# Homework 3: Caught in the Net

*Due: Monday, April 21 @ 11:59 pm EDT*

## Overview and instructions

This homework has only two problems. There is no extra CS1620/CS2660 component.

### Note on collaboration

You are welcome (and encouraged!) to collaborate with your peers, but the solutions you write down must be **your own work** (ie, written by you). You are responsible for independently understanding all work that you submit—after discussing a problem as a group, you should ensure that you are able to produce your own answers independently to ensure that you understand the problem. For more information, please see the course Collaboration Policy.

In your submission, we ask that you include a brief *collaboration statement* describing how you collaborated with others on each problem—see the next section for details.

### How to submit

You will submit your work in PDF form on Gradescope. Your PDF should conform to the following requirements:

- **Do not** include any identifying information (name, CS username, Banner ID, etc.) in your PDF, since all homeworks are graded anonymously
- Each problem (where "problem" is one of the Problems 1–2) should start on a separate page. When you submit on Gradescope, you will be asked to mark which pages correspond to which problem
- At the start of each problem, write a brief *collaboration statement* that lists the names and CS usernames of anyone you collaborated with and what ideas you discussed together
- If you consulted any outside resources while answering any question, you should cite them with your answer

## Problem 1: Local network eavesdropping

Relevant lectures: Lectures 18–19

Consider the network represented in Figure 1, a subnet whose addresses all take the form `192.168.1.*` (ie, the subnet 192.168.1.0/24.) Each router and host is labeled with its IP address and MAC address. In all parts of this problem, assume that all hosts (Host A, Host B, and Host C) and the router have entries for all other hosts on the subnet in their respective ARP tables, and thus no device in Figure 1 is actively sending any ARP messages.

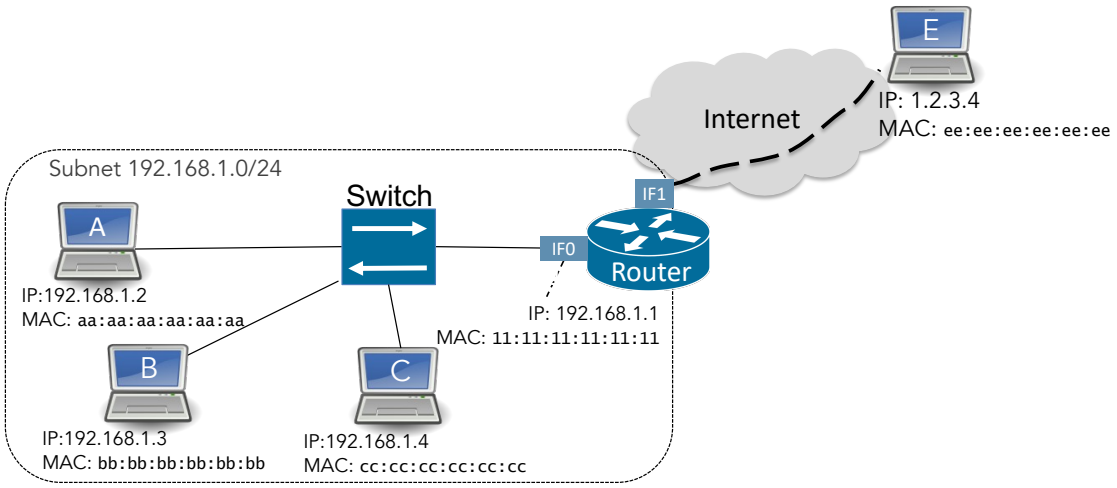


Figure 1: An example network.

**Question a)** (4 pts) Host B wants to intercept traffic between Host A and Host C. How could B use ARP to cause A to send its traffic to B instead of C? In 1–2 sentences, explain what ARP message(s) you would send, what they would contain, and to which hosts you would send them.
**Note**: In carrying out this attack, B needs to be careful not to accidentally break other hosts' connections. How can B make sure that their attack only targets A?

**Question b)** (4 pts) Host B's attack is successful and it's now intercepting traffic sent by Host A to Host C, but this means that Host C isn't getting the traffic. As a result, A may notice that something is going wrong and stop sending data, which won't give you the information you need. How could B make sure that the communication between A and C is retained, while B can intercept it?

**Question c)** (4 pts) Assume that Host B's attack was successful and it is now intercepting Host A's traffic, and Host C is now receiving the traffic and responding as normal—so A is not aware of the attack. However, B would also like to intercept **Host C's responses to Host A**. How can B accomplish this? Similar to part (a), explain your attack and why it works.

**Question d)** (4 pts) How would these techniques differ if Host B wanted to intercept Host A's communication with Host E, which is somewhere on the internet (ie, not on this subnet). Once again, we don't want to break the communication between Host A and Host E (ie, similar to part (b). In your response, try to be precise about the content of any attack packets you would send.
**Hint:** Since Host E is not on B's subnet, it will not suffice to spoof E'S MAC address.

## Problem 2: Thinking about DNS

Relevant lectures: Lecture 20

**Question a)** In the DNS lecture, we discussed how each DNS query has a `request identifier`, also called a transaction ID, to identify responses to individual queries. Consider the following questions about how DNS transaction IDs are used (your answers should be no more than 100 words for each):

(i) Why is using a transaction ID more secure than using *no* transaction ID?

(ii) Why is using *randomized* transaction IDs is more secure than having *sequential* transaction IDs.

**Question b)** Imagine you control the default DNS server for an Internet Service Provider (ISP). In 2–3 sentences, briefly explain how you can attempt to leverage DNS to block access to sites you don't want your customers to access.