



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Security Issues with Internet Routing – Border Gateway Protocol (BGP)

Introduction:

Richard Clark, Special Advisor to the President of the United States for Cyber Security, stated in the Draft of the National Strategy to Secure Cyberspace, that there are three major security flaws with the Internet; Domain Name Service (DNS), Internet Protocol (IP) and Border Gateway Protocol (BGP). This research paper will explain the security issues and present current work/research that is addressing one of these security flaws, the Border Gateway Protocol (BGP).

To assist those of us that have limited exposure to the world of routing, I will start this paper with an explanation of routing basics. The specific type of routing we will be discussing is Internet Protocol (IP) routing. We will discuss what is meant by routing, some of the terminology used in routing and how routing fits into the overall scheme of the Internet Infrastructure. After we complete a discussion of basic routing we will move on to BGP basics. We will discuss what BGP is, specific terminology used with BGP and how BGP works. Once we have an understanding of routing and BGP, we will move on to the security issues with BGP. In our security issues discussion we will identify the vulnerabilities of BGP and explain some of the possible exploits. Finally, we will discuss some current projects trying to identify ways in which these security issues can be reduced. The conclusion of this paper will be a summarization of what we have discussed and will provide the reader with some URL's that will point to other resources discussing BGP.

Routing Basics:

Routing is a term we have all heard, but not all of us truly know what is meant by it. A definition of routing as provided by Cisco Systems Inc.; "Is the act of moving information across an internetwork from source to destination". The key term in this definition is "internetwork". This term implies that there is a least one routing device between the source and the destination.

To understand what a router is and how it fits into the world of networking, we must have a fundamental knowledge of the Open Systems Interconnect (OSI) Reference Model. The OSI Reference Model is a standard developed to provide guidance to the manufacturing community in developing products that will operate together. There are 7 separate layers to the OSI Reference Model (see figure 1). Each layer of the OSI model provides its own functionality and standards. The OSI Reference Model is like a ladder, each layer of the model must provide the necessary information to the next layer for a data packet to be transported to its end destination. When a source device wants to send some information to a destination device the information is first broken up into data

packets. The data packet is handed off to each layer of the OSI Reference Model from top to bottom. At the bottom layers, the data packet is encapsulated into a frame or cell and is sent to the destination device. The destination device takes the frame/cell and starting from the bottom of the OSI Reference Model starts to unencapsulate the data packet. Once the data packets are unencapsulated they are reassembled into the information message that the source device wanted to send.

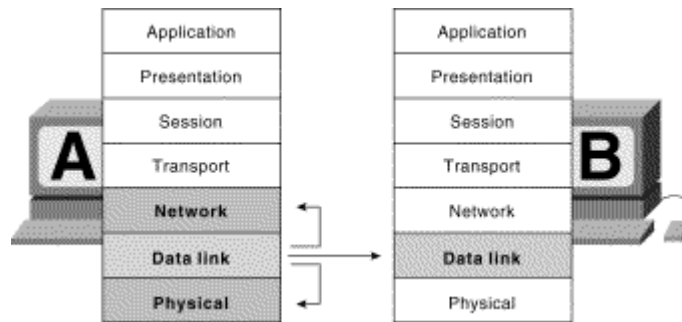


Figure 1
(Property of Cisco Systems, Inc.)

(http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/introint.htm#xtocid8)

The specific layer that we are interested in is Network Layer, layer 3. This is the layer where the functionality of routing is added to the data packets. The exact functions of this layer as identified by SANS Institute, 1.1 SANS Security Essentials I: Networking Concepts are: "The Network Layer handles interaction with the network address scheme and connectivity over multiple network segments. It describes how systems on different network segments find and communicate with each other". This simple description of the network layer of the OSI Reference Model provides us with enough information to move on with our subject of routing.

Routers are physical devices that operate at layer 3 of the OSI Reference Model. Based on the above description, of layer 3 of the OSI Reference Model, there are two required pieces of information for routing; an Internet Protocol (IP) address and identification of Routed/Routing Protocols.

You can not have a discussion of routing without a quick description of IP addressing. Let's think of an IP address as a street address. Every building on a street in a town has a specific street address, just like every interface on a router has an IP address. For someone to send a package to a specific building they must know what the street address is for that building. For device A on the network to send a data packet to device B on the network, device A must know the specific IP address of device B. An IP address is made up of two parts which are the network part and a node part. The network part identifies the domain where a device is located which is similar to the town, state, zip code for which our building is in. The node address is the specific address for the device itself

which is similar to the specific street address for our building. An IP address is constructed using binary numbers and is broken into four octets. An octet is a series of binary numbers that when added up are equal to a decimal number from 0 – 255. We have all heard of the different classes of IP addresses. These classes were once used to determine the network part of our specific IP address assigned to our router interface. Figure 2 is a table which lists the different IP address classes.

