



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>

Is The Border Gateway Protocol Safe?

Sargon Elias
GSEC 1.4b
Saturday, 05 April 2003

Abstract

This paper is about the security issues of organisations that are planning to run their own Border Gateway Protocol (BGP) router to provide a redundant internet connection. It is aimed at a wide audience from the non technical management to the technicians who will be implementing the BGP router. For those not familiar with Internet Protocol (IP) routing and the BGP process a high level description is included in section two.

This paper includes the following sections; description of the scenario, a brief description of IP and interdomain routing, the risks when using BGP, mitigation steps and future developments.

It also aims to draw together the current thinking on BGP security. BGP has had some recent press attention as people have been looking at critical infrastructure post Sept 11. This paper should give you an understanding of the concerns that these people have in relation to BGP.

1. Introduction

BGP is like the glue that connects the thousands of individual networks that make up the internet. It is a protocol that is used by ISPs and other people on the internet to share routing information amongst each other. It allows the network to find any computer in the world and send and receive data from it. We all use BGP whether we know it or not, we can either do it ourselves or our ISPs will be using BGP on our behalf. BGP and IP routing is discussed in a bit more detail later on in the text.


1.1 The scenario for our discussion

As was mentioned in the abstract, this paper will look at the security of BGP from the perspective of an organisation that is using it to provide a redundant internet connection. This section describes such an organisation, the decision making process to use BGP and the infrastructure they will need.

As an organisation's use of the internet expands it can easily become a critical part of the organisations activity. This can entail anything from just being able to send and receive emails to hosting complex web sites and internet services. Even though failures can occur in any part of the internet, for most organisations the most critical single point of failure is the physical

link to its Internet Service Provider (ISP). To mitigate this risk there are three main approaches;

- 1.1.1 ISPs often offer service guarantees, and it is also possible to buy some special insurance which achieves a similar result. But these are of limited value as the compensation that the ISP offers for loss of service may be much less than the financial loss to the organisation when their internet connection is down. The loss of reputation can be difficult to quantify, for example if you run an on-line shopping site you could lose valuable clients to your competitors when your web site is unavailable.
- 1.1.2 Improve the reliability of the connection to the ISP by adding redundancy to the link. Many ISPs will offer redundant circuits and equipment and may even offer redundant connections to different Points Of Presence (POPs) in their network. Again these measures can be effective and are relatively easy to implement. The main downside is that you are held hostage to the ISPs pricing policy. It can be very expensive and if your ISPs connection to the rest of the internet fails you still have an outage.
- 1.1.3 Get several connections to the internet via different ISPs. This can protect you from physical connection problems as well loss of service via one of the ISPs. This is assuming your ISPs don't share the physical lines out of your building or critical infrastructure near to the connections point. This is something to watch out for as ISPs often subcontract the physical circuit to local telcos. This is the best solution and the one I will discuss in this paper. The main issue we need to address with this approach is routing.



If you purchase internet provision from two ISPs, they will both allocate you some IP addresses from their own ranges. It means that if you provide

services to the internet, you will have to provide them on both address ranges. This may be relatively simple as you could duplicate your web site onto two web servers, and have one DNS (Domain Name Service) pointing to both addresses. But it may be impractical if you have a complex interactive web site or other internet services. It may be particularly impractical if your service relies on statefull connections. In this case the loss of one line will still leave your clients disconnected and having to re-establish a connection to the service on the other address.

As for outbound traffic, you will have to decide which of your two lines to use as your default route out. If you use Network Address Translation (NAT) 'hide' in your firewalls each firewall will have to hide behind an address that is specific to the line. Firewalls cannot hide behind two different addresses and will not be able to match traffic coming in from one line to outbound traffic on the other line with a different address. So you will need two firewalls one for each line. This means that if one line fails the traffic cannot flow back over the other line as it will have a different IP address. When such a failure occurs you are faced with having to reconfigure your internal users to use the other firewall as their default router. You may be able to automate this using internal routers and VRRP, but you need to be careful not to create a single point of failure inside your network.

The alternative and much better way to achieve the redundant connection with multiple ISP is to get your own address space and use BGP to route traffic to it. Then no matter which line is lost the traffic will automatically switch to use the other line without any manual intervention. There are a couple of additional advantages.

- a) Any existing TCP/IP sessions you have when the line fails will not be broken. The BGP router will simply reroute traffic via the working connection. This happens relatively fast, the convergence time for BGP routes around the globe is typically less than 30 seconds.
- b) It makes no difference where the break occurs. If it is your local physical connection to your ISP, their connection to the rest of the internet or anywhere else on the internet between your customers and your systems, the BGP routers will simply route round the problem.

