10/03/12 - DNS Security

Richard Xia

October 3, 2012

1 Ethics of Publishing Exploits

1.1 Bellovin paper

- Not clear at what point did Bellovin decide to not publish it
- Wrote it in 1990, eventually published it in 95
- May have circulated it to a small group of people
- Leaked anyway

1.2 Kaminsky Attack

- Critics didn't believe him and thought he just wanted publicity
- Picked a small set of respected security experts along with critics and explained the attack
- People now believed Kaminsky
- One guy accidentally published the attack on his blog

1.3 Limited disclosure

- Hard to disclose exploits to a limited set of people
- One vendor pushed an update with a way to detect an attack and people reverse engineered it to discover the attack
- Hard to get people to believe that the attack is possible

1.4 Example: Voting machines

- State of California asked security researches to audit the security of the voting machines
- Turned out there was the possibility of spreading a virus across many voting machines
- Told Secretary of State office to immediately pressure voting machine vendors to make a fix and get it certified
- Harry Hursti wanted to publish the exploit, David Wagner and group ask him to not publish enough information for people to write an exploit

- Vendor did not make a fix
- Group at Princeton got a hold of a voting machine and secretly worked on it. The week before the election they published the same exploit.
- The efforts to keep the exploit unknown were subverted

1.5 Summary

Two camps: Full disclosure people believe that the only way to create pressure to fix an exploit is to publish it. The responsible disclosure people believe that it is irresponsible to publish an exploit before letting people work on fixing them.

2 Kaminsky Attack

2.1 DNS Overview

If you make a DNS query for amazon.com, what you get back may not be the actual mapping, but instead "Here is the name server that knows where it is", such as ns.amazon.com (the name server for Amazon) and provides the IP address for it. The problem with DNS is how do you get the IP address of a name server without knowing the address of a domain server and vice versa? The protocol forces you to trust that the address that you get back for the nameserver is correct.

2.2 Attack #1

I force Alice to make a request to google.com, and before Google responds, I spoof a response which tells Alice that the name server for google.com is at an address that I own. Spoofing the IP address and port was easy, at least before Kaminsky, forcing the target to make a DNS request is easy, so all I need to do is to guess the QID. The trick is to guess correctly and send the spoofed response before Alice get the real answer or else Alice will cache the real answer.

2.3 Kaminsky Attack

Instead of making Alice query google.com, you make her query some random subdomain of google.com which does not exist. The attack is to supply the additional information and make Alice think that the nameserver for Google points to one of our own addresses. Now you can make more attacks against Alice and now has a much higher chance of winning on at least one of the attacks. More severe attack is to attack the DNS server that Alice uses.

3 DNS Sec

- Hierarchical authentication
- root signs .com which signs amazon.com which signs www.amazon.com
- Need a way to be able to sign negative responses (e.g. asdf.amazon.com is not a real subdomain)

4 Transitive Trust

If berkeley.edu lists ns.uoregon.edu as a possible name server, then an attack on uoregon could affect berkeley.edu. Even if there are multiple name servers listed, you could try to DoS the other name servers to force them to use yours. Transitive trust: Berkeley trusts UOregon, UOregon trusts Arizona, therefore Berkeley trusts Arizona. The average DNS server depends on 46 other servers. 45% of DNS servers depend on at least one server with a known vulnerability.