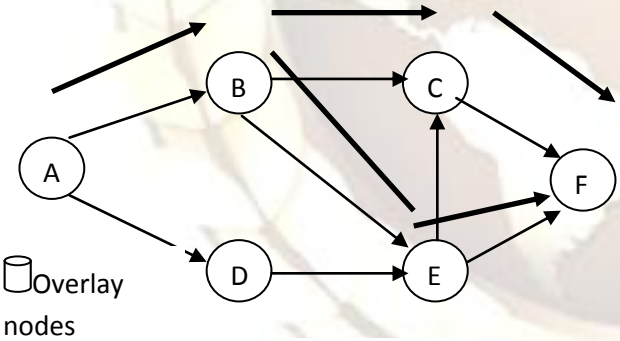

Fig 2 (C) Overlays Networks

Inter AS routing proposals

In this section, we present an overview of the current BGP protocol, source routing, and overlay networks. Here we represent each AS as a single router, as illustrated in Figure 2 where five ASes are selecting routes to a destination in AS F. In BGP, each AS selects a single best route (indicated by an asterisk) and advertises it to all neighbors. In source routing, each end host has complete knowledge of the entire topology and can choose whatever paths it wishes. In overlay networks, several overlay nodes connect to the physical network to form a virtual topology; each node can direct traffic through other overlay nodes en route to the destination [2].

BGP has several features that limit flexibility In path selection [2]:

- Destination-based: BGP distributes reachability information by performing a longest-prefix match on the destination address. As such, packets from different sources going through the same router would follow the same downstream path.
- Single-path routing: A router learns at most one BGP route from each neighbor and must select and advertise a single "best" route. This limits the number of paths advertised and poses severe restrictions on flexibility.
- Path-vector protocol: BGP is a path-vector protocol where routers learn only the AS paths advertised by their neighbors. This improves scalability at the expense of visibility into the possible paths.
- Local-policy based: BGP gives each AS significant flexibility in deciding which routes to select and export. However, the available routes depend on the composition of the local policies in the downstream ASes also.

Source Routing

In source routing, the end hosts or edge routers select the end-to-end paths to the destinations. The data packets carry a list of the hops in the path, or flow identifiers that indicate how intermediate routers should direct the traffic. Source routing maximizes flexibility, several difficult challenges remain:

• Limited control for intermediate ASes: Under source routing, intermediate ASes have very little control over how traffic enters and leaves their networks. This makes it difficult for intermediate ASes to engineer their networks and select routes based on their own business goals.

• Scalability: Source routing depends on knowledge of the network topology, at some level of detail, for sources to compute the paths. the sources must receive new topology information quickly when link or router failures make the old paths invalid.

• Efficiency and stability: In source routing, end hosts or edge routers adapt path selection based on application requirements and feedback about the state of the network. in some cases, a large number of selfish sources selecting paths at the same time may lead to suboptimal outcomes, or even instability.

Tunnels for Forwarding Data Packets

Under multi-path routing, routers must be able to forward the packets along the paths chosen by the upstream ASes. In MIRO, the two negotiating ASes establish a tunnel for carrying the data packets. The downstream AS provides a unique tunnel identifiers to the upstream AS, independent of which AS initiated the negotiation. In Figure 3(b), when AS A and AS B agree on the alternate route BCF, AS B assigns a tunnel id of 7 and sends the id to AS A. In the data plane, AS A directs the packets into the tunnel and AS B removes the packets from the tunnel and forwards them across the link BC. Then, AS C forwards the packets based on the destination IP

address along the default path to AS F. The upstream AS may apply local policies to direct some traffic along alternate paths

and send the remaining packets along the default path .In Figure 3, suppose BCF has lower latency than BEF. Then, AS A may want to direct its real-time traffic via BCF while sending best-effort traffic along BEF, especially if AS B charges for using alternate routes.

