

Security of Inter-Autonomous Systems Routing

Iikka Väkiparta
Helsinki University of Technology
Iikka.Vakiparta@iki.fi

Abstract

This article studies the security of inter-autonomous systems. The main focus is on Border Gateway Protocol (BGP), the protocol used to share information between Autonomous Systems, and the security vulnerabilities in it. However it must be kept in mind that most outcomes of the security flaws in the BGP protocol affect the routing service. In addition to attacks against BGP I discuss briefly another type of attack that can be used to launch a denial of service attack or to force sensitive data into the hands of the attacker.

KEYWORDS: BGP, Security, MD5, Routing, S-BGP

1 Introduction

Routing *per se* is not the part that should be secured. The information that is transmitted by the network is what counts. However routing is one critical part in ensuring that the information flows from the sender to the receiver.

The aim of this paper is to look at inter-autonomous systems routing based on its security. Therefore the focus is mainly on BGP, a protocol that is used to distribute routing information between autonomous systems, and other attacks than can be used to launch denial of service attacks or force sensitive information into the hands of the attacker. In order to do this, I begin by defining what is meant by both routing and security. Then by using the defined framework I analyse the security of BGG and the proposed changes to BGP to increase its security. I also briefly discuss the use of link-cuts to affect the routing. I focus on the routing in Internet where we can assume the routing information to be public information and the transmission path is out of the control of the parties sharing routing information. As last I conclude the findings.

2 Definitions

In this section I define the concepts and the framework in which this paper is written.

2.1 Security

There are numerous ways to define security. In this contexts I combine definitions from Gollman [6] and RFC3552 [10]. Gollman [6] defines security by dividing it into three aspects:

- Confidentiality
- Integrity

- Availability

Confidentiality is used to mean that unauthorized users should not be learning insensitive information. [6] Basically this means, that only those individuals for whom the information is intended to, should be able to obtain it.

With **intergrity** we refer to data remaining accurate. It means that only authorized users or systems are able to modify the information and only in a way that it stays available and accurate.

The last item, **availability**, refers to unauthorized prevention of authorized users from using resources or services. Nowadays often discussed attack, *denial of service* -attack is an attack against availability.

RFC3552 [10] enforces also both confidentiality and intergrity. In addition to these it brings forth **peer entity authentication** and **non-repudiation**. Peer entity authentication is an important consideration. In order to ensure confidentiality and intergrity, we must be able to make sure, that the party we are discussing with is who he claims to be. By non-repudiation we mean that we have an ability to demonstrate that the authenticated user has done what he has done.

2.1.1 Internet Threat Model

The threat model describes the capabilities that an attacker is assumed to be able to deploy against a resource [10]. The generally assumed threath model of the Internet expects that the communication endpoints have not been compromised. However, the attacker has nearly complete control of the communications channel by which the endpoints communicate. By this we mean that the attacker can read any Protocol Data Unit (PDU) on this network and undetectably remove, change or inject forged packets on the wire [10]. In this threat model an attacker can also forge packets to appear to be created by the trusted machine.

There exists also two limited threath models derived from the above mentioned threath model. In a **passive attack** we assume that the attacker can read arbitrary packets but cannot write them, whereas in an **active attack** we assume to opposite (write but not necessarily read). The next section describes attacks on network resources based on the this categorization.

2.1.2 Passive attacks

Passive attacks can be used to read information transmitted on the network. With these attacks the user can compromise the confidentiality of the information transmitted.

RFC3552 [10] describes three types of passive attacks: *Confidentiality violations*, *password sniffing* and *offline cryptographic attacks*. Confidentiality violations refer to an attacker reading confidential data while it traverses in the network. Password sniffing is in fact a special case of confidentiality violations where the attacker attempts to acquire passwords that travel in cleartext over the network. Offline cryptographic attacks refer to the attacker capturing a cryptographic message and attempting to break it offline. An example of such an attack is an attempt to break a password that has been encrypted.

2.1.3 Active attacks

In order to be able to perform an active attack the attacker must be able to write packets onto the network. However the attacker must not always be able to read the traffic. These attacks are referred to as **blind attacks**. RFC3552 [10] describes five types of active attacks: *replay attacks*, *message insertion*, *message deletion*, *message modification* and *man-in-the-middle*.

In a replay attack an attacker records messages from the network and retransmits them to the original target. This can be used to cause inconvenience or even gain access to a system. Message insertion attack instead refers to an attacker inserting messages to the traffic. Such messages can be commands inserted into a telnet session or sending TCP SYN packets to a server to mount a denial-of-service attack. Message deletion refers to an attacker removing a message from the network, whereas in a message modification attack an attacker modifies a message in transmission. A man-in-the-middle attack combines the above mentioned techniques. In a man-in-the-middle attack an attacker converts all of the traffic of an conversation to flow through his computer and plays as the conversation endpoint to both of the parties of the conversation.

2.2 Routing

Routing is a fundamental service of Internet. Routing is used to forward data from a source into its destination. Routers exchange information between each other to ensure that routers know about connections to networks and to ensure that routing is efficient. To preserve bandwidth the routing in Internet is partitioned into *autonomous systems* [3]. Between autonomous systems only routers dedicated as *border routers* exchange information, whereas inside an autonomous system all routers participate in sharing information about connections inside the autonomous system. Because of this division different sets of protocols are used inside autonomous systems and between them. Mostly a single protocol, *Border Gateway Protocol (BGP)* and its security enhanced version (Sometimes referred to as *Secure Border Gateway Protocol (SBGP)*), is used between autonomous systems. On the other hand, inside autonomous systems a number of protocols are used [3]. This paper focusses on the protocols used inter autonomous systems, i.e. on BGP.

2.3 Routing Security

As mentioned in section 2.2 routing is a service and by nature routing information is not secret. On the contrary, in order to ensure correct routing of packets routing information should be distributed. However information about some networks inside autonomous systems might be something that the administrator of that network wants to conceal. However, this is not the focus of this article.

When studying routing security it is important to make a distinction between routing *per se* and the information, that is conveyed by routing. Whereas information exchanged between routers to make routing optimal is generally not confidential, the information conveyed through the network often is. Because of this and to ensure optimal routing the availability and integrity of the routing data is a concern. So we have to remember, that by attacking the routing protocols an attacker can launch an attack against the service provided by routing. For example by attacking BGP an attacker can force the information to be transmitted through networks that he can access and thus be able read the information (confidentiality) and/or modify it (integrity). He can also disrupt routing of packets to make it impossible to transmit information from one network to another (availability). This paper focusses on the attacks on the routing protocols but generally the outcomes of the attacks are on the service provided by routing.

3 BGP (Border Gateway Protocol)

BGP is an inter-Autonomous Systems routing protocol. The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems [12]. This information is generally transmitted as a list of autonomous systems (ASs) that the system sending the information can reach. From this information a graph of AS connectivity can be constructed.

Each of the BGP messages contain a 19-octet header. This header contains a 16-octet marker, 2-octet length and a 1-octet type. The marker is a code agreed by both parties. In an OPEN message it is set to all ones but after that it is a code agreed by both parties and it can be used to authenticate parties and to detect loss of synchronization [12].

BGP uses four different types of messages. These types are:

- OPEN
- UPDATE
- NOTIFICATION
- KEEPALIVE

OPEN message is used to initiate the communication. Open message can contain additional parameters. In the specification (see [12] for more information) there is also an additional parameter for authentication. This authentication information contains a one octet field *auth.code* followed by variable length authentication data. However no use for the parameter has been specified.

UPDATE messages are used to inform peers about changes in routes. Update messages can be used to inform about new routes, withdraw or inform changes in existing routes.

KEEPALIVE messages are transferred between peers if there is no need to transmit route information. The keepalive messages informs the peer that the connection is still alive and that the routes advertised by the peer are still in available.

To inform peers about errors a NOTIFICATION message is used. After the transmission of an error message the connection is closed and all routes learned through the BGP communications are cleared from the memory.

3.1 BGP in Operation

In the beginning of a BGP session the BGP speaker attempts to initiate a TCP connection. During the connection setup there is no method for peer-entity authentication. Even the use of MD5 Signature Option (see 3.2) does not help the BGP here. The BGP signature only tells the BGP speaker that the party who he is talking with knows the shared key. Therefore the peers are just forced to believe that the TCP/IP layer is working correctly and the connection has been made to the right peer.

After the TCP connection is established the BGP speaker sends an OPEN message where the marker-field is set to all ones. After the OPEN message the marker is set to a code agreed by both parties. The peer replies to the message by sending a BGP OPEN message back to the initiator of the communication. This message is acknowledged with a KEEPALIVE message where the marker is set to the previously agreed code. The current edition of BGP allows for addition of parameters for authentication of peers. However, as discussed above, no method for authentication has been specified.

Once the connection has been established, the peers exchange their information about the networks they can access by sending UPDATE messages to each other. By default the UPDATE messages are not authenticated or validated, however the messages are checked for consistency. If the message does not conform with specifications it is discarded and a NOTIFICATION message is sent to the peer. This results in the BGP speaker clearing all resources used for the connection and deleting records of all connections learned from the peer that send a faulty message.

3.2 Use of MD5 Signature Option (SBGP)

RFC2385 [7] describes a method to protect the BGP sessions. This method is the use of the TCP headers MD5 signature option. In some literature the use of MD5 Signature Option is also referred to as SBGP (Secure Border Gateway Protocol). All current BGP implementations must be capable of using this option but the use of it is optional. The method calls for addition of 16-byte MD5 digest to the TCP header. This digest is calculated from the TCP pseudo-header (source IP-address, destination IP address, zero-padded protocol number, and segment length), TCP header excluding options, TCP segment data and an

independently-specified key or password. RFC2385 [7] does not describe any method for the negotiation of the use of this option, neither for the key exchange. Instead the BGP speakers must be configured to use this method and the key (password) included in the digest must be exchanged separately.

Eventhough this method provides considerable increase to the security of the BGP there are few issues that make this less than optimal. According to Dobbertin [4] the MD5 hashing algorithm is vulnerable to a collision search attack and a collision could be found with 1996 pentium-PC in about 10 hours. However, this only means that another message containing the same hash can be created in this time. Exploiting this vulnerability in operation circumstances is still somewhat difficult, as a colliding message must be created quickly enough. Basically, before the attacked peer realizes that a message is missing.

Another weakness from the use of MD5 raising from its status as optional. As long as the attacker can find one BGP speaker that does not use MD5 on all of its communications, it can use that to launch his attack. This speaker will then under right circumstances feeds that information to the MD5 secured BGP hosts.

As another drawback of the use of MD5 signature option we can say that it consumes a considerable portion of the of the available size of the TCP header. The total size of the TCP header is 60 bytes, out of which 20 bytes is used by the mandatory fields. The MD5 signature option requires 18 bytes, which is almost half of the available space. It is still currently enough for the other options currently needed in BGP implementations but there is not much left.

Anyhow, the main disadvantage of the MD5 signature option is its lack of key exchange. The key must be exchanged manually between the administrators, which easily leads to the key being changed seldomly and increases the possibility of a weak key being used for the protection of the BGP session. This raises the possibilities of using dictionary-attacks or similar to find the key.

3.3 Security of BGP

3.3.1 Vulnerabilities of BGP

BGP has been designed for environments that are not perilous. Therefore BGP has three fundamental vulnerabilities that compromise its security [11]:

- BGP has no internal mechanism that provides strong protection of the integrity, freshness and peer entity authenticity of the messages in peer-to-peer BGP communications.
- No mechanism has been specified within BGP to validate the authority of an AS to announce network level reachability information
- No mechanism has been specified within BGP to ensure the authenticity of the path attributes announced by an AS.

The BGP messages are very vulnerable to attacks. Errors in message header, receipt of OPEN message in wrong

state, errors in OPEN message, receipt of a NOTIFICATION message and receipt of faulty UPDATE messages (length of unfeasible routes field or total path attribute length) will result in the BGP peer closing the connection, releasing all associated BGP resources, deleting all routes learned through that connection and a rerun of its decision process to decide new routes [11]. Modification of the withdraw routes field in UPDATE message could lead to elimination of existing routes. However this results only in the peer removing routes associated with the peer that communicated the change. An attacker could also modify the path attributes of routes to affect routing. For example he could claim his route to be better in quality than alternative routes and thus manage to divert traffic through his network.

As BGP runs on top of TCP, it is also vulnerable to TCP's vulnerabilities like syn flooding or denial of service attacks.

3.3.2 Possible Outcomes of Attacks on BGP

Attacks on BGP can lead to following situations [11]. The damage can be intentionally (malicious user/system) or unintentionally (misconfigured BGP speaker) generated. In Parenthesis I have added the aspect of security (see section 2.1) that it violates from the routing service.

- *Starvation*: The traffic cannot be routed into its destination because it is forwarded to a part of the network that cannot deliver to the final destination. (Availability)
- *Network Congestion*: More traffic than needed is forwarded through some portion of the network (Availability)
- *Blackhole*: Too much traffic is forwarded through a router than, due to heavy load, drops many/most/all of them (Availability)
- *Delay*: Traffic is forwarded through a suboptimal route causing unnecessary delay (Availability)
- *Looping*: Traffic is sent on a path that loops, leading to it never being delivered (Availability)
- *Eavesdrop*: Data traffic is forwarded onto a route where it normally would not take, giving access to it being eavesdropped (Confidentiality)
- *Partition*: A partitioning attack makes some portion of the network believe that it is partitioned from the rest of the network. (Availability)
- *Cut*: A portion of the network believes that it has no route to some network even though they are connected (Availability)
- *Churn*: The routing information in the network changes rapidly. This can lead to congestions, changes in transmission delays of the packets, some packets being delivered to networks where they can be eavesdropped or never being delivered. (Availability / Confidentiality)
- *Instability*: BGP never converges on a global forwarding state. (Availability)

- *Overload*: BGP itself generates a significant portion of the traffic the network can handle. (Availability)
- *Resource Exhaustion*: The BGP messages themselves exhaust critical resources such as routing table space. (Availability)

3.3.3 Passive Attacks on BGP

BGP is vulnerable to passive attacks. The routing data transmitted by BGP is carried in cleartext. This makes it possible for an attacker to eavesdrop the data. However, as discussed in section 2.3 the confidentiality of the routing data normally is no concern. Anyhow as the messages are transferred in cleartext the information contained in the messages cannot be used to authenticate peers, unless that information is encrypted. Use of MD5 signatures [7] does not provide any protection against eavesdropping. But, again, the confidentiality is not the main issue.

3.3.4 Active Attacks of BGP

Using the definition in section 2.1.3 there are five types of active attacks. In this chapter I will study on how the BGP-4 is protected against the attacks with the use of MD5 signature option (SBGP) [7] and without it.

BGP provides no protection against *replay attacks*. Even MD5 signatures do not protect against a replay attack, however, the TCP's sequence numbers do give some protection.

BGP itself provides no protection against *Message insertion* or *message deletion*. However, TCP's sequence numbers do provide some protection against both attacks. Message insertion would require accurate prediction of the sequence numbers and the receiver would notice the deleted message, as that messages TCP sequence numbers would be missing.

On behalf of *message modification* BGP again fails. There is no protection against this attack. As long as the length of the TCP payload is unchanged the attack would most likely go undetected. Change in the size of the payload would require the attacker to also modify the TCP header and, if the length would grow too large to fit into original TCP packets, creation of new TCP packets. Again use of MD5 signatures provides protection against this, in the extent of the MD5's capabilities. (see section ?? and [4] for discussion regarding the security of MD5).

Man-in-the-Middle attacks can also be performed easily against BGP hosts. As BGP has no method for authenticating peers, implementing a man-in-the-middle attack is very easy. MD5 signatures can also be used to protect against this.

3.3.5 Increasing the Security of the BGP

Literature shows a number of ways to secure BGP. The current secure version uses MD5 Signature Option (see 3.2). This is discussed more in RFC2385 [7] and is currently mandatory in all implementations of BGP. Implementation of MD5 signature option ensures integrity of the message transmitted and peer entity authentication as long as we can assume that the key has not been compromised. This requires the MD5 algorithm to be secure (see 3.2 for discussion of the

security of MD5) and that the key used to secure the communication is well protected and difficult to guess.

However the MD5 signature option is not enough for all situations. Garcia-Luna-Acaves and Smith [5] have suggested an improvement in the security which entails 5 changes. These changes are

- Encryption of all BGP messages between peers using keys exchanged at BGP link establishment time.
- Addition of message sequence number
- Addition of UPDATE sequence number of timestamp
- Addition of PREDECESSOR path attribute indicating the AS prior to the destination AS for the current route
- Digitally signing all unchanging UPDATE fields at the point of origin.

If these changes were to be implemented the security of BGP would increase considerably. Use of key exchange during the link establishment time would guarantee that good keys are used and that they are changed often enough. Also encryption of all messages would provide confidentiality (eventhough it is not always needed). Message and UPDATE sequence numbers would protect against replay attacks. The PREDECESSOR path attribute would allow verification of the path information and digital signatures of unchanging UPDATE fields would not only provide authentication and integrity between BGP peers but also of the full AS_PATH. However, the uptake of these changes would require all of the BGP speakers to be updated. Therefore, at least possibility to negotiate the options used should be possible between BGP peers to get over the transition time.

The fact that BGP is run on top of ordinary TCP/IP allows the use of any security methods available on TCP/IP. Mainly this would allow the use of IPsec, which could be used to authenticate and secure BGP sessions. IPsec could be run in ESP (Encapsulated Secure Payload) mode, which provides both authentication and integrity and could also be used to encrypt the entire payload. However, again it would take time to implement IPsec in all BGP speaking hosts. [1] discusses briefly the questions that should be answered if it would be decided that IPsec should be used to protect BGP sessions.

Actually Kent *et al.* [8] has described a method they call S-BGP (Secure BGP), which uses IPsec to provide authentication. The method also includes two other means of securing BGP. These are a new path attribute, *attestations*, which establishes that the subject of the attestation is authorized by the issuer to advertize a path to the specified blocks of address space and the use of PKI (Public Key Infrastructure) certificates. They describe in detail how a PKI certification tree could be build and how the attestation attribute and PKI certificates could be used to validate routes. IPsec would be used to prevent an active wiretapper from spoofing route withdrawals or replaying intercepted UPDATE messages. In a later paper Kent has revised the S-BGP architecture. [9] This revision mainly adds fourth element, route attestations (RA), to the architecture. A route attestation would be created by a S-BGP speaker and describes to which ASs UPDATES with the RA can be sent.

Introduction of these changes would address a number of BGP's vulnerabilities. However, a malfunctioning BGP speaker could still disturb internet traffic and the lack of UPDATE sequence numbers could still cause BGP peers to re-assert routes that have been withdrawn earlier. Again also the implementation of this method will require changes. First of all a method to distribute PKI certificates to BGP peers must be established as well as the entire certification tree should be implemented. In addition to this IPsec should be implemented on all peers and the attestations path attributes should be added. Therefore this cannot be implemented overnight but work must be done to make this possible. This technology is developed and is available as opensource. Anyhow a large scale deployment will require coordination of all parties to implement required changes and to build the PKI certificate tree.

3.4 Link Cuts

Above I have focussed on affecting routing by attacking the BGP protocol. However, if an attacker wishes to obtain sensitive information, he can also use link cuts to force the traffic through portions of network where he can intercept it. Or by cutting all of the links between to hosts he can also create a denial of service attack. Bellovin and Ganster [2] show that an attacker can make links appear to be dead by flooding them with denial-of-service attacks. This will lead the traffic to be routed around this link. By cutting the right links the traffic can be forced to travel through paths where the attacker can intercept them.

This type of attack requires knowledge of the network topology and an algorithm to calculate the necessary links to cut. Bellovin and Ganster [2] show that the calculations needed to launch this kind of an attack are quite efficient. In their example the calculations needed to alter the traffic in network of several hundred nodes took less than half a second. In addition to the algorithm enough resources to flood the connections is needed. Anyhow attacks of this type are difficult to counteract as link cuts can also happen by themselves and the network should be able reroute the traffic around the cut.

