

Homework 4: The Final Homework

Due: Friday, April 28 @ 11:59 pm ET

Overview and instructions

This homework has 4 problems:

- Problems 1–4 are required for all students
- There is no extra CS1620/CS2660 part for this assignment. Have fun with Dropbox!

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 1–4) 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

There is only one Gradescope submission for this assignment. All students should submit Problems 1–4 to the assignment labeled “**Homework 4**” on Gradescope.

Problem 1: Thinking about DNS

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.

Problem 2: TLS vs. Blue University

Blue University worked out a really good deal with a the CA AwesomeTrust to get a TLS certificate for its main website, `blue.university`. However, based on what you know about Blue University and the reputation of their business partners, you know something is going to go very wrong here...

Consider the following scenarios, which examine what happens when different parts of a public key infrastructure are affected if a certain key is compromised. For each part, write a brief explanation (1-2 sentences) of what capabilities you (as an attacker) can do in each scenario. Specifically, for each question, consider:

- Who could you impersonate?
- Who would believe you? (ie, For what set of users would you be able to impersonate X? All users in the world, or just a specific group?)

Note: For considering who you can impersonate, you don't need to worry about *how* you'd actually perform the attack at the network layer—instead, just **focus on how the TLS authentication process would be affected if you control the key in question.**

Question a) You compromise a University webserver and obtain the private key for `blue.university`. If you have this private key, what capabilities do you have? Give your answer by considering the two bullet points above and briefly explain your reasoning.

Question b) It turns out that the CA AwesomeTrust is pretty shady (big surprise...) and you're able to obtain AwesomeTrust's CA private key. If you have AwesomeTrust's CA private key, what capabilities do you have? Give your answer by considering the two bullet points above and briefly explain your reasoning.

Question c) After you compromised AwesomeTrust, Blue University gives up on the idea of delegating trust to a third-party: it decides to *make its own CA* and issue its own certificate for `blue.university`. To make this work, the University's IT policy requires all users to install the Blue University CA certificate on their systems.

- (i) Why do Blue University users need to install the CA certificate? What happens if a browser **doesn't** have the CA certificate installed and connects to `blue.university` anyway?
- (ii) If you're able to obtain the Blue University CA private key, what capabilities do you have? Give your answer by considering the two bullet points above and briefly explain your reasoning.

Problem 3: TryHackMe Lab: Nmap

Go to <https://tryhackme.com/jr/nmap02uv> to complete a lab using `nmap`, a port scanning tool. This lab will expand on what we learned in lecture about IPs and ports and show you how to scan for open ports and services using `nmap`.

As a reminder, these TryHackMe rooms are graded based on *completion*, not correctness. As long as you have answered all of the questions on TryHackMe, you will get full credit—you do not need to submit anything in the PDF you upload to Gradescope. This assignment should not take more than 1 hour, so if you are stuck or are dealing with technical issues, please post a question on Edstem.

Problem 4: Onion-Flavored Handins

Note: We will have covered all of the material for this problem after the lecture on Anonymous Networking on Tuesday, April 25.

The CS666 course at Blue university has its own Tor network to allow students to submit their code “super-anonymously” to a custom testing server. Unlike Gradescope, the testing server doesn’t require students log in and instead just sends back a test report so students can check their work against the autograder. This way, only students’ final, best submissions need to be uploaded to Gradescope and recorded with their names.¹

Figure 1 depicts the CS666 Tor network and a specific circuit established between you and the testing server. Lines depict the specific Tor circuit established by your client using the nodes highlighted in blue. In this network, each Tor node i has a public key PK_i , which is known to all other Tor nodes and the client².