Once you decide that running your own BGP router is the way to go, you will need to set up the BGP router and connect it to the border routers of your ISPs. These are known as your 'peers' in BGP speak. You will then need to request your own address space and your own BGP Autonomous System (AS) number from the local internet registry organisations, in Europe this is RIPE¹. See the footnote for links to the other registry organisations. Of course as you are trying to make your internet connection redundant you should get two routers to work in parallel, so you don't move the single point of failure to your router.

¹ RIPE is one of the Local Internet Registry (LIR) organisations, there are currently four geographical registries which are defined by IANA, see <http://www.iana.org/ipaddress/ip-addresses.htm> for details

Your ISPs may well charge extra to set up BGP, although this is always negotiable. But once it is in place you should be able to negotiate your bandwidth rates down as it will be much simpler to swap to another ISP. Also your ISPs reliability is less important as you will have redundancy through your use for BGP, so you could go for cheaper ISPs. This is a real benefit in reducing your internet costs. Many ISPs rely on continuing business at higher than market rates because it is difficult for a customer to switch ISPs as they face all the hassle of reassigning their IP addresses.

A word of caution about selecting your ISPs, ISPs are split into tiers. Tier 1 are large organisations that have their own global network. Tier 2, are large regional ISPs. Tier 3, are small local ISPs that will use the tier 1 or tier 2 networks for transit. The bandwidth provision business is very cut throat and many of the ISPs are barely able to stay in business. Make sure your ISP has good technical knowledge in supporting BGP customers. Some of the small ISPs simply resell bandwidth and may not have much experience with providing BGP service to its clients.

The main disadvantage in using BGP is that it requires you to purchase and run your own BGP router. There are the security issues, which I will discuss later, but also day to day management work that would otherwise be done by your ISP.

Now that I have discussed the benefits of BGP, we will briefly look at how BGP works and then look at the security implications. We introduce risks because we manage our own BGP router and there are the risks that have always been there because our ISP uses BGP on our behalf.

2. How does BGP work?

2.1 The internet Protocol

The internet is based on the TCP/IP (Transmission Control Protocol/Internet Protocol) which has two interesting properties that you don't see in more traditional, centrally controlled networks such as the telephone system.

- a) There is no established channel between two end points during communication session, data is broken in packets and each packet gets routed to the destination independently. This means packets can take different paths through the network to the same destination.
- b) The second is there is no central exchange; instead the internet is made up of a large number of independent networks linked together in a mesh structure. There is a common understanding that they will carry traffic for others on their own network in exchange for other networks carrying the traffic for its user. They may connect at what are known as internet exchange points such as LINX in the UK or they may just link directly to each other through bilateral agreements.

2.1 IP Routing

As there is no central exchange, all these independent networks need a way to know where any other computer is on the internet. They need a kind of map to decide which way to send traffic at each inter-connection. This is known as routing in the IP world. Every end point on the internet has a unique IP address which is allocated via the internet registry organisations, see footnote¹. However they don't dictate where on the network these IP addresses are used. The key to making the routing work is for each network to know how to forward the packets in the right direction.

To send data between two computers the TCP/IP protocol specifies that the data is broken up into packets. An IP header is added to the data packet which contains the destination address. The packet is then sent out onto the network that the computer is attached. Specialist network equipment called routers will then forward these packets on to their final destination. The forwarding decision is based on its routing table. These are built up either by manual configuration or by receiving routing information from other routers. To find out more about IP routing see "TCP/IP Illustrated Volume 1" by Richard Stevens² chapter 9 or "Internetworking with TCP/IP" by Douglas Comer³ chapter 8. Both are good references to not only routing but the whole TCP/IP protocol.

Each router does not necessarily know where all other computers are on the network, but knows instead which direction to forward the packets to somebody who will know. A good analogy is the postal service. If a packet has a local address the router will deliver it directly, if it has an address that is served by one of the routers directly connected it will pass it to that router for final delivery. In our analogy with the postal service, this would be a letter destined for the same town or district; the local sorting office will recognise the address and place it in a sack for delivery by the appropriate postman. If it recognises the network part of the address it will pass in the direction of the target network even though it does not know where the actual address is. In our analogy, this would be a letter for another town, where the sorting office knows the town but not the street or building and will put it into a sack destined for a sorting office in that town.

Some routers in the same way as some out of the way post offices will send all their non local traffic to a larger main post office. In the router world this is a default route i.e. any traffic I don't know how to deliver I will forward to my default router. It will then route it to the correct destination or pass it on to its default router. At the core of the internet there are a set of routers that know about all other networks on the internet. These are the BGP routers; they have a table of all networks and how to get there.

