Network Security

Mobin Javed

October 5, 2011

In this class, we mainly had discussion on threat models w.r.t the class reading, BGP security and defenses against TCP connection hijacking attacks.

1 Takeaways from the class reading

- Paper was written about 17 years ago, and the main aim for assigning the reading was that you get to see how well the predictions turned out in time.
- Discussion framework: paper discussed attacks on different protocol layers:
  - Link layer e.g. Ethernet
  - IP
  - TCP/UDP
  - DNS (To be discussed in detail in future lectures)
  - Other application layer attacks exist but were out of scope for this paper
- In security, it is said that if you have got the wrong threat model, it is impossible to build a secure system. Interestingly, the crypto community and the network security community follow two different threat models. Consider that Alice wants to talk to Bob using the Internet as a communication medium.

  Crypto threat model: The cryptographer's mindset is that I'm not going to assume the security of *any* element that lies in the communication path. They have sort of a coarse-grained and simplified view of the attack landscape i.e. the bad guy can control any portion of the Internet.

  Network security threat model: The network security threat model backs off from the crypto model, and says that the attackers can can control some machines e.g. set up a web-server, set up a domain and control DNS for that domain but they can't control the whole Internet for e.g. they cannot control the routing in the Internet. The attacker cannot eavesdrop or modify packets, until the packets get routed through him.

So, why do we need to defend against the network security threat model when it seems that the crypto model is a superset of the threats covered by the network security model? There are a couple of reasons:

- Crypto doesn't solve the whole problem for e.g. it cannot defend against traffic analysis attacks.
- We haven't been able to deploy crypto as a complete solution for e.g. all TCP sessions on the Internet aren't encrypted and issues like key distribution challenge remain.

Let's try to analyze a bunch of protocols under the two threat models, and try to see where they are secure:

Protocol	Crypto threat model	Network security threat model
HTTP	Insecure	Insecure (TCP spoofing)
HTTPS	Secure if no flaws in the CA infrastructure	Secure
SSH	Secure if the user is careful with warnings	Secure
IMAP	Insecure	Secure
IMAP over SSL	Secure	Secure

The network security threat model changes a bit in the scenario of WiFi and if the attacker is on the same LAN, because tapping allows him to eavesdrop and consequently modify all the traffic. The class reading discusses techniques that allow a limited network security class attacker to have a crypto class attacker capabilities.

2 BGP Security

2.1 Primer on BGP routing

The standard protocol for routing in the wide area network is BGP. The Internet is divided into autonomous systems (ASes), and each AS has it's own BGP router. The BGP router knows how to reach all the addresses within it's AS. For e.g., Berkeley's BGP router knows how to get to all addresses within Berkeley.

BGP enables computers across different ASes to reach each other for e.g. it allows Stanford to reach Berkeley and solves the issue of how to get Stanford's BGP router to send packets destined for Berkeley to be sent in Berkeley's direction. It provides a mechanism for Stanford to build a tree, and decide which way to send the packets. BGP provides this capability for the whole Internet, by maintaining this tree in a distributed fashion.

BGP is a flooding algorithm. So, for example, Berkeley announces to it's neighbour that whenever you have a packet for this address range, send it to me. Announcements cause

the routing tables to be updated. Anytime a BGP router updates its routing tables, it sends a corresponding announcement to its neighbors. An announcement has some metric that measures of the quality of the path.

This is not quite how it works, for if it were, the flooding wouldn't ever stop. To make sure that the flooding doesn't keep going on forever there is a field called 'AS Path' in the announcement packet. Each BGP router, before forwarding the announcement, modifies it to include itself in the AS Path field. AS Path is used to silence announcements in the following fashion: If Berkeley sent out an announcement, and the announcement gets back to Berkeley, Berkeley see's itself in the AS Path and simply ignores the announcement.

How are peering relationships taken into account? A router doesn't send the announcement to all of it's neighbours. If an AS X has a policy that it doesn't want to carry traffic from the neighbouring AS Y due to economic and policy reasons, then X simply doesn't announce it's routes to Y.

