Paxson
Spring 2017

CS 161
Computer Security

Discussion 8

Week of March 20, 2017

Question 1 ***IP Spoofing*** **(15 min)**

You are the network administrator for a large company.

(a) Your company will be held liable for any spoofing attacks that originate from within your network and are sent out to the global Internet. What can you do to prevent spoofing attacks by your own employees?

(b) You now want to evaluate the risk your employees face from spoofed IP packets originating from outside the network. Assess the likelihood and dangers of spoofed IP packets that use TCP as the transport layer protocol. What applications might be vulnerable to such attacks? How does this change with UDP?

(c) What can be done to prevent parties outside your network from sending your employees spoofed traffic that impersonates your own employees?

Question 2 ***DNS*** **(15 min)**

Recall that in a *blind* DNS spoofing attack, the attacker tries to guess the identification number of the DNS request sent by the victim.

(a) For the following, assume that the victim's DNS resolver cache does not contain a record for the domain the attacker is targeting:

i. In a blind spoofing attack, the attacker must know when a DNS request is about to be made by the victim so that the attacker can respond with their attack responses. Recall from lecture how an attacker might learn this information.

ii. What can an attacker do if they successfully get a victim to believe their bogus DNS mapping?

iii. How can an attacker avoid having to carry out this attack for every time the victim visits the targeted domain?

(b) Now assume that the victim's DNS cache has a genuine `NS` record for the domain the attacker is targeting.

i. Can the attacker still be successful at poisoning the `A` records for some of the names belonging to the domain?

ii. Can the attacker poison the `NS` record of this domain? If yes, how?

Question 3 ***Sniffer detection*** **(10 min)**

As the security officer for your company, your network monitoring has observed a download of a "sniffer" tool. This tool passively eavesdrops on traffic, and whenever it sees

a web session going to a server in a `*.yahoo.com` domain, it extracts the user's session cookie. It then uses the cookie to create a new connection that automatically logs in as the user and exfiltrates all of their `*.yahoo.com` activity, such as their emails if they use a `yahoo.com` email account.

You become concerned that one of your employees may have installed a network "tap" somewhere among the hundreds of links inside your building, and will use it to run this tool. How might you determine whether such a sniffer is in operation?