² Stevens, Richard. "TCP/IP Illustrated Volume1". Addison-Wesley. 1994

³ Comer, Douglas. "Internetworking with TCP/IP, principles, protocols and architectures, Fourth Edition". New Jersey. Prentice Hall. 2000

If you want to send a large amount of information between two computers, the data is broken up into packets and each one is addressed separately. The routers on the way do not remember the previous packets and will make new routing decisions for each packet they receive. It's a bit like sending a long letter by sending a string of postcards, the postal service will deliver them all individually and some might take other routes to the destination than others such as the postman walking a different route to your house. It is up to the computer at the other end to reassemble the packets into the original data. This might seem strange, but it provides a number of benefits. Each router in the network does not have to keep track of connections between machines, thus making them much simpler to build routers. And the network can dynamically reroute traffic around failed connections making it very robust to failed lines or routers.

2.2 Routing Protocols

One of the problems with running a network is how the routers find out about the routes. For routers at the periphery, they simply use a 'default route' to send the traffic to the centre. However at some point we need some routers to know where all the other networks are. At its simplest a network administrator can manually configure each router with a table of all the networks and how to get to them. However, this soon becomes a big overhead and difficult to maintain. To make this more manageable 'routing protocols' were invented. These allow routers to share routing information with each other and automatically distribute changes as they occur.

These protocols are split into two types, internal routing protocols and external routing protocols.

- a) **Internal routing protocols** are used when a network is under the administrative control of one organisation. Here each router will learn all the routes in the network from the other routers on the network, effectively building a complete topographical map of the network. It can then use this to accurately route packets to the final destination. There are number of different protocols to achieve this but the difference between them is primarily based on how they communicate the information and how efficient the routes are that they produce. Examples of Internal routing protocols are RIP2 (Routing Information Protocol version 2) and OSPF (Open Shortest Path First).
- b) **External Routing protocol**, this is what we are interested in here. The one that is currently in use on the internet is BGP version 4, known simply as BGP-4. This protocol is used by organisations to exchange routing information with others outside their own network. The basic unit for external routing is a network managed by a single organisation, this is known as the Autonomous System (AS) in BGP speak. Each AS may well use an internal routing protocol within its own network to distribute routes, but when it wants to share routes with other AS's it will use BGP. These networks come in all shapes and sizes from very

large geographically dispersed networks to a single site with two connections.

2.3 BGP

Unlike internal routing protocols, it soon became clear it would be impossible for one router to know about all the machines on the internet and so BGP does not attempt to build a map of the internet. Instead each Autonomous system (AS) is allocated a unique number and advertises the routes it knows how to get to in terms of these AS numbers. The IP addresses are aggregated up as much as possible in these advertisements so that blocks of addresses are advertised.

For example, a BGP router may know that a local route exists to address range 10.0.0.0/8. It can send a single advertisement to its peers which will then pass this on to their peers with their own AS number attached and they to theirs and so on. Eventually all BGP routers know that if they have a packet destined for any address starting 10 they can pass it to the BGP router who made the advertisement via all the other networks that have appended their own AS numbers. There may be several routes to this network and the BGP router will decide which to use based on the least number of ASs.

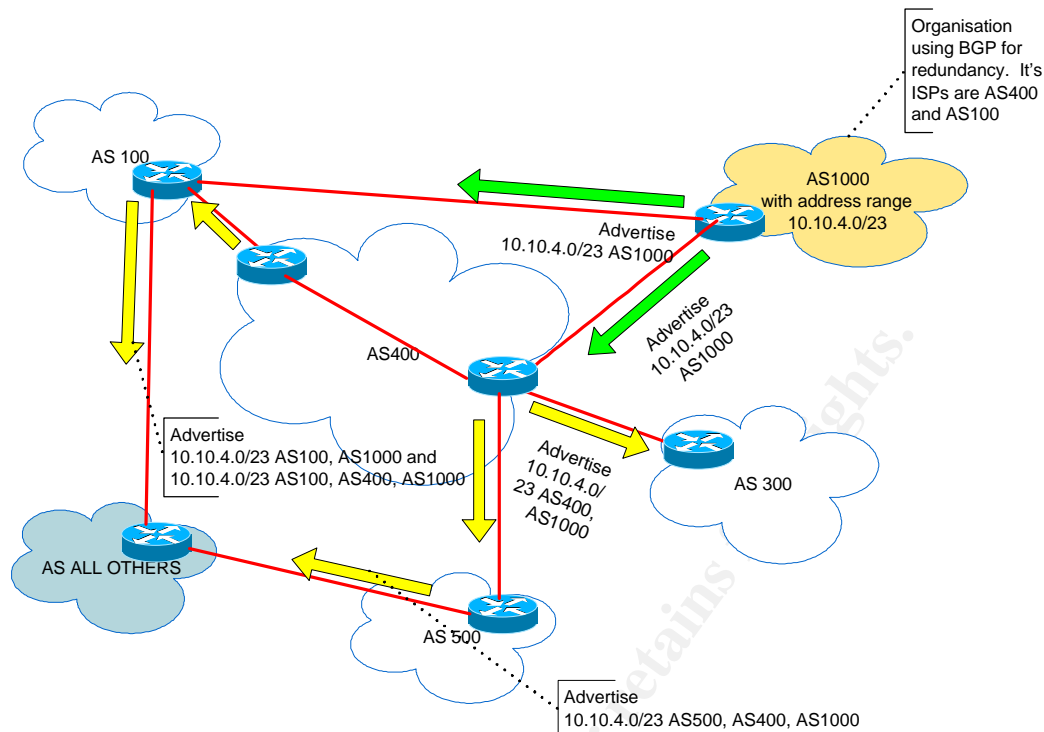
When two network providers decide to peer, they agree to set up a BGP link between each others border routers. The routers are configured to establish a BGP session between each other. Once the session has been established the BGP routers will exchange all the routes that they know about and send updates as and when they happen. The BGP routers stay connected so that if there is a failure on the link, the BGP routers will notice and stop using the routes advertised from this link. It will also withdraw those routes it has advertised based on the routes from the failed link.