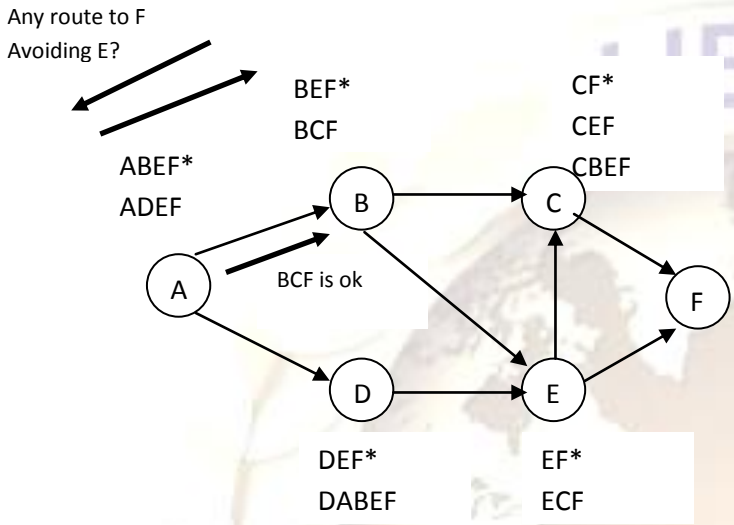

Fig 3 (a) Route Negotiation

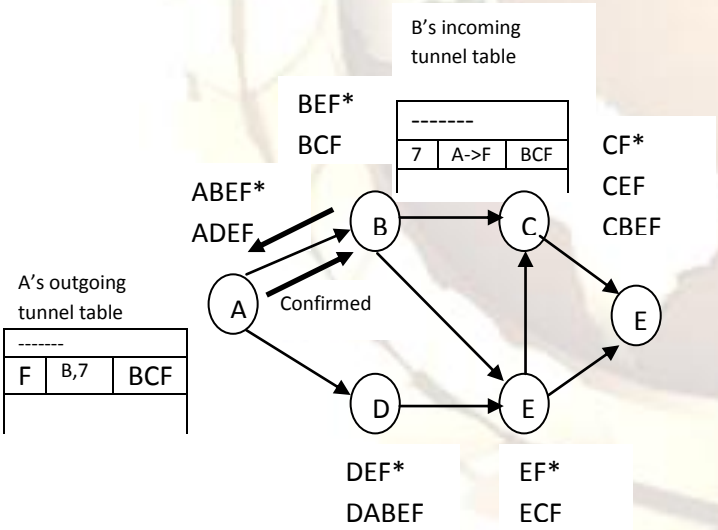

Fig 3 (b) Routing tunnel establishments
Fig 3 Multi Path Routing Examples

Packet Forwarding and Route Setup

The route setup protocol establishes a logical connection along the sequence of PGs that instantiate a policy route. Consequently, when an outbound packet arrives at a PG, the source PG looks to see if the packet is part of a source - destination association that has already been bound to a policy route. In order for the association to match, it must have

compatible type of service, user class identifiers, and other global conditions, in addition to the same source and destination addresses. If the association is bound, then it has already been assigned a unique path ID by the source PG. In this case, the packet is encapsulated in an IDPR data packet header, and the packet is forwarded to the next PG on the already established policy route. The header includes the path ID, which is comprised of the source AD, PG, and a unique number assigned by the source PG. It means, many transport sessions may share a policy route, and consequently will use the same path ID. Finally, if no active policy route applies, the PG invokes the route setup [6].

2. Motivation and Problem Definition:

Eras in which we are living, people of the world are growing themselves against the security of the network. Networks must be secure, on which the users are doing trust and continuously sending their sensitive information. So to make secure the networks I want my contribution through this paper.

Problems:

Policy-Induced Problems

Introducing policy into interdomain routing cause's two main problems [1]:

1. Protocol oscillations
2. Weak security

Protocol Oscillations Instability results from two main causes:

Inter-AS oscillations (caused by policy disputes) and

Intra-AS oscillations (caused by non-monotonic ranking functions).

