CS155: Computer Security **Spring 2023**

Final Exam

Instructions:

- **Please answer all six questions. You have two and a half hours.**
- The exam is open book, open notes, and open laptops. However, you are expected to do the exam on your own. You may not interact, collaborate, or discuss the exam with another person during the exam window. **You may not connect to the Internet during the exam window** and may not use an AI tool.
- If you are taking the exam on campus, please write your answers on the exam sheet. We will scan the exam sheets into gradescope, so please do not write on the back side of the exam sheets. If you are taking the exam remotely please do the same, and upload a picture of your exam sheets to gradescope (course code `E7NGZK`).
- Students are bound by the Stanford honor code. In particular, you are expected to do the exam on your own and without access to the Internet.

I acknowledge and accept the Honor Code.

_____

(Signature)

_____ _____

(SUNet ID) *(Print your name,* legibly!*)*

Graduating this quarter? ☐ (mark X if yes)

1. (*10 points*) ..... True or False

For each question, please write `T` or `F` in the space provided. No explanation needed.

_____ (a) The hash function used in a password authentication system needs to be slow to prevent dictionary attacks.

_____ (b) The Server Name Indication (SNI) extension used in HTTPS is needed when the server has no certificate or secret key.

_____ (c) A use-after-free vulnerability cannot happen in a language that uses garbage collection, such as Java.

_____ (d) The Log4j vulnerability was discovered by a software engineer in China.

_____ (e) Alice visits a web site on Monday and the site sets a cookie named `sessionID` in Alice's browser with an expiry date one year in the future. Alice revisits the same site on Tuesday. On the second visit the site cannot delete the cookie `sessionID` from Alice's browser.

_____ (f) `https://cs.stanford.edu/crypto/index.html` and `http://cs.stanford.edu/data.html` are in the same origin.

_____ (g) One of the goals of Intel SGX is to provide a secure hardware enclave so that even if the host operating system is controlled by an attacker, the attacker cannot read data in the enclave's memory.

_____ (h) It is important for a website served over HTTPS to use DNSSEC for name resolution.

_____ (i) The HTTP `Referer` header provides the domain name of the linking website, but not the URL path or query parameters.

_____ (j) Because the TLS protocol is used to encrypt communication between Tor nodes, users can safely visit HTTP-only (i.e., no HTTPS) websites over Tor.

2. (*20 points*) ..... Questions from all over with a short answer

(a) (*4 points*) What is a race condition and why is it a security problem?

Your answer:

(b) (*4 points*) Suppose you find a memory leak in a third-party library, say the library's `cleanup` function forgets to call free on some pointer `p`. You notice that some applications fix the problem by freeing the memory themselves after calling `cleanup`. How would you fix the library's `cleanup` function to free `p` without introducing a double-free vulnerability in applications that free the memory themselves? You are only allowed to change the library. You may assume that calling `free(null)` does nothing.

Your answer:

(c) (*4 points*) The same origin policy (SOP) for DOM access is based on the triple (protocol, host, port). Suppose SOP did not include protocol (i.e. SOP was defined using only host and port). What goes wrong? For example, explain how a network attacker could steal a gmail secure cookie (i.e. a cookie sent only over HTTPS). Note that reading document.cookie in an HTTP context does not reveal secure cookies.

Your answer:

(d) (*4 points*) Explain why only UDP-based protocols are used for amplifying DDoS attacks.

Your answer:

(e) (*4 points*) Many companies are moving away from having internal privileged applications secured by a single organizational VPN. Explain (1) why organizations are moving away from this model, (2) what security model organizations are adopting, and (3) how it protects against (1).

Your answer:

3. (*15 points*) ..... HTTPS

(a) (*3 points*) Suppose an HTTPS page includes an HTTP iframe where both are loaded from the same domain and port. Clearly, the iframe is vulnerable to being corrupted by a network attacker. Is the outer frame (the one loaded over HTTPS) vulnerable to being corrupted by a network attacker?
Hint: are the inner iframe and the outer frame in the same origin?

Your answer:

(b) (*3 points*) Consider a country XYZ and suppose that all Internet traffic to and from XYZ must pass through a single government gateway. Citizens of XYZ that visit the `nytimes.com` site, which resides outside of XYZ, are protected from snooping because all of the NY Times content is served over HTTPS. Suppose the NY Times certificate is issued by a certificate authority CA1. The intelligence services of XYZ are able to bribe or compromise some certificate authority CA2 and obtain a certificate for `nytimes.com` signed by CA2. Can they now snoop on all the traffic from inside XYZ that is headed to `nytimes.com`? If so explain how, if not explain why not.

You may assume that CA2 is a certificate authority that is trusted by all browsers in the world. You may also assume that the NY Times does not use certificate pinning, and does not check the certificate transparency logs.

Your answer:

(c) (*3 points*) Continuing with part (b), which of the security mechanisms we discussed in class would help detect that XYZ's intelligence services were able to obtain a rogue certificate for `nytimes.com`?

Your answer:

(d) (*6 points*) A corporate admin wants to monitor all web traffic originating from inside the corporate network and headed towards an external web site (e.g., to `nytimes.com`). The monitoring needs to be done transparently, i.e. without any change to the user experience. The admin can route all web traffic through a transparent web proxy and the web proxy will monitor the traffic. The problem is that for sites that use HTTPS, the proxy sees encrypted traffic and cannot do its job.

Propose an architecture that will enable the proxy to monitor HTTPS traffic in cleartext. Again, this should be done without affecting the user experience, i.e. without causing any security dialogs to popup on the employee's browser.