All internet protocols are described in what are known as Request For Comments (RFC). The BGP protocol is no exception and is described in RFC1771⁴ which can be found at <http://www.ietf.org/rfc/rfc1771.txt>

To illustrate the BGP process the following diagram shows an example of 5 Autonomous Systems and how they are connected together. Each cloud represents a separate AS that may contain lots of computers and one or more internal router. However all traffic in and out of the network will go via that networks border router running BGP. In this example we are AS1000 and use AS100 and AS400 as our ISPs.

Note that ISP AS400 has two BGP routers, these would typically exchange information via an internal routing protocol or a version of BGP called IBGP.

⁴ Rekhter, Y; Watson, T J; Li, T. "A Border Gateway Protocol 4 (BGP-4)". March 1995. <http://www.ietf.org/rfc/rfc1771.txt>. (April 2003)



When there are multiple routes, as in the one that AS100 advertises for our address range 10.10.4.0/23 the router will decide on the best route based on the number AS hops. Note, there may be local preferences that override this decision making. As you can see in the diagram, the green arrows show us advertising our address range to our ISPs, AS100 and AS400, the yellow arrows show them propagating those routes on to their peers with their own AS number attached.

3. What are the security risks?

The risks can be split into two main categories. Attacks related to us running our own BGP router and attacks directed at the BGP protocol.

3.1 Attacks on our Border Router

By definition the BGP router will be outside our firewall and needs to be available at all times for internet connectivity to be maintained.

The BGP router can be a dedicated device such as a Cisco router or a general purpose computer running a BGP routing package such as Linux with Zebra. It could even be our firewall itself; the Nokia Firewall platform (IPSO) for example has BGP built in. However from a security perspective it is better to separate services and not use one machine for several tasks. Any of these platforms can have vulnerabilities in the Operating System and services they run. Such a vulnerability could lead to the usual risks when one of our external hosts is compromised - loss or destruction of data and loss of service. This is compounded because it will effect our entire internet

connection. The attacker could disrupt all our traffic or target particular traffic. It would also be an ideal place to intercept traffic, as all our inbound and outbound traffic will flow through that point.

Should our BGP router be compromised it could also allow an attacker to poison or corrupt the routing tables of our peers. If they don't filter their routes carefully this could create black holes in the internet for their own and other connections.

3.2 Attacks on the BGP protocol

The other set of risks concern the attacks against the BGP protocol itself.

Information that is received from our peer router is simply trusted to be correct. If it says it can get to a particular network our router will simply trust that information is correct. We also trust that any routes we advertise will be forwarded to other networks that our peer is connected to, so that those networks know how to get back to us.

One of the fundamental designs of the TCP/IP is that routers make independent routing decisions for every IP packet and so can dynamically re-route traffic around a network failure or congestion. This has been one of the strengths of the TCP/IP protocol and came from one of the early desires by the military that the destruction of a single node in the network would not lead to complete network failure. However, this does mean it is possible for an attacker to inject data into an open TCP/IP stream. This is not a trivial attack and not something you would see a script kiddie do, but can be achieved.

The most widely discussed attack is spoofing of the BGP peer. In the book "Building Reliable Networks with the Border Gateway Protocol", Ijitsch van Beijnum describes it as follows, when discussing attacks against BGP:

"...to take over a peer's IP address and present themselves as your peer. The routers will then set up a BGP session, and the attacker can inject disruptive information in your routing tables, unless filters for this peer are strict. The attacker may even route some of your network address space to himself and present himself to the outside world as a host on your network, so he can receive your email and web requests.

...⁵

The second method that he discusses is breaking the BGP connection by injecting RST packets into the TCP/IP stream. As we discussed earlier, BGP maintains an open connection to each peer via the use of 'heart beat' messages. If the connection is broken the router will assume that any routes advertised via that link are no longer available. It removes them from its

⁵ Van Beijnum, Ijitsch. "Building Reliable Networks with the Border Gateway Protocol". Sebastopol. O'Reilly & Associates Inc. September 2002. p135-136

routing table and sends a withdrawal update to its peers. If an attacker disrupt the BGP connection via rogue RST packets this it would be a simple and effective Denial of Service (DOS) attack against the whole site.