Policy Disputes:

Because BGP's path selection is based on an AS's local preferences, rather than shortest paths, a group of ASes can have preferences that cause 0BGP to oscillate forever. These "policy disputes" occur because there is no possible path assignment for which at least one AS in the system does not have a better path available thus, that AS would switch to the better route. That act of switching creates a different path assignment that is also unstable.

Even when given stable inputs, BGP might never converge! Griffin et al. showed that, in general, determining whether a set of ASes would experience a policy dispute is an NP-complete problem. They also defined the concept of a "dispute wheel", which describes a circular relationship among a group of ASes where each AS prefers an indirect route via another AS in the group over a direct route to the destination. It might seem that the dispute problem is "solved" because the Gao/Rexford constraints are realistic and they guarantee convergence. We disagree. First, it may be difficult to guarantee that the constraints are satisfied (the proposed

constraints on inter-AS relationships are a global property). Second, there may be legitimate reasons to deviate from the

guidelines: an AS may decide to provide transit between two peer ASes (which violates the Gao/Rexford constraints) as part of a special business relationship.

Weak Security

Interdomain routing involves thousands of competing ASes, each of which sets policy about the routes it is willing to accept and propagate. Interdomain routing security is

particularly difficult because interdomain routing must support complex policy. Today's infrastructure provides scant support for either preventing or detecting invalid routes. BGP does not allow an AS to verify that a route it learns is valid and provides no guarantees about where packets will actually go.

Control-Plane Security BGP does not provide any support for controlling route announcements. Specifically, BGP does not prevent an AS from advertising arbitrary prefixes. One of the most fundamental problems in interdomain routing is determining whether an AS is authorized to announce a certain prefix. S-BGP proposes using certificates to bind IP address space to the AS that owns the space, but this solution requires a public key infrastructure, expensive cryptographic operations, and relatively high message overhead. [1].

Data-Plane Security

Even if an AS could verify that the routes it receives were authentic and policy-compliant, it still cannot verify that packets actually traverse the same ASes as those in the routes AS path. Since policy should ultimately dictate the path that data packets take, we ask [2]:

• Verifying the forwarding path. How can an AS verify that a route's AS path matches the actual forwarding path? A router should reject packets from sources that should not have a valid route through this router to the destination. Deploying packet filters only at routers in large ASes in the Internet "core" could eliminate most of these packets, but an AS must be able to construct these filters in the first place, which would require discovering the routes from the source to that AS.

Scalability Induced Problems

Interdomain routing must scale too many ASes, routers, and destinations. The three main scaling techniques

1. Representing an AS as a single node (e.g., in the AS path),
2. Route reflection, and
3. Prefix aggregation

Reduce overhead by hiding routing information. In this section, we review how this obfuscation causes serious problems (e.g., slow convergence, forwarding loops, persistent oscillation, and network partitions) and pose open questions related to solving these problems.

1. **Missing Topology Details** - BGP abstracts the routing details inside each AS and aggregates information about routes to individual destinations. These techniques allow BGP to scale, but they also make it difficult to determine the cause of a routing update, which can slow convergence, prevent problem diagnosis, hide fine grained information about the reachability of destinations, and reduce an AS's control over incoming traffic.
2.
3. **Missing Routes** - A route reflector selects a single best route for each destination and advertises this route to its clients, obviating the need for each pair of routers in an AS to exchange routes. Route reflectors reduce the number of BGP sessions in an AS and the total number of BGP routes that each router must learn but can cause forwarding loops, protocol oscillations, and partitions.

4. Methodology

Secure and Reliable Interdomain Routing

Interdomain routing is the process by which different ISPs' networks share information about how to reach destinations on the Internet. However, the information contained within BGP is not authenticated, meaning that an attack or mis-configuration by a router anywhere on the Internet can affect the global flow of traffic to any destination, rendering the destination unreachable [4].

While proposals for securing BGP have been around for quite some time, no protocol design has been adopted and deployed on the Internet due to significant adoption hurdles. A major aspect of our work is to explore designs that reduce these deployment hurdles and to explore how different protocol designs affect adoption. We also look at how different approaches (e.g., multi-path routing) can also protect traffic against attacks or errors from routers already on a legitimate path, which BGP cannot handle at all [4].

Improving the Adoptability of BGP Security

Deploying a more secure version of BGP is fraught with adoption hurdles that are tied deeply into the design assumptions of any solution. For example, the efficiency of the cryptographic primitives used to authenticate secure routing data determines whether routers will need to include new crypto acceleration hardware to support secure BGP. The Secure Path Vector (SPV) proposal uses efficient symmetric key cryptography to significantly reduce the cost of signing and verifying routing announcement. The creation of a PKI to establish public keys to authenticate address space ownership and identify ASes is another case where BGP adoption faces a large one-time cost. Our "Grassroots PKI" proposal offers a novel mechanism that lets the PKI start out in a simple manner and grow more secure over time. Finally, the exact type of protection offered by a routing protocol affects the level of protection it provides during partial deployment. [5].

Multi-Path Availability Centric Routing

Unlike traditional secure interdomain routing research, which focuses on cryptographically securing the contents of the BGP protocol to avoid invalid announcement, we explore the possibility of having the infrastructure expose many possible paths and allowing end-hosts to select among those paths to determine which path "works". Since most end-host traffic that needs strong security is already capable of recognizing the valid destination using end-to-end mechanisms like SSL and

IPSec, this approach offers powerful robustness with only minor changes to the infrastructure, and none of the cryptographic and management overhead of securing BGP. We refer to this simple and light-weight approach as Availability Centric Routing, because the infrastructure is focused on making sure at least one legitimate path is available, not on the correctness of all routing information [2].

Inter-domain Routing Security in the Internet

Inter-domain routing in the Internet is managed using BGP4 .This was originally developed for use in a trusted environment, and so provides little security against attackers or misconfiguration. Current BGP operations depend completely on peers trusting one another not to inject bad information into the routing updates. This is coupled with limited filtering (e.g. to filter out advertisements of unallocated address space, and to ensure that downstream customers only advertise their own address prefixes). In addition to such filtering, there is some use of TCP-MD5 to provide integrity protection for the protocol between peer routers [9].

Approaches to BGP security which avoid the use of cryptographic components by relying on BGP policy tools have also been proposed. One solution, pgBGP (Pretty Good BGP), simply adjusts BGP policies to provide some additional cautiousness in accepting new routes. New origin ASs for a prefix are regarded as suspicious for a period of time, and then accepted as normal. This reduces the likelihood of a (short-lived) prefix or sub-prefix hijacking being successful when used in conjunction with appropriate monitoring systems.[9].

Routing Security

"Good housekeeping" practices that are intended to prevent the local routing infrastructure from being subverted. After all, routers represent one of the more vulnerable points of weakness in the entire framework of routing security, and if an attacker can gain control of a router within a network, or gain control of its configuration, then using this platform to inject false information into the routing system is a logical next step. So the first element of the security framework is protection of the routing elements themselves [8].

The vulnerability here to protect the integrity of operation of the routing protocols, is that if a third party can inject packets into the routing protocol exchange, then, at the very least the attacker can disrupt the operation of the routing protocol and cause various forms of disruption and denial of service. There is also the potential to hijack a routing peering session and inject false information into the routing system. So the second element of the security framework is that of protection of the routing protocols [8].

The basic questions that need to be answered in a secure routing framework include establishing the bona fides of the party that originally injected the information into the routing system, as well as the bona fides of the routing prefix itself. This identity information is not of any intrinsic value in itself, but a means to answer the more fundamental question of whether this party has the necessary credentials or permission to inject this information. Have they been authorized to originate a reach ability advertisement for this particular address prefix? The associated question is: Is the address prefix valid? In other words, has this prefix been duly allocated for use through the established address distribution procedures, and is it valid to use this particular address range in the context of the public Internet? [8].

