Security in Internet Routing Protocols

Niklas Alberth
nikal407@student.liu.se

Rickard von Essen
ricvo248@student.liu.se

Linköping Institute of Technology
IDA Department of Computer and Information Science
Sweden

May 8, 2006

Abstract

Routers are used to route traffic between different networks. They are a vital part of all large networks, especially the Internet. Since it is hard to run a large network with only static routing, there is a need to do this dynamically and that is where routing protocols come in.

In this paper we investigate two of the most common routing protocols, BGP and OSPF, from a security perspective. We investigate what weaknesses that exist and how to protect against these.

1. Introduction

In order to connect several networks together to form an internet a router is needed. The router forwards packets to the appropriate network. To decide onto which network a packet will be forwarded to, the router looks in its routing table. The routing table is a list of which networks are connected to which ports on the router. To connect small intranets, static routes (manually configured) can be set up by the administrator, for larger internets this is not feasible. This is mainly because of the dynamic changes, adding new networks or when links go up and down, of a large internet. The solution is to use a dynamic routing protocol to automatically calculate and set up routing tables for all routers on the internet.

Correct routing is vital for a secure network since routing controls how the data in the network flows. If a malicious user can change the routing tables, he can forward data to his own node where he can read and manipulate the data, or he can detach hosts and networks creating a denial of service situation.

In this paper we have investigated some aspects of two internet routing protocols from a security perspective. We have collected known weaknesses of the two routing protocols BGP-4 and OSPF-2. We also point out how these weaknesses can be avoided or restricted by proper configuration. In the case of BGP we also make suggestions on how to make the next version of the protocol more secure.

This paper only investigates the routing protocol, this is the mechanism for distributing network topology data and choosing the best route. It does not include the actual forwarding of packages.

To show how vital proper functional routing is we quote Vetter et. al.: *"It has been pointed out that blackhole routers will be a performance killer for distance-vector routing protocols like RIPv2. On April 27, 1997, a router from MAI Network services in Virginia absorbed about 50,000 network addresses which caused much of the internet to be disconnected from 20 minutes to 3 hours. A technical bug was blamed to the MAI's Bay Network router, but the same attack is very feasible from an evil insider."* [19]. This clearly shows the magnitude of damage that attacks against routing infrastructure can do.

1.1. Classification of Threats and Damages

To better understand the threat against routers we divide the treats and damages into different classes. Threats to routing security can be divided into three classes, according to *An experimental study of insider attacks for OSPF routing protocol* by Vetter, Wang and Wu [19]:

- **External:** an intruding device, joining the router domain (collection of routers exchanging information). This could be a computer running routing software.
- **Internal:** a compromised router. This class includes compromised routers either hacked or with stolen access. It could also be an insider who has legal access but with malicious intent or just a misconfiguration.

These threats are much harder to protect against, since the source is a legal participant in the router domain.

- **Not participating:** a host which does not join the router domain and therefore is invisible to all routers. If an attacker can insert false data into the routers without joining the routing domain detection is much harder than detecting a illegal participant.

Below is a list of damages that a network can suffer when being attacked [12].

- **Starvation:** data is sent to a part of the network that can not deliver it to the correct final destination.
- **Network Congestion:** a link on the network is overloaded with traffic and is forced to drop packets.
- **Blackhole:** large amount of data is directed at a router that can not handle it. The traffic will be dropped.
- **Delay:** data may go over a slower path.
- **Looping:** data may end up in a routing loop and will never arrive at its final destination.
- **Eavesdrop:** a malicious user can route traffic through a part of a network which he is in control of.
- **Partition:** part of a network thinks it is partitioned off from the network when it in reality is not.
- **Cut:** part of the network is not aware of routes to another network even though it actually has a physical connection.
- **Churn:** network topology changes rapidly resulting in a large variation in how the data is delivered.
- **Instability:** the best route keeps changing back and forth and convergence on a global forwarding state is not achieved. A global forwarding state is when all routers agree on the best path to select.
- **Overload:** the routing protocol traffic takes up a significant amount of all network traffic.
- **Resource exhaustion:** some routing resource is exhausted and the router can not perform its duties.
- **Address spoofing:** data traffic is forwarded through some router or network that is spoofing the legitimate address.