Other operational aspects such as route flapping can be turned into security risks. If a route advertisements changes frequently in a short period of time it is said to flap. Many ISP have flap dampening policies that will ignore such changes after a certain threshold is reached. This could again black hole our site if an attacker can cause frequent route changes even if they are corrected quickly by our own BGP router.

The infamous AS7007 case shows what can happen if incorrect BGP routes are advertised. This was not an attack but illustrates what can happen if incorrect routes are advertised. The case was described by Rik Farrow <http://www.spirit.com/Network/net0102.html> as follows:

“On Friday morning, April 25 of 1997, a small ISP in Florida made a mistake in the configuration of the router that joined their small network to Sprint. This ISP, known by their AS number, 7007, allowed all the routes learned from Sprint using BGP to be exported back to Sprint as their own routes. This actually is easy to do, as BGP implementations can take routes from IGP and convert them into EGP routes. In this case, the IGP converted CIDR routes into classful routes.

The Sprint BGP speaker was not filtering properly either, and began sending out updates that added AS7007 as the correct route for a portion every CIDR block (essentially, the first class C, or 24 bit long, network prefix).

This misinformation first spread through Sprint's network, then to neighbouring NSPs, including ANS, MCI, UUNet, and other NSPs. Many routers crashed, as their routing tables suddenly doubled in size (an additional route added for each CIDR block), and the routing instability spread throughout the Internet. Remember that when a router crashes, it drops its BGP connection with its peer, who then sends out an update withdrawing all the routes previous announced by the crashed router. It took over an hour for the Internet to gradually become stable again. Network managers added filters that blocked routes that included AS7007, fixing the problem until the ISP solved their local problem and Sprint reconnected them to the Internet.”⁶

This has sometimes been used to illustrate the dangers to BGP. Post September 11 a lot of discussion has gone into the risks to the internet which is seen by many as “critical infrastructure”. An attack on the lines of the AS7007 incident has luckily not happened yet, however this does not mean that it can't or won't happen in the future.

⁶ Farrow, Rik. “Network Defence, Routing Instability, Border Gateway Protocol, the routing glue of the Internet, Lacks Strong Security”. 2002. <http://www.spirit.com/Network/net0102.html>. (April 2003)

4. Mitigation

As we discussed during the previous section there are two areas to consider when protecting our BGP router. Protecting your BGP routers against host attacks and protecting against attacks of the BGP protocol.

4.1 Protecting our BGP router

The first is arguably easier to deal with. It is similar in many ways to building a bastion host that runs any other internet based services such as mail or web servers. Most administrators are familiar with this work and a lot of organisations will have policies in place that describe how to build and maintain a host exposed to the internet. However the mechanics and syntax may be unfamiliar to those administrators who deal only with Linux, Unix and Windows based hosts.

4.1.2 Host level protection

As with any machine that provides a service on the internet it is important to harden the OS and applications. Only run the minimum number of services and use the security facilities that have been built in. Often the BGP routers are dedicated devices such as those produced by companies like Cisco. However more and more organisations are using general purpose computers running Linux or BSD based UNIX to provide these services. This has many cost advantages but will need even more careful configuration as there are more services and packages that can lead to compromise.

When using a general purpose computer a package such as Zebra can be used to provide the BGP routing facility. Some versions of Unix will also allow the OS to do this natively. Zebra <http://www.zebra.org> is an Open Source based routing package that is specifically designed to work with internal and external routing protocols. This has a user interface which is very similar to the Cisco IOS interface and uses similar security facilities.

When using a general purpose machine instead of a dedicated router device there is an additional task to make sure the OS is locked down using the normal guidelines for any internet facing service. However, using sensible rules this can be just as effective as a dedicated router device.

In general the things to consider from a security perspective are:

1. How the administrator is going to access the router.
2. How the administrator authenticates with the router.
3. How passwords are stored.
4. A patching process for when vulnerabilities are discovered.

Routers are traditionally managed remotely via telnet, the problem with telnet is that it is unencrypted and passwords are passed across the connection in clear text. For this reason all remote access should be switched to secure

shell (SSH) and telnet disabled. This was not always possible on older versions of Cisco's IOS. Therefore it is important to use one of the latest versions of the IOS. On Unix and Linux, ssh has been available for a while and recent distributions make ssh the default shell access method moving away from telnet.

