the Availability Digest

www.availabilitydigest.com

Eavesdropping on the Internet

March 2009

At the 2008 DEFCON hackers' conference, security researchers Anton Kapela and Alex Pilosov demonstrated a fairly simple technique to divert Internet traffic to an eavesdropping site on its way to the intended recipient. Though this vulnerability had been earlier predicted, no one had ever been able to demonstrate it. Kapela and Pilosov showed that they could eavesdrop on DEFCON traffic with their own server to the cheers of the hacker crowd.

Eavesdropping can be used by corporations for competitive purposes and by governments for surveillance purposes. Messages can even be modified in transit by the eavesdropper.

There is no easy way to detect that your traffic is being monitored, nor is there any easy way for the Internet to prevent this sort of attack. Your best bet is encryption so that your traffic has no value to an eavesdropper.

The problem is in a vulnerability of BGP, the Border Gateway Protocol. BGP is the routing protocol used to distribute global routing information throughout the Internet. It is used by all major ISPs as well as by many smaller providers and other organizations.

In this article, we review those elements of BGP that lead to the possibility of eavesdropping. But first, we review some pertinent characteristics of the Internet.

Internet Addressing in Review

If you mail a letter in Germany to an address in Boston, the German postal service doesn't care about the Boston address. All it cares about is that the letter is going to the U.S. and forwards it there. The U.S. distribution center doesn't care about the street address. All it cares about is the zip code and forwards it to that post office. The post office sorts the mail and loads all mail for a particular route onto the mail truck servicing that route, but it doesn't care where the mailbox is located. The mail-truck driver is the final link in the delivery process and puts the letter in the appropriate mailbox.

The delivery of messages over the Internet uses the same hierarchical strategy. The Internet hierarchy is three levels - networks that comprise subnetworks (subnets) that serve hosts. Each host has a unique IP (Internet Protocol) address that in IP Version 4 (IPv4) is a 32-bit address comprising four 8-bit octets (an *octet* is Internet-speak for bytes).[1] IPv4 addresses are usually expressed in *dotted decimal notation*, such as 161.35.1.19.

[1] A 32-bit address space provides for four billion users. However, this address space is nearing exhaustion. The new IPv6 extends the address space to a 128-bit address – 16 octets – that provides 10^{38} addresses. Though most vendors' equipment now supports IPv6, less than 1% of the ISPs currently do.

In the IP address, the two higher-order octets specify a network. The two lower-order octets specify a subnet in the higher-order bits and a host in the lower-order bits. Often, the third octet is a subnet address; and the forth octet is the host address; but the number of bits used by each is determined by the network administrator. A 32-bit subnet mask containing all higher-order one bits specifies which bits are the network/subnet addresses (the first sixteen bits are always the network address) and which bits are the host address (the zero bits). For instance, a subnet mask of 255.255.255.0 indicates that the subnet address is octet three; and the host address is octet four.

The network and subnetwork addresses (i.e., those bits specified by the subnet mask) are called the IP address *prefix*.

Internet Routing in Review

At the simplest level, the Internet can be thought of as a series of networks interconnected by routers. A message is passed from network to network by routers until it reaches its destination.

For instance, when you request a web page from your browser, you are asking for a page from www.anypage.com. URLs are textual and easy to remember. However, the Internet needs an IP address to forward your request to the appropriate web server. It is the Domain Name Service (DNS) that provides this translation. Your browser accesses a DNS server to obtain an IP address for the web server.

This address is forwarded to a router to which your browser is connected. This may be a corporate router or a router provided by your ISP. The router looks at the network address and decides the best route, based on its internal routing tables, to send the request. This may send the request to another router that may forward the request to yet another network.

Eventually, the message will arrive at the network serving the web server that you are trying to reach. The network's internal routers will forward your request to the appropriate subnet and then to the destination server. The response to the message, if any, is returned to you in a similar fashion.

The power of the Internet is in its flexibility and resilience. Its network topology is always changing. Hosts are added and removed. New routes are added, and existing routes fail and are removed so that they can be routed around. To keep track of this dynamic network topology, routers periodically exchange their routing tables with immediate neighbors, which update their routing tables. These changes are then exchanged with their neighbors, and so on. Like a rumor, changes in the Internet's topology is soon reflected in all routers worldwide.