2. BGP

The Border Gateway Protocol (BGP) [17] is an *Exterior Gateway Protocol – EGP*. As such BGP is mostly used to route traffic between different networks. To connect these different networks the internet is broken into several ASes. *"An* **Autonomous System (AS)** *is a group of networks and routers under the authority of a single administration."* [3]. BGP can also be used for routing within a large AS, using private AS numbers. The traffic can flow directly between two physically connected ASes or transit through one or more ASes before reaching its destination. This is called transit traffic.

BGP has its root in the first EGP protocol [10] which was used to connect ASes in the early DARPA Internet. It evolved with experience gained from the use of EGP in the NFSNET Backbone [1, 16]. Version 4 (BGP-4) is the most current BGP revision used and is the de facto standard for connecting autonomous systems together on the Internet.

BGP is a full mesh protocol, all participating routers must be configured to be peers. The mesh configuration is done manually and not through any automatic discovery of nearby routers. When two routers need to exchange routing information this is done over a TCP connection on port 179, which is the default port for BGP communication. By using TCP, BGP does not need to handle error control, retransmission or reliability of the connection, these services are provided by TCP. After the connection is initialized it is maintained with the BGP KEEPALIVE messages every 60 seconds. Changes to the routing topology are done with UPDATE messages. BGP only sends notifications when changes occur. BGP does not send complete routing tables periodically.

A BGP UPDATE message consists of several *type-length-vectors – TLVs*. First comes the withdrawn routes, then a set of attributes and last a list of network prefixes. The list of network prefixes are called *Network Layer Reachability Information – NLRI*. In an UPDATE message the set of attributes apply to all network prefixes in that message.

These attributes are mandatory:

- **ORIGIN:** how the route was learned. It can be
  - **IGP:** learned from an internal gateway protocol
  - **EGP:** learned from an external gateway protocol
  - **INCOMPLETE**: it is not known how the route was learned
- **AS_PATH:** is a ordered list of ASes that the traffic will need to pass through to reach the network prefixes announced.
- **NEXT_HOP:** contains the ip address used as the next hop for the network prefixes in the BGP message

- **LOCAL_PREF:** local preference, a value that can be used to select one route over another. This is used within an AS.

BGP works by informing neighboring BGP capable routers which network prefixes it can access. This information is passed via the NLRI field. This can be network prefixes the router is directly connected to or network prefixes that can be reach via another router. In the case of advertising network prefixes that the router is not connected directly to, it supplies the list of ASes the traffic have to pass through in the AS_LIST.

When presented with several ways to reach a network prefix BGP uses a selection algorithm to only save what it thinks is the best route. How this is done exactly is not specified in the current BGP RFC [17]. Often this is based on the shortest AS_PATH to the end destination. But it can also depend on other factors like preference to send transit traffic over a cheaper transit provider.

2.1. Potential Vulnerabilities

As with many Internet protocols that have been around for a long time, there were initially no security features built into the BGP protocol. The reason for this was that when the protocol was designed, the networks where small and the administrators of the different networks trusted each other. Today the picture is different. The Internet has grown significantly and it is not possible to trust all the participating networks and users.

BGP provides the service to connect isolated island of networks together to form one large network, like the Internet we are using to today. An attack or even a misconfiguration can cut parts of the network off from the rest of the global network resulting in degraded communication. These types of cuts can be the result of a faulty route being announced resulting in routers starting to send traffic the wrong way. The traffic might end up in a blackhole or in a loop. The traffic can also pass through a network controlled by an attacker, where it is possible to eavesdrop on the traffic or kidnap the traffic and run a fake server. Imagine redirecting the traffic for online banking.

Whilst not being a vulnerability in BGP it is possible to cause instability in the routing table, by announcing and withdrawing routes quickly (route flapping). As with all network connected devices denial of service is also a problem.

In order to make BGP more stable and secure there needs to be secure communication between BGP routers. Also there needs to be a way to make sure that the announced network prefixes are actually announced by its owner.

2.2. TCP MD5 Signature

In an effort to protect the communication between two BGP routers the latest BGP standard enforces the use of a TCP MD5 [18] signature when exchanging BGP messages.