<u>IP Address Classes & Ranges</u>		
Class A	0.0.0.0 - 127.255.255.255	mask 255.0.0.0
Class B	128.0.0.0 - 191.255.255.255	mask 255.255.0.0
Class C	192.0.0.0 - 223.255.255.255	mask 255.255.255.0
Class D	224.0.0.0 - 239.255.255.255	reserved multicast
Class E	240.0.0.0 - 255.255.255.255	future use

Figure 2

The different classes made the routing of information easier, because once the routers knew the network part of the destination device it could route the data packet. The problem with this classful routing was that it wasted a lot of IP addresses. Very large organizations would be assigned a class A address range and not use all the addresses. While small organizations were not able to receive even a small block of IP addresses, class C, because they could not justify the need. Plus each IP address range assigned had to have its own entry in the routing tables of our routers. This made the routing table very large. To mitigate these problems a process known as classless interdomain routing (CIDR) was developed. CIDR introduced and uses the concept of subnet masking. When routers are configured to use CIDR, they use the subnet mask to determine the network part of the IP address. The subnet mask tells the router how many addresses are used, allowing large blocks of IP addresses to be broken down into smaller blocks of addresses. When using CIDR, the subnet mask is like a zip code from our building in a town example. The zip code tells the post office the geographic area of the building which the package is being sent. The subnet mask tells the router the general area of the IP address. Subnetting a block of IP addresses can be a very complex and confusing task. Some people even consider it an art form. The important factor for this paper with regards to IP addressing is that it is required for layer 3 of the OSI Reference Model and that each used interface on a router requires an IP address.

Let's move on to Routed/Routing Protocols. Protocols are the software components within the router. Protocols are used to determine the path the data packet is going to use. The easiest way to remember the difference between a routed protocol and a routing protocol is that routed protocols are routed by routing protocols. Examples of routed protocols are Internet Protocol (IP), Transmission Control Protocol (TCP), Appletalk and Novell NetWare. Routing Information Protocol (RIP), Open Shortest Path First (OSPF) and Border

Gateway Protocol (BGP) are examples of Routing Protocols. Routing protocols use what are known as routing algorithms to determine optimal paths for movement of data packets and to fill in routing tables. The way routing protocols optimize paths is with the use of a metric. A metric is a unit of measure that is assigned to each path known by the router. The metric of a path determines its use. Each routing protocol has its own algorithm to determine the metric of a path. Some typical components of metric's are the routing protocol used, the amount of bandwidth available for that path and the number of routers between the source and destination devices (hop count).

Routing within an internetwork is a very complex and daunting task. This section of the paper is presented at a very high level of understanding to ensure that all readers have the necessary information to understand the following sections. The key points to remember are that routers operate at layer 3 of the OSI Reference Model, each used interface on a router has an Internet Protocol (IP) Address, CIDR is used to make more efficient use of IP addressing blocks and that routing protocols use algorithms to apply metrics to known data paths.

BGP Basics:

As we know, the Internet is an extremely large and complex network. There are thousands of routers spread throughout the world connected together to create the Internet. There needs to be a scalable and stable way for all these routers to communicate with each other and exchange information. The Border Gateway Protocol (BGP) is one such mechanism. BGP is the primary routing protocol used in the Internet today. In this section we will discuss what BGP is and how it accomplishes this daunting task of keeping the traffic moving on the Internet.

BGP is not new the current version in widespread use is version 4. BGP-4 is defined in Internet Engineering Task Force (IETF) RFC 1771, March 1995. This RFC defines BGP's primary function as: "For systems speaking BGP to exchange network reach ability information with other BGP systems. This network reach ability information includes information on the list of Autonomous Systems (AS)".

BGP sees the Internet as a group of Autonomous Systems (AS). An AS is defined as a group of systems (ie, routers, network devices) under the control of a single policy/administrator(s). Internet AS numbers are assigned by either the ISP an end user connects too, or by the Internet Registry Organization. An ISP will assign a customer an AS number if that customer is only routing to the Internet through them. Once an organization decides to have multiple routing paths to the Internet, or wants to have their domain routed on the Internet, that organization must apply for their own AS number through an Internet Registry Organization.