Autonomous Systems

Actually, the Internet is more highly structured than simply a set of interconnected networks. The Internet comprises a network of interconnected *autonomous systems* (AS). An AS is a collection of networks controlled by a single (or in some cases, more than one) entity, such as an ISP or a large corporation. More specifically, it is a collection of IP routing prefixes under the control of one or more network operators.

Each AS is assigned a unique autonomous system number (ASN). Currently, the ASN is a 32-bit number. However, this address space is about to be exhausted; and the ASN is being extended to 64 bits.

BGP Routers

Autonomous systems (ASes) are interconnected by routers. Typically today, these routers use the Border Gateway Protocol (BGP) and are called BGP routers.

Routers that share a direct connection (a single hop) are *BGP neighbors*. A message propagates through the Internet by being relayed from one BGP router to its neighbor until it arrives at the destination network. There, the destination network takes over and routes the message to its appropriate subnet and host. Within a network, routers may use internal BGP routers or, more likely, routers using RIP, the Routing Information Protocol.

It is the connectivity between the external BGP routers that describes the relationships of the various ASes and therefore the topology of the Internet.

To maintain a view of the current Internet topology, BGP routers exchange messages with their neighbors advertising new routes and withdrawing unfeasible routes. These changes quickly propagate through the Internet so that all routers have a reasonably current view of the Internet topology.

The BGP Protocol

The BGP protocol[2] is the core routing protocol of the Internet. It provides the mechanism for BGP routers to maintain routing tables that designate network reachability among the ASes.

In the early days of the Internet, there was only one backbone network - the NSFNET, managed by the National Science Foundation initially on behalf of a handful of universities. As the Internet grew, it became clear that more flexibility was needed in the Internet. BGP was introduced to fully decentralize routing to create this flexibility. Its use to interconnect ISP and other networks can be likened to Signaling System 7 (SS7), used for the interprovider core call-setup protocol for the public switched telephone network.

With respect to the eavesdropping problem, we focus our attention on the way in which BGP routers advertise their routes to their neighbors. They do so by periodically sending UPDATE messages to all of their neighbors. The messages can be sent whenever a router experiences a routing table change, or they may be sent periodically – such as every thirty seconds.

The UPDATE message contains route information to a certain subset of one or more IP prefixes (networks plus subnets). In the UPDATE message are two important fields that relate to eavesdropping:

- Network Layer Reachability Information (NLRI): This field contains a set of IP prefixes to which the UPDATE applies. Any of these IP addresses can be reached by any one of the paths specified in the Path Attributes field.

- Path Attributes: This field contains a list of AS paths that can be used to reach the IP addresses listed in the NLRI. Each AS path is a list of ASNs that must be traversed to reach any one of the IP addresses in the NLRI.

An AS path is built by the routers as a new path propagates through the Internet via router advertisements. When a router receives a new route, it adds the ASN of its autonomous system to the path before it advertises it to its neighbors. For instance, if a router in AS 22 receives a new path that is [345, 2078], it will advertise a path of [22, 345, 2078] to its neighbors. This is called *AS path-prepending.* A neighbor then knows that it can route a message to an IP prefix in the

[2] Y. Rekhter, T. Li, S. Hares, A Border Gateway Protocol 4 (BGP-4), *IETF RFC 4271*; January, 2006.

NLRI by sending it to autonomous system 22, from where it will traverse AS 345 and AS 2078 before reaching its destination.

In addition, the UPDATE message may include a list of withdrawn routes to the IP addresses listed in the NLRI.

Thus, a router's routing table may contain several paths advertised by different neighbors to a particular IP prefix. When choosing a route, the BGP protocol chooses the path that routes a prefix most directly to the destination address. For instance, assume that the message is to be delivered to the IP address 161.45.210.10. It finds one entry that will route to IP prefix 161.45 and another that will route to 161.45.210. It will choose the latter route as being the more direct route to the destination.

If a router ends up with more than one route to a specific prefix, it will then choose the route that goes through the least number of ASes. If there are still ties, it uses additional conditions to choose a route.

It is the updating of a router's neighbors plus the construction of routes via AS path-prepending that leads to the eavesdropping vulnerability.