As BGP makes use of unencrypted TCP when it talks to its neighbors it is possible to inject fake messages. To perform an attack on TCP the attacker need to guess the TCP sequence number. With a modern TCP/IP stack this is hard but it is theoretically possible.

The MD5 hash is constructed from: the TCP header, a TCP pseudo-header, the data segment and a pre-shared secret word. The TCP pseudo-header consist of fields from the ip header. By using a pre-shared secret word an attacker attempting a *man in the middle attack* would need to not only guess the TCP sequence number but also to know the secret word.

The drawback of using a TCP MD5 signature is that both peers needs to be manually configured to enable the feature and to setup the pre-shared secret. Also, in order to be cryptological strong the secret word have to be renewed at least once every 90 days [7]. The secret word should be unique for a pair of peers. Adding the use of a MD5 hash puts additional load on the processor of the router.

2.3. Detecting Invalid Route Announcements

In April 2001 a small ISP announced, by mistake, that it accepted traffic for 9177 network prefixes [24]. This announcement got picked up by the ISPs BGP peers and spread into the Internet causing one of the first major routing incidents.

Since a small mistake can be noticed all over the Internet it is necessary to employ methods to ensure that only valid routes becomes announced and propagated. Another reason that there needs to be a way to control authentic routes, is that it is hard to discover who initially made the faulty announcement, this can cripple Internet traffic for hours.

The most basic method is to use static filter rules on the incoming BGP announcements. Managing this filter list is only feasible if the network topology is stable and does not change often.

Usually there is only one AS that announce that it is the origin for a network prefix. But there are cases when two or more ASes announce the same network prefix. This is commonly called *Multiple Origin AS – MOAS*. This can either be intentional for doing multi homing without a AS, for a range of ip addresses that are reachable by more than one route. It can also be due to misconfiguration or an attack.

One non cryptographic suggestion to solve the MOAS problem is an addition to the BGP protocol that carries a list of ASes that are allowed to originate the network prefix in question [24].

If a router sees two ASes announce that they handle a network prefix and both are in the MOAS list everything is looking correct. If the router for example see three ASes announcing a network prefix but the MOAS list is different it detects there is something wrong. It cannot detect what is wrong, only that there is a discrepancy. MOAS does not add any more protection than this but it does add a method to detect routing anomalies.

2.4. Route flapping

Since BGP uses TCP to communicate with its peers it is dependent on the TCP implementation not to be vulnerable. The TCP connection can be shutdown by a TCP RST (reset) attack. A TCP reset attack is when a third malicious host tries to close a TCP connection between two hosts. By sending a TCP packet with the RST bit set. Interruptions in the TCP connection causes the router to release all routes associated with the peer whose connection is lost. The withdrawn routes then propagate further on to the Internet. When the connection later is resumed new routes will be inserted into the routing table. This will cause route flapping when there is an instable connection between two peers. There exists methods to try and suppress route flapping by waiting an exponentially longer time to announce routes learned on a unstable connection.

Another way to cause route flapping is to send a broken BGP message. The behaviour specified by the RFC [17] is to drop the connection and clear the routing entries associated with that connection when a broken message is received.

2.5. S-BGP

S-BGP [6, 8] aims to be a complete solution to secure BGP. In *Secure Border Gateway Protocol (S-BGP)* [6] Kent et. al. claim that other available security solutions only addresses parts of the problem where S-BGP aims to be a complete solution.

S-BGP uses two *Public Key Infrastructures – PKIs*, a new path attribute and IPSEC [5] in order to provide a complete security solution. The first PKI is used to sign the ip ranges assigned to a organization. ICANN is at the top level of this verification hierarchy with a self signed certificate claiming ownership of the entire ip address space. This is divided up on a local registry like ARIN (American Registry for Internet Numbers) and RIPE (Réseaux IP Européens); which handle the assignment of network prefixes and ASes for a specific geographic area. Which in turns signs the network prefix it hands out until a certificate for a network prefix reaches an end user. This procedure make sure that it is possible to verify the owner of a network prefix.

The second PKI is used for assignment of ASes and router associations. This certificate tree also have ICANN at the root and then the local registries. The third level is organizations that owns ASes, the next is a tier of AS numbers and routers.

These PKIs are used to verify that the AS that are announcing a network range are the actual owner and that the router announcing it belongs to the AS in question.

IPSEC is used to secure the communication between two router peers. Similar to the TCP MD5 checksum method. The advantage of using IPSEC is that it can negotiate the cryptographic keys automatically via IKE (Internet Key Exchange) without any manual configuration of a shared secret word. It also adds protection to TCP against RST attacks. IPSEC counters this by encrypting the TCP layer.

2.6. Additional Security Suggestions

There are a few more ideas on how to secure BGP [4]. One of these ideas is to add *Denial of Service – DoS* protection it is called the *BGP TTL Security Hack – BTSH*. It works by setting the TTL field in the TCP packet to 255. BGP routers then discards packets arriving at port 179 with a TTL lower than 254. When a packet is forwarded through a router the TTL field is decreased by one. BTSH make sure that the packet comes from a host directly connected to the receiver and not through any other routers. An attacker would have to be directly connected to the router in order to be able to perform a DoS attack with a high value in the TTL field. This adds protection for routers that are not located on the edge of the network as they are at least one hop away from the attacker. Routers often have special hardware to do the packet filtering so the offending packets never reach the routers CPU, when BTSH is deployed.

Cisco have implemented a extension to BGP they call *secure origin BGP – soBGP* [15, 23]. With the goal to provide a method to authenticate an AS that is originating a network range. It is designed to add a low overhead to existing BGP and not to rely on a central authority handling certificates. soBGP is using X.509 certificates to authenticate an AS to be authorized to announce a network prefix. It also checks to see that a given AS_PATH is valid by building a topology map of the network and asking each router pair on the way if they can reach each other. It does not suggest a method to secure communication between routers. IPSEC or TCP MD5 can be used. soBGP transports the certificates over a new BGP message type. This is to avoid using central repository of certificates.

There is also an idea that suggests that DNS shall be used to house the mapping between AS and network prefix. This suffers from the ”chicken and the egg” problem. If the routing is not working properly it is not possible to reach the DNS servers.

3. OSPF

The *Open Shortest Path First – OSPF* [11] is an *Interior Gateway Protocol – IGP*, it is used to route traffic within an AS. OSPF is a link-state routing protocol that sets up routing paths based on the state of all links within the AS. In OSPF all links are assigned a cost and all OSPF routers share this information and calculates the shortest path to every other router within the AS.

The current standard for OSPF is version 2, OSPF-2 [11] for IPv4 and OSPF-3 [2] for IPv6. This paper will only cover OSPF-2.

Here we will describe some important features and mechanisms in OSPF. The OSPF protocol is large and complicated, as such only the most important parts or parts of particular interest for this analysis are included. For a complete definition of OSPF please refer to RFC 2328 [11].

All OSPF routers monitor the status of their interfaces and periodically transmit the status and a metric cost for the interface to their neighbors. This information is also transmitted upon status change. The neighbor router receives the information and floods it to its neighbors.

The cost of the link is only for outgoing data never for incoming data. When an OSPF router starts up it has to find its neighbors and form adjacency. This is done with a simple hello handshake where the router sends a multicast message that all other routers in the network reply to. After this is done the router transmit a *Link State Advertisements - LSA* with information for each of its interfaces. These are received and if they are found to be correct, are acknowledged by adjacent routers. These routers flood the data onto all its interfaces except the one the LSA arrived on.

After a short period of intense data interchange, all routers have the same database of valid LSAs. From this database, a shortest spanning tree is calculated. Using Dijkstra's shortest path algorithm [22], each router calculates the spanning tree with itself as the root. This information gives the shortest path to all networks and is used to build a routing table.

The LSA is the heart of the distribution of network topology data in OSPF and most of the features and vulnerabilities can be found in it, therefore we will go into some detail describing it.

The LSA carries information between routers. The LSA is built by the advertising router and sent to its neighbors. The neighbors check if the LSA is valid and if it is valid insert it into the routing database and flooded onto all interfaces except the one that received the LSA. Multiple LSAs are packed into a *Link State Updates – LSU* before sending.