There are a number of mechanisms that BGP uses to provide the exchange of information. Two of the most important mechanisms are, classless interdomain

routing (CIDR) and attributes. BGP uses CIDR to summarize routes. From our discussion above, CIDR allows routers to read IP addresses by blocks instead of class. This ability to read blocks of IP addresses provides the routers the necessary information to combine small blocks of addresses into one entry in the routing table. Here is an example of this summarization. Let's say that an ISP owns the IP address range 200.50.0.0 through 200.50.255.255. Anytime that ISP issues a portion of the IP address range to one of its customers, it shows up in the routing table of the internet routers. These routers determine that this new routing table entry is part of a larger range of addresses and incorporates the new route into the existing route. Now the new customer is attached to the Internet without adding a completely separate route to the Internet router's routing table. This process dramatically reduces the size of the routing tables in the routes.

BGP uses seven different attributes to determine the metric of a route. These attributes as identified by Cisco Systems are listed below.

- Weight
- Local Preference
- Multi-exit Discriminator (MED)
- Origin
- AS_path
- Next Hop
- Community

Now let us take a look at the path decision process that BGP uses. This process is extracted from Internet Routing Architectures, by Bassam Halabi.

1. If the path specifies a next hop that is inaccessible, drop the route.
2. Prefer the path with the largest weight.
3. If the weights are the same, prefer the path with the largest local preference.
4. If the local preferences are the same, prefer the path that was originated by BGP running on this router.
5. If no route was originated, prefer the route that has the shortest AS_path.
6. If all paths have the same AS_path length, prefer the path with the lowest origin type.
7. If the origin codes are the same, prefer the path with the lowest MED attribute.
8. If the paths have the same MED, prefer the external path over the internal path.
9. If the paths are still the same, prefer the path through the closest interior route that originated within the local AS.
10. Prefer the path with the lowest IP address.

Finally, let us discuss how BGP peers communicate. There are four different types of messages that BGP peers send each other. The message types are:

1. OPEN

2. KEEPALIVE
3. NOTIFICATION
4. UPDATE

These message types are fairly self explanatory. When a new BGP peer comes on line the router sends out OPEN messages. While the router is running it will send out KEEPALIVE messages to let all his routing peers know that he is still around. If a router detects an error, a NOTIFICATION message is sent and the BGP connection is terminated. If one router has a change to its routing table it sends out UPDATE messages to all its peers.

With this minimal understanding of the workings of BGP, we need to move on to the primary purpose of this paper.

Security Issues with BGP:

What are the Security issues with BGP? As Information Technology (IT) professionals we all know that one of the main problems in security is human error. Everybody makes mistakes and a configuration mistake can leave our environments open to a security breach or cause systems to crash. The only real solution to this issue is training and a conscious effort to detail by us, the IT professionals.

Other than human error there are two primary security issues with BGP:

1. Sessions can be hijacked
2. Injecting incorrect information into the BGP tables

Session hijacking is when an attacker places himself/herself in between the source device and the destination device. This is also known as the "Man in the Middle Attack". This is not easy to accomplish. The attacker must know the TCP sequence number of the TCP session.

Injecting incorrect information into the BGP routing tables is also not easy, but can be done. Most ISP's incorporate filters on the routing information from their customers, but few filter information from other ISP's. An ISP that does not correctly filter its customer information or wishes to inject bogus data into the routing tables of BGP routes can do so.

Current Work in Securing BGP:

There are a few different projects under way to address these security issues with BGP. Even as early as 1998, an RFC was developed to address the issue of session hijacking (RFC 2385 "Protection of BGP Sessions via the TCP MD5 Signature Option"). This extension to BGP uses the MD5 hash to encrypt these items; TCP pseudo-header, TCP header, TCP data segment and an independently-specified key or password. So the bottom line with this protocol extension is that not only does the attacker need to know the TCP sequence number, but he must know the MD5 hash key as well.

There are three solutions that are currently being worked. These solutions are Secure BGP (S-BGP) by Bolt, Beranek and Newman (BBN), Secure Origin BGP by Cisco Systems, Inc. and BGP Scalable Transport (BST) by Packet Design.

Secure BGP (S-BGP)

Secure BGP (S-BGP) is a solution purposed by Bolt, Beranek and Newman (BBN). BBN is a research and development arm of Verizon. In December 2001, an Internet Draft, draft-clynn-s-bgp-protocol-00a.txt, was issued on S-BGP. This draft was authored by Charles Lynn, Joanne Mikkelsen and Karen Seo. The draft has expired and is no longer available.

S-BGP incorporates 3 distinct security mechanisms:

1. Public Key Infrastructure (PKI)
2. A new optional attribute, Transitive Path Attribute
3. IPSec is used to provide data and partial sequence integrity

Public Key Infrastructure (PKI) will be used to authenticate ownership of IP Address blocks, ownership of AS numbers, an AS's identity and a BGP router's identity and authorization to represent an AS. The new Transitive Path Attribute will be used to carry digital signatures authenticating the router information in a BGP update message. Using IPSec for data and partial sequence integrity allows the BGP router to authenticate each other before they exchange BGP control traffic.

All of these security mechanisms can be accomplished and provide the needed security for BGP, but there are major obstacles. One obstacle is the participation of several organizations. As you can see from the quick description given of the project above, efforts in time and money are need from Internet Registries, router vendors and ISP's. No one group is going to apply the man power or money to do this without the other groups applying man power and money also. Another obstacle is the cost of implementation. There will be a need to increase the amount of memory in the routers that are used for BGP. One estimate from www.bgpexpert.com claims as much as 4 times the memory currently used. The routers will also need more CPU power. This CPU power will be needed to do the encryption. Depending on the router this may even require dedicated crypto hardware to be added to the system.

Secure Origin BGP (soBGP)

Secure Origin BGP (soBGP) is an Internet Draft Standard purposed by Cisco Systems. There are two Internet Draft documents describing soBGP can be found at:

Extensions to BGP to Support Secure Origin BGP (soBGP)

<http://www.ietf.org/internet-drafts/draft-ng-sobgp-bgp-extensions-00.txt>

Deployment Considerations for Secure Origin BGP (soBGP)

<http://www.ietf.org/internet-drafts/draft-white-sobgp-bgp-extensions-00.txt>

soBGP is an extension to the original BGP protocol. There are six different pieces to the soBGP solution drafted by Cisco Systems. Below is a list of these six pieces as identified in the draft standard:

1. Carrying security information within BGP
2. Certificates used in the security system
3. Authenticating the identities of entities within the routing system
4. Authorizing entities to advertise given blocks of address space
5. Aggregation of prefixes within the routing system
6. Verifying the path of any given advertisement.

All six of these pieces will be implemented by the introduction of another communication message type. This new message type will be called a SECURITY message type; it will be of type 6. Type <number> is a designation that when a BGP message is received from a peer, the type is how the peers know what kind of message it has received. Under the proposed standard the SECURITY message will be used to carry three types of certification and a request format for requesting security certifications. The three types of certifications are: Entity Certification (EC), Policy Certification (PC), and Authorization Certification (AC).

The soBGP has a four step process:

1. Each Entity Certificate must be validated
2. Each Authorization Certificate must be validated
3. Information contained in the Policy Certificate must be correlated with the information in the Authorization Certification Database
4. Each prefix must be validated against the Authorization Certification Database

Deployment of the soBGP solution is not as elaborate as what is required for S-BGP. As a matter of fact, the soBGP solution can be implemented incrementally, meaning that only the two peers that will be communicating with the SECURITY messages need to have all the certification databases configured. There is still a significant amount of coordination needed between multiple organizations.

BGP Scalable Transport (BST)

BGP Scalable Transport (BST) is a solution developed by a company named Packet Design. Packet Design has taken a different approach to solving the security issues with BGP. They have developed a solution that replaces the current transport protocol for BGP, Transmission Control Protocol (TCP), with a propriety transport protocol.

TCP requires that a connection must be kept open between every pair of routers. This requirement uses up router resources by having the same data travel across the network several times. Packet Design's protocol will use a technique they are calling "flooding" to transport this data. Flooding works by sending connection

messages only to its immediate neighbors, instead of connection messages being sent from one router to all the other routers in the network. This cuts down on the number of required connections between routers and uses fewer resources. With fewer connects there will be greater scalability, less connection loss, fewer security breaches, faster convergence times and less complex network configurations.

BST, as with the other security solutions for BGP, has some obstacles to overcome. There is skepticism about Packet Design's approach and claims. Not everybody agrees that finding a solution for TCP will fix BGP. BST also does not address the issues of authentication of the origin of a BGP route or whether the data was changed in route. Finally, as a start-up, Packet Design has to sell the ISP's on its solution.

Conclusion:

In this research paper we have discussed one of the security issues facing the Internet, the Border Gateway Protocol (BGP). To understand what these security issues are, we needed to start at the beginning and discuss routing.