4 Conclusions

This article has focussed on the security of the inter-autonomous systems routing. The discussion has showed that BGP, the protocol used to share network reachability information between autonomous systems, is vulnerable to all types of attacks which can lead to many kind of disruptions on routing. By attacking BGP an attacker can attack any three aspects of security on the information transmitted by the network. The attacker can divert information to travel through a network where he can intercept it, in order to obtain the information (confidentiality) or to modify it (integrity). He can also disrupt routing considerably by creating routing loops, entering bogus routes, diverting traffic to suboptimal routes or causing congestions by diverting considerable amounts of traffic on links that cannot tolerate it (Availability).

This paper also found that the use of MD5 signature option considerably improves the security of the BGP (SBPG). If MD5 security option is configured well and good concern with the keys is kept (they are changed often enough and stored in secure place), reasonably good security can be obtained. However it must also be noted that MD5 algorithm is not unbreakable and the lack of method for key exchange makes it difficult to maintain the keys.

Suggestions to improve BGP were studied. Kent *et al.* [8] and [9] introduce a 4 part method to make BGP considerably more secure. This method includes the use of IPsec, creation of PKI certification tree and two new path attributes, route & address attestations. The implementation of these changes will take time but would increase the security of BGP considerably.

As last an alternative way to attack routing, using link cuts to force traffic to routes where the information can be intercepted, was studied. This attack requires resources to calculate the links to be cutted and to flood the links. However attacks of this type are difficult to counteract as the Internet by default is defined to react to problems in the network.

I can conclude that when inter-autonomous systems routing is somewhat protected when MD5 hashing is used but there still exists vulnerabilities that should be addressed. Possible security enhancements would be welcome to BGP. The suggestions introduced by Kent *et al.* are a good starting point for improvement. The disadvantage in the use of them, is that it need cooperation of numerous parties. However, these steps to create BGP more secure should be taken.

References

- [1] S. M. Bellovin. Guidelines for Mandating the Use of Ipsec IETF Internet Draft, October 2003
- [2] S. M. Bellovin and E. R. Ganster. Using Link Cuts to Attack Internet Routing Draft 2003 <http://www.research.att.com/~smb/papers/reroute.pdf>
- [3] Douglas E. Comer. Internetnetworking with TCP/IP, Principles, Protocols and Architectures. Prentice Hall, 2000
- [4] H. Dobbertin. The Status of MD5 After Recent Attack RSA Labs' CryptoBytes, Vol. 2 No. 2, Summer 1996. <http://www.nullify.org/docs/crypto2n2.pdf>
- [5] J. Garcia-Luna-Acaves and B. Smith. Securing the Border Gateway Protocol <http://www.cs.ucsb.edu/~rsg/Routing/references/smith96securing.pdf>
- [6] Dieter Gollman. Computer Security. John Wiley & Sons, 1999
- [7] A. Heffernan. Protection of BGP Sessions via the TCP MD5 Signature Option. RFC 2385, IETF Network Working Group, August 1998.
- [8] S. Kent, C. Lynn and K. Seo. Secure Border Gateway Protocol (S-BGP) IEEE Journal on Selected Areas In Communications, Vol. 18, No. 4, April 2000. <http://www.comsoc.org/sac/private/2000/apr/pdf/18sac04-kent.pdf>
- [9] S. Kent. Securing the Border Gateway Protocol: A Status Update http://www.ir.bbn.com/sbgp/S-BGP_CMS-2003-Kent.pdf
- [10] B. Korver and E. Rescorla. Guidelines for Writing RFC Text on Security Considerations RFC 3552, IETF Network Working Group, July 2003.
- [11] S. Murphy. BGP Security Vulnerabilities Analysis IETF Internet draft, IETF Network Working Group, June 2003
- [12] Y. Rekhter. A Border Gateway Protocol 4 (BGP-4). RFC 1771, IETF Network Working Group, March 1995.