To make the protocol more scalable, each autonomous system can be divided into several areas. There is one special area: area 0 – also called the backbone. All routers belong to just one area except *Area Border Routers – ABR*

LS Age		Options	LS Type
Link State ID			
Advertising Router			
LS Sequence Number			
LS Checksum		Length	
|V|E|B|		#Links	
Link ID			
Link Data			
Type	#TOS	TOS 0 Metric	

Figure 1. Link State Advertisement packet [9]

which belong to an area and the backbone. Routers can also be connected to external networks in other AS, these are called *Autonomous System Boundary Routers – ASBR*. This gives OSPF a two-level hierarchy. See figure 2.

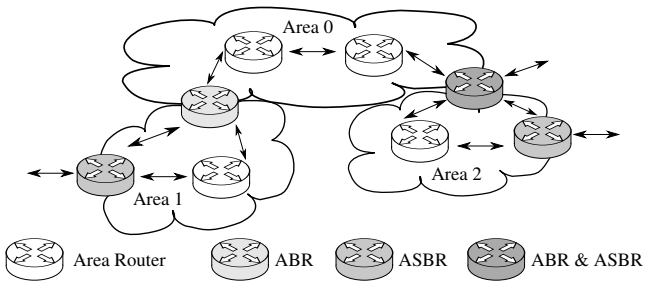

Figure 2. Three areas with different types of routers

ABRs run two versions of OSPF, one for each area it belongs to. After calculation of the shortest spanning tree an ABR will aggregate information and advertise it into the other area. ASBR use information of routes received from other sources, e.g. a EGP, and advertise them into all areas.

There can be several ASBR in the same area, if they connect to the same network they will be considered a redundant link. ABRs are always connected to exactly one area and the backbone.

There are five types of link state advertisements [13]:

- **Type 1:** *Each router advertises a Router Links Advertisement to its area, describing the state of each of the router's interfaces in the area.*
- **Type 2:** *Each multi-access network selects a Designated Router through which to communicate, in order to reduce traffic on the network. The Designated*

Router advertises the list of routers connected to the network in a Network Links Advertisement.

- **Type 3:** *Each ABR advertises a Summary Link Advertisement to each of its attached areas, describing routes to networks outside that area (but within the autonomous system).*

- **Type 4:** *An ABR advertises a Summary Link Advertisement to each of its attached areas, describing routes to ASBR's outside that area.*

- **Type 5:** *Each ASBR advertises many AS External Link Advertisements, each describing a route to a destination in another autonomous system.*

When an LSA arrives at a router, some properties are checked before it is inserted into the router database and flooded on to other interfaces. If one step fails the LSA is discarded. See figure 3.

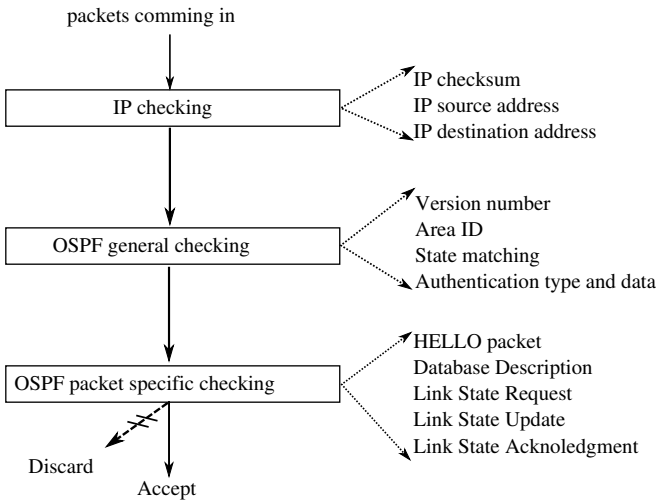

Figure 3. OSPF packet procedure checking [21]

Two fields are of special interest, the *Age* and the *Sequence* fields. The Age field is incremented with each hop and for every second stored in a routers database. If it reaches `MaxAge`, usually after 1 hour, it is deleted from the database and the spanning tree is recalculated without this link. After that the LSA is flooded to the neighboring routers, which will remove the LSA from all router's databases. The flushing of max aged LSAs will make the databases converge very quickly.

