CS5430 – System Security Spring 2018

Assignment 0

Due: February 12, 2018

Problem 1: Security Policies

Consider the following web-based email system. Users login by visiting a pre-specified URL for the system and then entering both an identifier (i.e., a name) and a password. This starts a session that is associated with the specified identity. The system then displays in a preview frame a list of messages that have been sent to that identity and have not been deleted during this or some prior session associated with that identity. Here, for each message, the name of the sender and the contents of the message are displayed.

During a session, a user can:

a) Click on an icon to generate a reply to the message the user is currently viewing. The user then types the body of the reply. That reply later becomes a message that will be available for viewing by the sender of the original message to which this serves as a reply.

b) Click on an icon to generate a new message. The user then enters an identity of some receiver and enters a body for the message. That body is incorporated into a message that will be available for later viewing by the intended receiver.

c) Click on an icon to delete the message that the user is currently viewing.

d) Click on an icon to end the session.

However, if 15 minutes elapses during which no action is taken by the user, then the system automatically terminates the session.

Threat model: The adversary is a user who desires to read email, generate bogus email, and/or alter email that has been generated by bona fide users. The adversary has access to the URL for the mail system and also can read, delete, and/or update network packets in transit. The adversary cannot physically access or run programs on a user's machine that is running a browser to access the mail system. And the adversary cannot physically access or run programs on the mail system server.

To Do: Identify a list of security policies for this system. Try to be comprehensive, but don't go over one page. For each security policy, label with one of: confidentiality, integrity, or availability.

Problem 2: Security Principles

The `su` command enables a UNIX user $u1$ to access the account of another user $u2$. Unless $u1$ is the superuser ("root"), `su` prompts $u1$ to enter the password of $u2$. Checking whether that password is correct requires `su` to open the password file, `/etc/passwd`. On a correctly configured UNIX system, that particular open operation will always succeed. Then su can proceed with checking whether the password is correct.

A CS 5430 student becomes concerned with what might happen if the UNIX system is not configured correctly —in particular, what if a misconfiguration caused the open operation to fail, and what if that led to the system becoming unusable?

So the student decides to build a new version of `su` that works as follows. If the open operation succeeds, then the password is checked. If it is indeed the correct password for $u2$, then $u1$ is granted access to the account of $u2$. But if the open operation fails, then $u1$ immediately is granted access to the account of the superuser ("root"). The student's intention is that $u1$ would then be able to fix the misconfiguration.

Discuss which of the following security principles the student's new version of `su` upholds, which principles it violates, and which are simply irrelevant:

- Accountability
- Complete Mediation
- Least Privilege
- Failsafe Defaults
- Separation of Privilege
- Defense in Depth
- Economy of Mechanism

You will be evaluated in part on how well you demonstrate understanding of each of the principles. You might find it helpful to review the discussions of them in [Schneider, chapter 1] and [Saltzer and Schroeder 1975].

Problem 3: Real-World Policies

Cornell's Policy Regarding Abuse of Computers and Network Systems (`https://it.cornell.edu/policy/policy-50-abuse-computers-and-network-systems`) is, in part, a security policy that stipulates appropriate usage of computer systems at Cornell.

As a student studying computer security, you obviously need to know your responsibilities with respect to that policy. And as a security expert, you might some day

be asked to write such a security policy or to evaluate somebody's actions relative to a policy. So study the policy, then consider the following problem.

Suppose that a CS 5430 student discovers a vulnerability that can be exploited to bypass the usual NetID authentication used to login to Cornell systems. Such exploitation would enable an attacker to login under any NetID of their choice, thus impersonating any Cornellian. The attack would yield access to all Cornell email, student grades, and student financial statements.

Discuss whether each of the following behaviors is permitted by the Cornell Policy linked above:

a) The student programs a tool that accomplishes the attack. The student uses the tool, but only to read files they are already allowed to access with their NetID.

b) The student programs a tool that accomplishes the attack. The student doesn't actually use the tool but posts it to a well-known website, along with instructions for use of the tool.

c) The student does not program an attack tool but does post a discussion of how the attack would work to the "Overheard at Cornell" Facebook page (i.e., a well-known public website). The discussion contains sufficient technical details to enable a CS major to program an attack tool.

Explain your reasoning. You will be evaluated in part on how well-supported your arguments are. It is to your advantage to quote specific excerpts from the policy that support your arguments.