That's basically how BGP works. Now to analyze the security of BGP, let's look at some real BGP security incidents.

2.2 BGP security incidents

MIT's blackholing: Someone in Florida misconfigured it's BGP routers to announce that it has a new very fast way to get to MIT. The announcements spread and consequently a significant fraction of traffic destined for MIT went to the T1 line in Florida.

Class Discussion:
Q: Is IP hijacking still happening?
We still see some instances of this happening. A recent example is the Pakistan - YouTube incident from 2008.
Q: Who can mount such an attack?

- Anyone who owns a BGP routers.
- Anyone who can compromise a BGP router. Often the BGP routers are professionally administered but sometimes an attacker can find one that hasn't been updated.

Q: How does one own a BGP router?
You could easily buy and set up one but the challenge is to get others to accept your routes because today a lot of ISPs configure their routers to do some sanity checks before forwarding the announcements.
Q: Is BGP over Ethernet/IP?
Long lived TCP connections between gateways.

2.3 Impact of black-holing incidents

How could you make money on the Internet using a BGP black-holing attack? You could advertise that you are Amazon or eBay and capture a lot of credit card numbers. Implications? You can get caught quite quickly. Routers send revocations for revoking the old paths before sending new announcements, and so many revocations going around on the Internet would serve as a big indicator.

So, the conventional wisdom is that black-holing attacks are possible, but anyone doing it would be so blindly obvious that it's possible to catch them.

2.4 Class discussion on defenses on these attacks/ new attacks

- Use SSL; even if BGP has gone bad, crypto at app layer will provide end-to-end security (Doesn't ensure availability).
- Iterative signing of announcements. For example, Amazon signs announcement saying University of Oregon is one hop away from me, Oregon signs Stanford is one hop away, and Stanford signs it's announcement that it is one hop away from Berkeley. Why do we need signatures from all along the path? If we only get signature from last hop, we would be trusting the last hop. The last hop might be evil/ compromised or might have been deceived by routers up in the path.
  Issues: Because routes change, we need a way to tell how long is the signature valid and the ability to re-generate signatures. One of the consequences is that there are a lot of signatures being generated and verified, and this is computationally expensive.
- Intrusion detection:
  - Detect routes going in strange directions.
  - Compare different router's routes to see if they are consistent.
  - Measure if the traffic going in to the router is equal to the amount of traffic going out of the router.

2.5 Stealthy BGP attacks

Conventional thinking was that you cannot do black-holing attacks in a stealthy way until a bunch of people came up with this attack:

Suppose Berkeley wants to steal Stanford's Amazon traffic. The key to launching a stealthy attack is to stop flooding. Berkeley sends out an announcement to Stanford for Amazon with an AS path containing the addresses of all other routers to whom Stanford is going to flood the announcement. These routers would see themselves in the AS path and ignore the announcement. This attack also allows Berkeley to serve as a man-in-the-middle for Stanford's traffic to Amazon, because other routers have a valid route to Amazon through which Berkeley will route Stanford's traffic.

BGP state of the art: Lot of reliance /trust on the people who own the BGP routers.

3 TCP Connection Hijacking attacks

Today, primarily unguessable Initial Sequence Numbers (ISNs) are used to raise the bar for an attacker for spoofing/ hijacking a TCP connection. Bottom-line is that since the ISN is random, chance of success in attack is $1/2^{32}$. If the attacker is on a extremely high speed link, then he might be able to send millions of requests per second, and succeed in hijacking the connection.

So, let's have a discussion on how are unguessable ISNs generated? Two crypto primitives are of use here: i) True random number generators ii) Pseudo-random functions. The idea is to get 128 bits of true randomness, and then expand the 128 bits to as much as you want but these expanded bits are only pseudorandom.

Pseudorandom function is sort of a keyed hash function. Bellovin proposed the use of a pseudorandom function based on the key *(sip, sport, dip, dport)*, and to use the output as an ISN. The key insight of using this tuple as a key is that sending a lot of unspoofed requests for reconnaissance about the ISNs will not help the attacker, because when he tries to initiate a spoofed connection, a totally different key will be used in the pseudo-random function.