In IOS and Zebra, there is no real concept of user accounts, just levels of privilege. To connect to a router you provide the 'virtual terminal' (vty) password, this will put you into a read only mode. You then provide an 'enable' password which is like a read/write mode where changes can be made. The enable password is similar to the root account on Unix and should be treated with the same care. A strong password that conforms to organisation password policy should be chosen for both of these functions. It can be tempting to have a simple password, as these machines may not be accessed very often but this could be a fatal mistake, as the damage that can be done by compromising your BGP router could be worse than a compromise on any single internet based service that you run. It is an ideal location from which to sniff all your organisations inbound and outbound traffic and would allow an attacker to completely stop traffic to and from your site with ease.

On the default IOS install no passwords are set, make sure you set a remote access passwords. Usually the configuration sets a maximum number of simultaneous logins to 5. To set the virtual terminal passwords for these, use the following commands:

```
Router 1(config) # line vty 0 4
Router 1(config) # login
Router 1(config) # password yourpassword
```

The IOS traditionally stored passwords unencrypted; you should make sure you configure your router to store them encrypted so that anybody gains access to the router or its configuration file cannot easily obtain the enable password. This can be achieved with the following command

```
Router 1(config) # service password -encryption
```

It is critical that we set up an effective security patch update procedure. The administrator of the BGP router needs to keep up to date with the latest vulnerabilities and software patches and implement these as soon as possible. This may involve having a test system to make sure the patch works and does not break the existing routing before applying it to the production router.

4.1.2 Network protection

Because the BGP router will sit outside the main firewall, it cannot rely on the organisations firewall for protection and will need to protect itself from the usual scans and attacks that we see on the internet everyday. This also means the traffic will be visible on our external network where it may be sniffed. In the previous section we discussed the implications of this for our

management traffic and we suggested using ssh instead of telnet so that this traffic is encrypted.

IOS and Zebra also allow us to set up Access Control Lists (ACLs) which are in effect IP filters working in the same way as a personal firewall. These are one of the main weapons in defending our router. These should be carefully configured to allow only the minimum amount of access. As a guideline the access should be along the lines of; allow ssh from management station, allow BGP access from peers, allow SNMP access from network monitors (if you use SNMP), allow certain types of ICMP messages to help diagnose problems and block everything else.

ACLs should be set up to achieve this, so we would restrict admin access to the administrators IP address, allow BGP access only from the IP addresses of our peers. If we use SNMP (see discussion in section 4.1.3) allow this only from the network monitor IP address. All other traffic should be blocked, with the exception of certain ICMP packets which are required for smooth operation. Such as pings Echo and Echo Reply (type 0 and 8), and router responses such as 'Destination unreachable' (type 3) and 'Time Exceeded' (type 11). All other ICMP packets should be blocked as these do not make sense for our BGP router and may allow an attacker to disrupt the normal operation. It is particularly important to switch off and block any other routing protocols we don't use, as this could accidentally become a backdoor for an attacker to inject bad routes.

For example to only allow access only to our peers on the BGP port we could use the following Cisco IOS command, where our peers are at 10.10.5.1, 10.10.10.1 and 172.17.70.2

```
!
! We protect TCP port 179 (BGP port) from miscreants by limiting
! access. Allow our peers to connect and log all other attempts.
! Remember to apply this ACL to the interfaces of the router or
! add it to existing ACLs.
access-list 185 permit tcp host 10.10.5.1 host 172.17.70.1 eq 179
access-list 185 permit tcp host 10.10.10.1 host 172.17.70.1 eq 179
access-list 185 permit tcp host 172.17.70.2 host 172.17.70.1 eq 179
access-list 185 deny tcp any any eq 179 log -input7
```

4.1.3 Network Monitoring via SNMP

It is common to monitor network equipment, such as routers, via Simple Network Management Protocol (SNMP), many network management packages rely on this protocol to obtain statistical information and detect problems via alerts. However vulnerabilities in the SNMP protocol are well known and heavily publicised. See the CERT advisory CA-2002-03 as an example of the type of vulnerabilities. (<http://www.cert.org/advisories/CA-2002-03.html>)

⁷ Thomas, Rob. "Secure BGP template version 2.0". 16 Oct 2002.
http://secinf.net/firewalls_and_VPN/Secure_BGP_Template_Version_20.html. (April 2003)

Many organisations have a policy to block all SNMP traffic at the firewall. This provides a dilemma for monitoring our BGP router which will by necessity be outside the firewall. The monitoring process is very important to help provide smooth operation of the router and detect problems and incidents quickly. However it's the very same SNMP that could make our router vulnerable to being compromised.