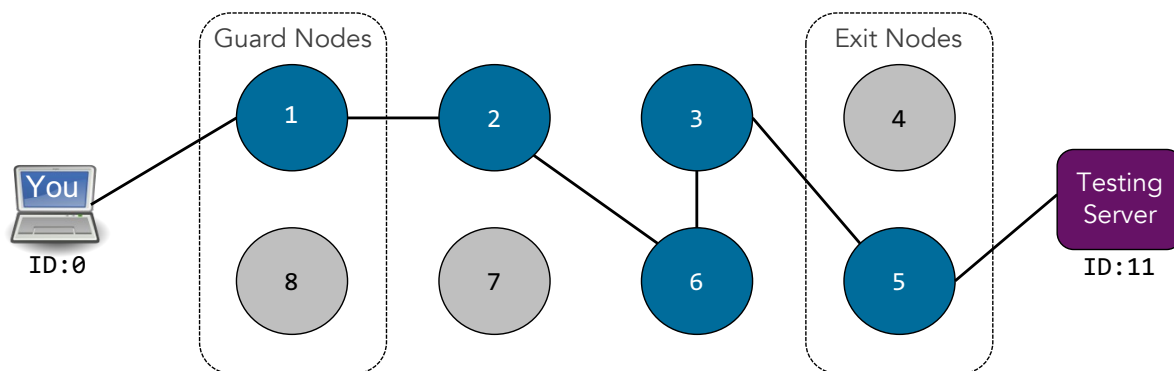


Figure 1: The CS666 Tor network. Note that the lines indicate connections in the Tor circuit, not individual network links.

Question a) To send a message M to the testing server using the circuit in Figure 1, what is the packet your Tor browser sends out?

You can express your answer in the format: `[dest-id, data]`

where `dest-id` is the ID of a node as labeled in Figure 1 and `data` is the packet content. For example, `[1, "Hello"]` sends the string “Hello” (in plaintext) to node 1. You can denote encryption of message M as $E(PK_i, M)$, where PK_i is the public key for a node with ID i . You can assume that each key is known to all Tor nodes and the client (this is established during circuit setup).

Question b) **True or False:** even if you connect to the testing server using plain HTTP (**not** HTTPS), no one can read M except the testing server—that is, TLS isn’t required in order to have *confidentiality* in the presence of an eavesdropper in the network because the traffic is encrypted. **Explain your reasoning.**

(Continued on the next page)

¹I promise you that we don’t actually care about your non-final submissions and don’t look unless we think there’s an autograder bug. It makes a nice case study for Tor, though. ☺

²This is a simplification of the actual Tor protocol, which uses a combination of asymmetric and symmetric keys to set up circuits.

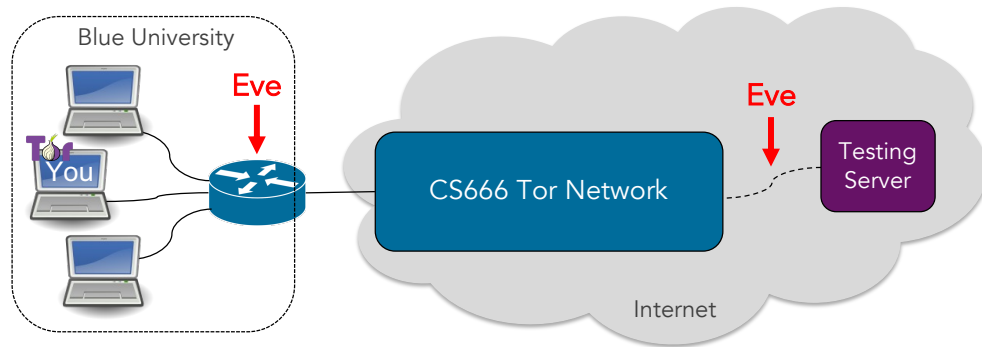


Figure 2: Eve's eavesdropping capabilities.

Question c) Suppose that Eve gains access to a bunch of Blue University systems and figures out how to view network traffic at the two points shown in Figure 2. Specifically, Eve can observe individual users' traffic leaving the Blue University network *and* traffic entering the testing server. Even if you use HTTPS, what, if anything, could Eve learn about your submissions to the testing server?

Explain your reasoning in 1–2 sentences—you don't need to provide specific details of an attack, just sketch what sort of information Eve could learn, and what is still hidden.