In our basic routing discussion we learned that routing is accomplished at layer 3 of the Open Standard Interconnect (OSI) Reference Model. We included in our discussion some information about Internet Protocol (IP) addressing. We learned that there are two parts to an IP address, the Network part and the Node part. As part of our discussion we learned about CIDR and subnet masking. From there, we discussed the differences between a routed protocol and a routing protocol.

Having this general knowledge about routing we were able to learn how Border Gateway Protocol (BGP) works. We learned a definition of BGP as presented by Cisco Systems. From this definition we started learning the functionality of BGP. In our discussions we were introduced to Autonomous System (AS) numbers and how they are used, how BGP uses CIDR for route summarization and the seven attributes of BGP. We looked at the process that BGP uses and how path selection is accomplished. Finally we discussed the communication messages that BGP uses to communicate between its peers.

Based on the general knowledge from the first two sections we were able to discuss the specific security issues with BGP. There are primarily two issues; Session Hijacking and false routing path injection.

The final section of this research paper presented the different solutions that are currently being worked to answer the security issues. One very important factor to any security solution is training and due diligence on the part of the staff performing the router configuration. Mistakes happen and the sooner mistakes are found and fixed the more secure systems will be. For complete security solutions there is one that was in Internet Draft form; Secure BGP (S-BGP). One that is currently in Internet Draft form: Secure Origin BGP (soBGP). Both of these

are solutions to the security issues, but both solutions require significant money and personnel, plus multiple organization coordination to implement. Finally, there is a solution that is being sold by a company named Packet Design, called BGP Scalable Transport (BST). BST only provides part of a solution.

In summary there is still a lot of work that needs to be done before the BGP security issues are resolved. No matter what solution(s) are finally implemented it will take a census of multiple organizations and a significant investment in both money and resources to complete the task.

References:

Y. Rekhter and T. Li. "RFC1771; A Border Gateway Protocol 4 (BGP-4)" March 1995. URL <http://www.ietf.org/rfc/rfc1771.txt?number=1771> (March 6, 2003)

P. Traina. "RFC1774; BGP-4 Protocol Analysis" March 1995. URL <http://www.ietf.org/rfc/rfc1774.txt?number=1774> (March 6, 2003)

A. Heffernan. "RFC2385; Protection of BGP Sessions via the MD5 Signature Option" August 1998. URL <http://www.ietf.org/rfc/rfc2385.txt?number=2385> (March 6, 2003)

James Ng. Internet Draft Standard. Draft-ng-sobgp-extensions-00.txt "Extensions to BGP to Support Secure Origin BGP (soBGP)". October 2002, Expiration Date: March 2003. URL <http://www.ietf.org/internet-drafts/draft-ng-sobgp-bgp-extensions-00.txt> (March 8, 2003)

Russ White. Internet Draft Standard, Draft-white-sobgp-bgp-extensions-00.txt "Deployment Consideration for Secure Origin BGP (soBGP)". October 2002, Expiration Date: March 2003. URL <http://www.ietf.org/internet-drafts/draft-white-sobgp-bgp-extensions-00.txt> (March 8, 2003)

Cisco Systems, Inc. "Routing Basics". http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/routing.htm (March 8, 2003)

Cisco Systems, Inc. "Border Gateway Protocol" http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm (March 8, 2003)

Carolyn Duffy Marsan. "Fed plan exposes 'Net's weak links". Network World October 7, 2002. URL <http://www.nwfusion.com/news/2002/1007security.html> (March 8, 2003)

Ilijtsch van Beijnum. <http://www.bgpexpert.com> , Web Site for Information and Research on Border Gateway Protocol.

BBN Technologies, "Secure BGP Project (S-BGP)". URL <http://www.net-tech.bbn.com/sbgp/sbgp-index.html> (March 8, 2003)

Packet Design, Inc. "Packet Design Solves Security, Reliability Problems of Major Internet Routing Protocol, BGP" Press Release November 4, 2002. URL <http://www.packetdesign.com/news/pressreleases/2002/pr-11-04-2002.html>

Bassam Halabi, "Internet Routing Architectures". Cisco Press/ New Riders Publishing; Copy Right 1997. ISBN: 1-56205-652-2

SANS Institute, "1.1 SANS Security Essentials I: Networking Concepts". Copy Right 2002.

The President's Critical Infrastructure Protection Board. "National Strategy to Secure Cyberspace; National Policies and Guiding Principles" DRAFT September 2002.

© SANS Institute 2003, Author retains full rights.