3.1. Cryptographic protection

OSPF use one of three authentication methods: none, simple or cryptographic. The standard authentication in most implementation is none. The second method, simple, uses a plain text password that has to match among all routers in the area. Since the password is sent in clear text with every LSA an attacker can sniff the traffic on a segment to read the password easily.

The cryptographic authentication method uses a public key signing of all LSAs. The specification for this is detailed in *RFC 2154: OSPF with Digital Signatures* [14].

We will give a brief description that will cover some details needed to analyze the advantages and limitations of signed LSAs.

A trusted entity is used to sign all routers public keys. All routers have the trusted entity's public key stored. Routers send out a special LSA containing its own public key. Since this package is signed by the trusted entity all routers can verify that it is correct and trusted. These public keys are inserted into a database used to check all signed LSAs. Currently signing is done with RSA or DSA and the hash algorithm used is either MD5 or SHA.

The signing of a LSA is always done by the advertising router, since it is the only entity with the private key, no other entity can alter the LSA. Since all routers must be able to age the LSA the LS Age field is *not* signed.

Any router could set the LS Age field to `MaxAge` and flood the LSA to its neighbors and thereby flush the LSA out of the system without being the advertising router. This is prevented by requiring a signed LSA with the LS Age field set to `MaxAge` to include the LS Age field in the sign as opposed to LSAs with a lower LS Age. The drawback of this is that every router have to age out the LSA by it self which makes the global forwarding state converge slower [13].

This authentication only protects against attackers connecting to the network with their own faked routers. This do not protect the confidentiality of the routing information nor does it prevent fake data from being inserted into the network by a compromised router, although it can only insert LSAs for which it is a legitimate advertiser of.

More about the protection and gain of cryptography is given with each vulnerability below.

3.2. OSPF Multicast

The default setting is to multicast (called broadcast in the OSPF RFC) routing information on multicast capable networks, e.g. ethernet. This sends all routing information to everyone that listens on the multicast addresses `AllSPFRouters` (224.0.0.5) and `AllDRouters` (224.0.0.6). This is unnecessarly open and can give a attacker valuable information about the setup.

The suggested solution to this is to shut off OSPF multicast communication and use unicast communication. The downside of this solution is that the administrator must define the addresses for all neighbors in each router. Normally multicasting is used to automatically detect them. This also gives some protection against some misconfigurations and slips like a user installing his own routing software on his desktop, breaking the routing. Disabling multicasting increases the security since it makes it harder to join the router domain for an external attacker.

3.3. MaxAge Attack

The MaxAge attack as described by Gong et. al. [20] abuses the LS Age field of an LSA. Since the age should be incremented by one with each hop and once per second in router databases, all routers can change this field. However, if a malicious router modifies the LS Age to `MaxAge` (usually one hour) and then re-injects it into the system, this will make all other routers flush the LSA out of their database. Eventually the original advertising router will receive the max aged LSA. It will recognize it as its own LSA and retransmit a new LSA with a fresh sequence number that will replace the LSA in all routers, this is called *fight-back*.

If the malicious router keeps sending max aged LSA for a link the network topology will become unstable since the link will go up and down in short intervals. This should be detected by network monitoring systems. Also if one or more malicious routers partition the area it can hinder the return of the bogus LSA to the original advertiser and thereby prevent it from *fight-back* [19].

The use of plain text key to authenticate the LSAs does not protect against this attack since an attacker only has to receive an LSA to get the key. It only helps to keep external attackers out if they cannot sniff the key. Using the cryptographic signing of packages requires a LSA to include the Age field if and only if it is equal to `MaxAge`. This protects against this type of attack since only the advertising router that holds the proper private key can flush the LSA from the system. All other routers can age the LSA and remove it from their own database if it passes `MaxAge`, but do not flush it as an unsigned version of OSPF would do. This makes the routing databases converge slower than normal OSPF.

We fail to see why a malicious router can not insert an LSA with LS Age = MaxAge $-x$ (where x is a small number)? This would make the LSA age out within x seconds. This is untested but will probably result in a slightly less effective attack. Maybe the attack will be useless since the advertising router will have time to *fight back* before the LSAs are discarded.