The risk should be evaluated in each case, and it is up to each site to decide if the risks are acceptable, if you do decide to use SNMP you could moderate the risks as follows:

- a) Make sure that you disable write access and only use SNMP for monitoring purposes. This means that even if the community name is compromised it will only provide read only access. Although this will not protect the router from exploits in the SNMP service itself.
- b) Set a good community string. Despite its name you should consider this as a password and treat it the same way. Don't use the same string for devices inside your network as your BGP router, again don't be confused by the name most network management software will allow you to specify different community strings for each device that you want to monitor. Above all make sure you remove all the default community strings such as 'public' and 'private' from your router.
- c) Set up ACLs to allow only the network management stations IP address to connect via SNMP. Obviously this still leaves us open to IP spoofing and as SNMP uses UDP packets spoofing is somewhat easier to accomplish than for TCP/IP. Despite this using any security measures will make the attacker's job harder and will discourage script kiddies who will simply move on to easier target. It should also stop attackers running SNMP based network scans from picking up your router.
- d) If possible use SNMP version 3 where you can use strong authentication and encryption. However this is still not available for all devices or network management software.
- e) Keep your router patched with the latest patches.

4.2 Protection from abuses of the BGP protocol

For our BGP router to work we trust our peers to provide all the routes that they know about. And we need to advertise our IP address range to them with our AS number. If we are an organisation that only has a couple of internet lines and uses BGP to provide redundant internet access, we would only be advertising our address space and expecting all other routes to be provided by our peers. In this case we do not provide any transit traffic so will never advertise other networks to our peers.

4.2.1 Invalid Routes

If a peer sends invalid routes to us it in effect will stop us being able to contact those IP addresses as we will pass traffic to a router that is not able to forward

it to the correct destination. In effect we 'black hole' those addresses ranges for our users. If our peer does not pass on our route, it will 'black hole' our address space.

The primary weapons against these abuses are filters. We can set up a filter to remove any route that does not make sense or should not be advertised by a particular peer. We should use both ingress and egress filters, cleaning what we send out as well as what we receive from our peers.

We can sanity check a number of the route advertisements as follows:

1. Check that we only send out advertisements containing our own address block. An egress filter can be set up for this.
2. Assuming we don't provide transit to anybody else, check we don't advertise any other address blocks. A similar egress filter as the one above.
3. In any advertisements we receive check that the IP address range belongs to the first AS in the advertisement. Some ISPs have automated scripts that check the AS numbers to address blocks based on downloads from the internet IP address registrations organisations RIPE, ARIN, LACNIC and APNIC. The filter then checks the advertisements against these lists. This means it is very important that any address range we have is properly registered and that the entry in the 'whois' database is correct. This ingress filter can be very effective, but assumes that an attacker is not able to infiltrate the registration organisations or affect the automatic download processes.
4. Discard routes for reserved address blocks such as (10.0.0.0 to 10.255.255.0), (172.16.0.0 to 172.31.0.0) and (192.168.0.0 to 192.168.255.0). These should never appear on the internet, an ingress filter can get rid of these.

You should discuss with your peer what their policies are and try to make sure for every egress filter we have, they have an equivalent ingress filter and visa versa. This will stop most attacks where an attacker spoofs our or our peers IP address and tries to inject obviously rubbish routes.

4.2 Authenticating your peers

The fact that BGP uses TCP instead of UDP makes it a little less prone to session hijacking, and spoofing of messages from attackers. However, a number of attacks are still possible that allow an attacker to pretend to be a trusted peer. One such attack is the man in the middle attack, where the attacker inserts himself in between the two peers, this is hard to achieve unless the attacker can insert himself close to your router or the peer's router as traffic can flow in many different directions across the internet. The only thing that helps the attacker here is that often the BGP routers will be at different locations under the physical protection of different organisations. It may be that the physical security of your peer or any subcontracted telcos is

not as good as your own. This may allow an attacker to insert a computer in the link without you knowing about it.

When setting up a connection with a peer you are always dealing with external organisations and getting it to work at all can be difficult and time consuming. Different organisations have different policies and procedures and it may be difficult to speak to one individual if you are dealing with one of the larger ISPs. This can lead to using the lowest common denominator when it comes to authentication and encryption. Both strong authentication and encryption can protect the BGP session from many of these attacks.

You should discuss this with your peer, to see what they are willing to do in this area. If you are paying them to have the peering set up, you should be in a better position to demand that strong authentication and or that encryption is used. There are some interesting proposals in this area, see the section 4.3 for details.

4.3 Future development

There is a constant development in this area, and particularly post September 11 a lot of discussion has been around protecting critical infrastructure such as the Internet. BGP is a critical component of this. We have been lucky up to now that there have been no significant security related incidents related to BGP, but the possible impact of an attack is so significant we should not be complacent. The internet community is coming up with a lot of new proposals regarding BGP security all the time; we discuss a few here that may be significant in the future.

4.3.1 S-BGP

S-BGP (secure BGP) is a proposed version of BGP that includes strong authentication and encryption using public key infrastructure. Read the work by the BBN Technologies, Internetwork Research Department <http://www.ir.bbn.com/sbgp/>⁸ for a detailed explanation of the proposal. This has been around for a few years in discussion stages and a number of prototype network have been successfully trailed.