Hijacking

We can now explain the first part of the eavesdropping vulnerability – hijacking. Let us assume that a nefarious AS installs a modified rogue BGP router. The AS is interested in all traffic to a particular subnet with an IP prefix of 161.35.1. Its rogue router is configured to send an update message to its neighbors advertising a route to IP prefix 161.35.1 in its NLRI field and a path to which it has prepended its ASN. This route will be stored in the neighboring routers' route tables. Within a few minutes, the rogue route will propagate throughout the Internet like a bad rumor.

If the advertised prefix of the rogue router is closer to the desired subnet prefix than any other advertised route along the path (say that the closest other route is to prefix 161.35), BGP routers throughout the Internet will select the rogue route as the preferred route. All traffic carried by them to be sent to prefix 161.35.1 will be routed to the rogue router and thence to the nefarious AS. There, it can be analyzed by the AS administrators.

In effect, BGP highjacking represents a denial-of-service attack. The intended destination will cease getting some or all of its traffic. Even worse, this traffic is now available for viewing by unauthorized personnel.

This is a real threat. In fact, it happened to YouTube quite accidentally in February, 2008. Pakistan decided to redirect YouTube traffic to a "black hole" via BGP hijacking because of what it perceived to be a blasphemous video clip.[3] However, a simple mistake by an engineer at Pakistan Telecommunications Authority caused the redirection to be propagated throughout the Internet. YouTube was inaccessible for about two hours.

Eavesdropping

Hijacking simply makes a site or subnet inaccessible. As with the YouTube outage, this can be quickly detected and corrected. In fact, the attacking AS can even be identified by its rogue path.

What is needed for eavesdropping is the capability that, once hijacked, a message can then be sent on its way to its intended destination. In this way, the recipient is unaware that his traffic is being viewed. This cannot be accomplished by simply having the rogue router send the message

[3] Pakistan lifts the ban on You Tube, *BBC News*; February 26, 1998.

to its neighbors for delivery, as they will route the message back to the rogue router per their routing tables.

Kapela and Pilosov solved this problem by adding a second redirection. By manipulating the AS path-prepending for their advertised path (in an undisclosed way), they were able to compromise only selected neighbors so that these neighbors would redirect their traffic. The other neighbors were unaffected.

As a result, once hijacked, traffic could be sent on its way to the rightful recipient via an uncompromised router. Consequently, Kapela and Pilosov were able to mount what is known as a Man-in-the-Middle (MitM) eavesdropping attack. This is what they demonstrated at DEFCON when they intercepted and viewed all traffic destined to the DEFCON network.

Eavesdrop Protection

It has been known for a long time that a MitM attack was theoretically possible,[4] but it had never before been demonstrated.

BGP hijacking and MitM attacks revolve around locating an ISP that is not filtering advertisements. That is, such attacks could be prevented if only authorized peer routers could advertise to each other. Although this is possible, it is rarely done because of the immense manual effort required to establish and maintain the lists of authorized routers. Furthermore, the memory and CPU requirements to do this are beyond the capability of many of today's routers.

BGP itself has no mechanism to validate the authority of an AS to advertise NLRI information nor to ensure the authenticity of the path attributes advertised by an AS.

BGP is the only routing protocol that uses TCP to exchange messages between neighbors. A move is underway to use a TCP option called TCP MD5[5] to correct this problem. Using such an option, each TCP segment carries a signature incorporating information known only to the connection end points. TCP MD5 is now required to be in all BGP routers. However, though the capability is now available, ISPs have been slow to adopt it.

The bottom line – MitM attacks are possible, and you may not even know when you are attacked. Your best (and maybe only) protection is to encrypt all sensitive traffic.

Summary

Internet protocols were developed in the 1970s under the assumption that the network is trustworthy. It is now known that, unfortunately, this assumption is not valid. In July, 2008, another serious vulnerability was discovered when it was demonstrated that traffic could be hijacked by fraudulently modifying DNS servers to provide an erroneous IP address. This proved to be a flaw in the DNS specification and was quickly patched by the DNS vendors.

The hijacking and MitM attacks described above do not exploit a bug or a flaw in BGP. They use BGP in the way that it is supposed to work. BGP assumes that when a router says that it is the fastest way to a destination, it is telling the truth.

The best protection against such hijacking and MitM attacks is end-to-end encryption of all sensitive traffic.

[4] BGP Security Vulnerabilities Analysis, *IETF RFC 4272*; January, 2006.
[5] A. Heffernan, Protection of BGP Sessions via the TCP MD5 Signature Option, *IETF RFC 2385*; August, 1998.