3.4. Seq++ Attack

The Seq++ attack abuses the LS Sequence Number. When the malicious router receives an LSA it increase the LS Sequence number, recalculates the checksum and reinject the LSA into the system. Since the new LSA has a greater sequence number it will replace all LSAs in the system with this Link State ID. As in the MaxAge attack this will cause the real advertising router to *fight-back*, resulting in a unstable network topology [20].

Signing LSAs will solve this vulnerability since the LS Sequence Number is included in the message digest and therefore the attacker can not change it without breaking the signing.

3.5. MaxSeq Attack

The MaxSeq attack is described by Gong et. al. [20] as: *"This attack differs from Seq++ attack in two aspects:*

- *Seq++ is a persistent attack where you have to modify sequence number constantly to keep the attack going; while MaxSeq attack is hit-and-run attack, that is you only need to modify once or twice, and this will then bring the system into a unstable state*
- *MaxSeq attack modifies the link state metric and set the LSA Sequence Number to `0x7FFFFFFF` (i.e., Max Sequence Number), not just doing simple increment.*"

This LSA will be considered the latest by all routers. When the advertising router receives the LSA it will recognize it as a faulty LSA and try to send out a new LSA.

The RFC specifies that a router that wraps the LS Sequence Number has to flush the max sequence numbered LSA before wrapping, otherwise all LSA with lower LS Sequence Number will be considered older and no router will accept them. During normal operation a router is not allowed to send out new link state updates with a new sequence number more frequently than one every five seconds [11] so it will take a minimum of 340 years before wrapping the sequence number.

Some implementations fail to comply with the mandatory flushing of the LSA before wrapping. This will make them extra vulnerable to this attack since without flushing the bogus LSA, it will persist in router databases until it ages out, which is usually one hour. This makes this attack very powerful, by inserting one bogus LSA a malicious router can break the routing for one hour.

3.6. Intrusion Detection

By monitoring the LSA traffic in the network it is easy to detect all of these attacks and also detect failing routers

and links. More on this topic can be found in *Intrusion detection for link state routing protocol through integrated network management* [20]. Since the routers are vital to a network infrastructure they should be properly monitored to detect attacks, misconfigurations and other failures.

4. Conclusions

In general the routing protocols we have examined are not designed with security in mind. All security measures were added after the fact. To achieve a more secure routing infrastructure routing protocols should be secured as much as possible and proper network monitoring should be used.

4.1. BGP

Examining BGP we have came to the conclusion that there are two main problems with the protocol that need to be addressed.

Security is needed when two routers talk to each other to avoid injection of false data. This can be done via the TCP MD5 method or by using IPSEC. The TCP MD5 method is easier to implement within a small network of BGP routers. The IPSEC solution requires a substantially larger infrastructure to manage the certificates necessary to exchange cryptographic keys. After setup, the IPSEC solution requires less manual configuration on each router as they can exchange keys automatically.

The second problem is that there is no way to know that a announced ip range under an AS is allowed to be announced by that AS. The S-BGP initiative solves that problem with an extensive public key infrastructure based on X.509 certificates. The PKI approach requires a multitude of descending certificates that needs to be validated everywhere. The soBGP approach to authenticating orignating ASes is similar to S-BGP. The main differense is how the certificates are managed. The use of cryptography requires more CPU power and significantly more memory in routers. Other approaches use the DNS system to verify the authenticity of an AS announcement.

4.2. OSPF

OSPF suffers from the unprotected age and sequence number fields in the link state announcements. The LS Age field can only be partially protected by signing, since it is modified in each hop. Without signing the LSA, OSPF is insecure as it only requires an attacker to be able to connect to the network where OSPF sends traffic and join the routing domain. If a plain text key is used it also requires the attacker to sniff it.

By turning of multicast in OSPF the administrator can make it harder to detect and join the routing domain. This comes at the rather small price of having to manually configure the address of all neighboring routers on every router.

OSPF should be used in unicast mode and with signing of LSAs to make it as secure as possible. Furthermore it should only be used on the routers interfaces where it is necessary. Good routines for monitoring routing should be used. An intrusion detection system that monitors the OSPF traffic and reports unusual behavior should be deployed.