However a technological chicken and egg problem has so far stopped it from being deployed in the real world. The main issues are that there are several parties involved that all need to agree. The Internet Engineering Task Force (IETF), the registration organisations and the ISPS. The ISPs will need to invest in new technology and spend money to implement this. As this will not offer new services to it clients they cannot directly charge their clients for this work. This makes them very reluctant to invest in these changes. The registration organisations need to agree to set up a digital signature Public Key Infrastructure (PKI) again without ISP buy in they are reluctant to invest the time and effort. The hardware manufactures have not really bought into

⁸ BBN Technologies, Internetwork research Department. "Secure BGP Project (S-BGP)". <http://www.ir.bbn.com/sbgp/>. (April 2003)

the new process and have been reluctant to include support for SBGP in their platforms.

The dilemma is that there is a high cost to secure BGP and avoid any major incident, but unless an incident occurs the parties involved are reluctant to spend that money.⁹

4.3.2 Ptomaine

When you advertise a route to a peer you have little or no control over how this is distributed or how other ASs will use it in route decision making. Ptomaine proposes a new external community mechanism that allows a route advertisement to include filters so that we can control how this is passed upstream to other ASs. As we know prefix filtering is one of the main security techniques and anything that allows this to be extended this will improve the security of BGP. It would be possible for example, to not only create egress filters to stop us advertising other IP addresses than our own, we add this filter to the route advertisement so that upstream ASs so they don't propagate invalid routes for our address space. There is an internet draft describing the proposal at <http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ptomaine-bgp-redistribution-00.txt>¹⁰

4.3.3 Sharing filter policies

As we have already discussed using ingress and egress filters are an important tool in securing the BGP protocol. There is a new proposal that automates this process of sharing your filters with your peers. The proposal specifies the syntax for allowing filters to be sent to a peer via the BGP session. See the internet draft <http://www.ietf.org/internet-drafts/draft-ietf-idr-route-filter-08.txt> for more details.¹¹

⁹ Vamosi, Robert. "Router security hole threatens web". 3 March 2003. <http://news.zdnet.co.uk/story/0,,t286-s2131302,00.html>. (April 2003)

¹⁰ Bonaventure, Olivier; De Cnodder, Stefaan; Hass Jeffrey; Quoitin, Bruno; White, Ross. "Controlling the redistribution of BGP routes". April 2002. <http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ptomaine-bgp-redistribution-00.txt>. (April 2003)

¹¹ Chen, Enke. "Cooperative Route Filtering Capability for BGP-4". Jan 2003. <http://www.ietf.org/internet-drafts/draft-ietf-idr-route-filter-08.txt>. (April 2003)

References

- 1 Barry, Green, "Is the Sky Falling In". June 2002, <http://www.cymru.com/Presentations/barry.pdf>. (April 2003)
- 2 Stevens, W. Richard. "TCP/IP Illustrated, Volume 1". Reading: Addison Wesley Longman, Inc. 1994.
- 3 Comer, Douglas. "Internetworking with TCP/IP, principles, protocols and architectures, Fourth Edition". New Jersey. Prentice Hall. 2000
- 4 Rekhter, Y; Watson, T. J.; Li. T. "A Border Gateway Protocol 4 (BGP-4)". March 1995, <http://www.ietf.org/rfc/rfc1771.txt>. (April 2003)
- 5 Van Beijnum, Ijitsch. "Building Reliable Networks with the Border Gateway Protocol". Sebastopol. O'Reilly & Associates Inc. September 2002. p135-136
- 6 Farrow, Rik. "Network Defence, Routing Instability, Border Gateway Protocol, the routing glue of the Internet, Lacks Strong Security". 2002 <http://www.spirit.com/Network/net0102.html> (April 2003)
- 7 Thomas, Rob. "Secure BGP template version 2.0". 16 Oct 2002. http://secinf.net/firewalls_and_VPN/Secure_BGP_Template_Version_2_0.html. (April 2003)
- 8 BBN Technologies, Internetwork research Department. "Secure BGP Project (S-BGP)". <http://www.ir.bbn.com/sbgp/>. (April 2003)
- 9 Vamosi, Robert. "Router security hole threatens web". 3 March 2003. <http://news.zdnet.co.uk/story/0,,t286-s2131302,00.html>. (April 2003)
- 10 Bonaventure, Olivier; De Cnodder, Stefaan; Hass Jeffry; Quoitin, Bruno; White, Ross. "Controlling the redistribution of BGP routes". April, 2002", <http://www.ietf.org/proceedings/02jul/I-D/draft-ietf-ptomaine-bgp-redistribution-00.txt>. (April 2003)
- 11 Chen, Enke. "Cooperative Route Filtering Capability for BGP-4". Jan 2003. <http://www.ietf.org/internet-drafts/draft-ietf-idr-route-filter-08.txt>. (April 2003)