4. Conclusions and future work

In this paper, we aim at providing a comprehensive approach at defining the security issues that concern interdomain routing. Then, by providing a few solutions, this will give us the ability to analyze and compare them. The solutions chosen are Secure-BGP (S-BGP) and secure origin BGP (soBGP).

This choice was made because of the high focus of the research community on these two protocols. They both aim at providing a certain level of security to the de-facto standard of interdomain routing: BGP. To sum up, although BGP was provided with a few security mechanisms, it has not shown that it is safe and secure. Moreover, these mechanisms are independent from the protocol and they represent measures applied only by those who want to. Thus, mechanisms inclusive to the protocol should be designed and implemented. However, research has brought us a few still debatable solutions to this issue. Interdomain routing has shown quite a lot of interest in the last decade. This is due to its importance to many organizations and the whole Internet community in general. The Internet has become a fundamental resource in academic institutions, government agencies and small to large businesses, as well as a vibrant part of our daily lives. This large network of networks requires the interconnection and collaboration of a significant number of autonomously controlled networks. The good functioning of communication in the Internet relies on routing, which is the component that determines feasible paths (or routes) for data to follow from a source to a destination. Today and for nearly two decades, the Internet has seen a new born protocol that could cope with its scale of growth. The Border Gateway Protocol relies on the exchange of messages. More precisely, the routing information provided in tables relies on UPDATE messages exchanged between bordering routers in Autonomous Systems. The way BGP-4 was designed excluded it from all security aspects. This led to an insecure interdomain routing protocol deployed in the entire Internet. BGP has shown many weaknesses and vulnerabilities to malicious behavior. Since BGP requires the use of a TCP session, it inherited all the issues that the Transport Control Protocol has. It became vulnerable to even a larger number of different attacks. BGP can be subject to eavesdropping, replay, message insertion, message deletion, message modification, man-in-the-middle, and denial of service attacks. If the routing infrastructure is attacked and apprehended, it can be used to attack other systems on the Internet such as DNS.

Many countermeasures were built to secure BGP. However, they are not part of the protocol and some of them employ weak security mechanisms. In order to provide a comprehensive

solution for interdomain routing, the security requirements of a well functioning BGP need to be defined. The major three issues that need to be emphasized on are hop integrity (peering session protection), origin authentication of ASes and speakers, and route validation. Many solutions for securing interdomain routing have been proposed. However, majorly only a few have been discussed over the last five years. We covered two of them: S-BGP and soBGP. Both of these protocols have a similar aim which is to protect BGP-4. However, through their design, they seek to secure different parts of the protocol through the use of the same cryptographic primitives.

5. References

1. Some Foundational Problems in Interdomain Routing, Jennifer Rexford AT&T Labs–Research.
2. Multi-path Interdomain Routing, Wen Xu and Jennifer Rexford, Princeton University
3. A New Inter-Domain Routing Architecture, Xiaowei Yang, Member, IEEE, David Clark, Fellow, IEEE, and Arthur W. Berger
4. On Inter-domain Routing Security and Pretty Secure BGP (psBGP) P.C. van Oorschot, Tao Wan, Evangelos Kranakis Carleton University, Ottawa, Canada
5. How Small Groups Can Secure Interdomain Routing, Ioannis Avramopoulos, Martin Suchara, and Jennifer Rexford Princeton University
6. Inter Domain Policy Routing : Deborah Estrin and Martha Steenstrup
7. How Secure are Secure Interdomain Routing Protocols? Sharon GoldbergMicrosoft Research, Michael Schapira Yale & UC Berkeley etc.
8. The ISP Column, An occasional column on things Internet March 2005
9. Inter-Domain Path Authentication in Tactical MANETs, Steffen Reidt and Mudhakar Srivatsa Royal Holloway, University of London IBM T.J. Watson Research Center