References

[1] U. Braun. *RFC 1093: NSFNET routing architecture*. IETF - Network Working Group, February 1989.

[2] R. Coltun, D. Furguson, and J. Moy. *RFC 2740: OSPF for IPv6*. IETF - Network Working Group, December 1999.

[3] B. A. Forouzan. *TCP/IP Protocol Suite*. McGraw-Hill, third edition, 2006. ISBN: 0-07-111583-8.

[4] G. Goth. Fixing bgp might be difficult - or not so though. *IEEE Internet Computing*, pages 7 – 9, May June 2003. 1098-7801/03.

[5] S. Kent and R. Atkinson. *RFC 2401: Security Architecture for the Internet Protocol*. IETF - Network Working Group, November 1998.

[6] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (s-bgp). *IEEE Journal on selected areas in communications*, 18(4):582 – 592, April 2000.

[7] M. Leech. *RFC 3562: Key Management Considerations for the TCP MD5 Signature Option*. IETF - Network Working Group, July 2003.

[8] C. Lynn, S. Kent, and S. K. *RFC 3779: X.509 Extensions for IP Addresses and AS Identifiers*. IETF - Network Working Group, June 2004.

[9] J. McEachen and R. Chesser. Vulnerabilities in the open shortest path first interior gateway protoco. In *MILCOM 2000. 21st Century Military Communications Conference Proceedings*, volume 2, pages 1224 – 1228, October 2000. 10.1109/MILCOM.2000.904121.

[10] D. Mills. *RFC 904: Exterior Gateway Protocol formal specification*. IETF - Network Working Group, April 1984.

[11] J. Moy. *RFC 2328: OSPF Version 2*. IETF - Network Working Group, April 1998.

[12] S. Murphy. *RFC 4272: BGP Security Vulnerabilities Analysis*. IETF - Network Working Group, January 2006.

[13] S. Murphy and M. Badger. Digital signature protection of the ospf routing protocol. In *Network and Distributed System Security, 1996 Proceedings of the Symposium on*, pages 93 – 102, February 1996. 10.1109/NDSS.1996.492416.

[14] S. Murphy, B. M., and B. Wellington. *RFC 2154: OSPF with Digital Signatures*. IETF - Network Working Group, June 1997.

[15] J. Ng. Extensions to bgp to support secure origin bgp (sobgp). Technical report, April 2004.

[16] J. Rekhter. *RFC 1092: EGP and policy based routing in the new NSFNET backbone*. IETF - Network Working Group, February 1989.

[17] Y. Rekhter, T. Li, and S. Hares. *RFC 4271: A Border Gateway Protocol 4 (BGP-4)*. IETF - Network Working Group, January 2006.

[18] R. Rivest. *RFC 1321: The MD5 Message-Digest Algorithm*. IETF - Network Working Group, April 1992.

[19] B. Vetter, F. Wang, and S. Wu. An experimental study of insider attacks for ospf routing protocol. In *Network Protocols, 1997. Proceedings., 1997 International Conference on*, pages 293 – 300, October 1997. 10.1109/ICNP.1997.64373.

[20] F. Wang, F. Gong, F. Wu, and R. Narayan. Intrusion detection for link state routing protocol through integrated network management. In *Computer Communications and Networks, 1999. Proceedings. Eight International Conference on*, pages 634 – 639, October 1999. 10.1109/ICCCN.1999.805585.

[21] F. Wang and S. Wu. On the vulnerabilities and protection of ospf routing protocol. In *Computer Communications and Networks, 1998. Proceedings. 7th International Conference on*, pages 148 – 152, October 1998.

[22] E. W. e. a. Weisstein. Dijkstra's algorithm. From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/DijkstrasAlgorithm.html.

[23] R. White. Securing bgp through secure origin bgp. Technical report, September 2003. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_6-3/securing_bgp_sobgp.html.

[24] X. Zhao, P. Pei, L. Wang, D. Massey, D. Mankin, S. Wu, and L. Zhang. Detection of invalid routing announcements in the internet. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN'02)*, 2002. 0-4695-1597-5/02.