Hint: You may assume that the employee's computer is installed by the corporate admin and that the admin installed a corporate CA certificate in the browser's certificate store. The proxy has the secret key for this corporate CA. You may also assume the browser sends the SNI client hello extension in the clear. Be precise in your explanation of how the proxy works. In particular, what does the proxy need to do the first time an employee on the corporate network visits a website such as `nytimes.com`?

Your answer:

4. (*20 points*) ..... Isolation and sandboxing

(a) (*3 points*) In lecture 5 we discussed how the freebsd jail program can be used to run an application in a jail. Suppose we run a PDF viewer in a jail. Before the jail is created, we create a symbolic link `temp.pdf` in the jail's top level directory that points to a PDF file `doc.pdf` that resides outside of the jail. We then start the jail and run the PDF viewer in the jail, asking it to display `temp.pdf`. Will the PDF viewer display the document? Justify your answer.

Your answer:

(b) (*3 points*) Another important isolation mechanism is called `seccomp-bpf`. Briefly explain how a `seccomp-bpf` policy can ensure that a compromised application cannot cause long-term damage to a system. For example, suppose we want to prevent a PDF viewer from accessing the network. How does `seccomp-bpf` enforce that?

Your answer:

(c) (*4 points*) Suppose we want to allow the PDF viewer to connect to a single remote site that provides font files. The PDF viewer should be unable to connect to any other remote site. Can one write a `seccomp-bpf` policy file to enforce this policy?

Your answer:

(d) (*4 points*) In our discussion of Docker confinement flags, we mentioned that one can start a docker container with the flag `--restart=on-failure:20`. This flag means that if the program running in the container crashes, the container will be automatically restarted. However, after the 20th crash, the container will no longer restart automatically. What attack is this meant to prevent?
Hint: think back to our discussion of control hijacking defenses.

Your answer:

(e) (*6 points*) In Lecture 5 we also discussed how security companies often detect web sites that serve malware: the security company crawls the web using a vulnerable browser running inside of a hypervisor. After visiting a web site, the security company checks if the browser has been damaged, and if so, it declares the web site as a malware site. Can you think of two ways in which a malware site can evade such a scan? That is, the site will infect regular visitors, but will not be detected by the scanner.

i. Method 1:

Your answer:

ii. Method 2:

Your answer:

5. (20 points) ..... Privacy and Tracking

(a) (3 points) Despite the near-ubiquitious use of HTTPS on the web today, some users continue to use Tor to browse the web. Provide one privacy guarantee that Tor provides against an on-path passive-eavesdropper compared to normal HTTPS. **Hint:** What data is leaked by HTTPS?

Your answer:

(b) (3 points) Explain what protection Tor provides that is stronger than what is provided by a standard VPN service.

Your answer:

(c) (4 points) Tor places extra emphasis on finding entry nodes (also known as guard nodes) that they know are operated by trusted parties. Explain how an adversary might be able to deanonymize users if they control the *entry node* and *exit node* for a user's session on a website but none of the nodes in between. How can Tor protect against this attack?

Your answer:

(d) (*4 points*) Tracking pixels are 1×1 pixel images that advertisers and other trackers include in order to generate a simple request to a third-party website. Imagine there are two independent advertisers (with different domains) that use tracking pixels, but cannot execute any Javascript. They want to synchronize the user identifiers with one another to broaden their tracking information. Devise a mechanism for how two providers can synchronize their unique identifiers for a user without using Javascript.
Hint: HTTP Redirects enable a web server to redirect the web browser to a different URL when loaded.

Your answer:

(e) (*3 points*)
Several browser extensions like Ghostery exist for blocking ads and other tracking mechanisms. They operate by blacklisting domains that belong to tracking sites, and blocking every request from the browser to a blacklisted domain. Explain why these extensions sometimes break websites.

Your answer:

(f) (*3 points*)
Many websites now check for ad blocking extensions in browsers. Explain how a website XYZ can determine whether ads from `doubleclick.net` are being blocked by the browser, despite the fact that the same origin policy prevents XYZ from viewing content loaded from Doubleclick.

Your answer:

6. (*20 points*) ..... Network and Internet Security

(a) (*3 points*) The Mirai Botnet compromised hundreds of thousands of IoT devices by attempting default credentials against Telnet and SSH services. The malware then used compromised devices to DDoS popular websites. Why couldn't victim websites simply drop traffic from the compromised devices?

Your answer:

(b) (*4 points*) Imagine that you are an administrator that controls a large network with thousands of devices that could potentially be compromised with the Mirai malware. How could you prevent devices from being compromised, even if you can't patch any of the devices? How could you detect devices that were infected?

Your answer:

(c) (*3 points*) Nearly all of the devices infected by Mirai were not behind Network Address Translation (NAT) gateways. Why weren't devices behind NATs infected?

Your answer:

(d) (*3 points*) Let's suppose that a network prefix (171.67.22.0/23) has been BGP hijacked by an imposter. What can you do to ensure that traffic is sent to you instead of the imposter?

Your answer:

(e) (*3 points*) Does TLS protect users from their traffic being intercepted during a BGP hijacking attack? Why or why not? If not, how can an attacker convince the user's browser that the connection is legitimate? (Assume no HTTPS key-pinning.)

Your answer:

(f) (*4 points*) Most home routers have an HTTP-based internal administration interface. Why is it important that you change the default password even though the interface isn't accessible on the Internet? What attacks can an attacker accomplish if they can change the configuration of your home router?

Your answer: