

Homework 4

CS161 Computer Security, Fall 2008
Assigned 11/12/08
Due 11/19/08

For your solutions you should submit a hard copy; either hand written pages stapled together or a print out of a typeset document¹.

1. *Attestation using Trusted Computing Hardware* [2 points]

Recall that the trusted computing architecture includes a hardware TPM in every machine. The TPM contains a secret signing key and the resulting signatures are used for attestation. Do all computer hardware TPM chips require to have a distinct signing key? If yes, why is it secure. If not, describe a possible attack on the attestation process.

2. *Trusted Boot using Trusted Computing Hardware* [3 points]

Recall that the trusted computing architecture includes a hardware TPM in every machine. Suppose the TPM does not have any persistent *protected* storage other than those needed for storing the cryptographic keys. This would still permit remote attestation applications, as the attesting process requires no protected persistent storage. Would it also permit secure boot? If yes, how would you design a secure system. If no, explain why.

3. *Inline Reference Monitors* [10 points]

Recall that *address sandboxing* is the policy of confining all memory reads and writes made by an untrusted module to a pre-determined segment in memory. It is possible to enforce address sandboxing using software fault isolation (SFI) – SFI embeds an inline reference monitor to confine all memory accesses to within one segment in the process address space.

Consider a target architecture has only two memory access instructions, `load [r1], r2` reads the value from memory pointed to by register `r1` into register `r2`, and, `store [r1], r2` writes the value of register `r2` to memory pointed to by register `r1`. The target architecture has only one unconditional control transfer instruction, `jmp [r1]`, that transfers control to the address stored in `r1`. All instruction on the architecture are 8 bytes in length. You may assume other instructions exist if you need them to solve this question, but be sure to define them before using them in your answers.

(a) [3 points] The following checks for the enforcing address sandboxing policy are inserted before each `load [r1], r2` and `store [r1], r2` instruction.

```
r1 = r1 & 0x00000fff
```

```
r1 = r1 | 0x00000000
```

The following checks are inserted before `jmp [r1]` instructions :

```
r1 = r1 & 0x00000fff
```

```
r1 = r1 | 0xbbbbb000
```

This implementation of the address sandboxing scheme aims to confine memory reads and writes to the memory range `0x00000000 - 0xbbbbbfff`, and the untrusted code with inline checks is placed in the memory segment `0xbbbbb000 - 0xbbbbbfff`. Explain why this implementation still allows a malicious code to write memory outside the intended range. Outline an alternate implementation that is secure, and explain why your proposed implementation is safe.

¹LaTeX is the most suitable tool for typesetting mathematical documents, but other use of other editors are perfectly acceptable

(b) In parts (b.1) and (b.2) you are asked to design an inline reference monitor scheme, similar to SFI, with a different policy than address sandboxing. The new policy is to prevent an untrusted module from reading/writing data within one segment, starting at address `0xabcde000` and ending at `0xabcdefff`.

(b.1) [3 points] Show the inline reference monitor checks, if any, that may be needed before the `load`, `store` and `jmp` instructions in the untrusted code. Clearly state use of any dedicated registers and/or memory in your scheme.

Extra Credit [2 points]. Fine tune your scheme to have at most 2 instructions, that use only bitwise AND, OR, NOT, XOR operations.

(b.2) [4 points] How does your scheme prevent a malicious module from bypassing the inserted checks using indirect control transfer instructions in its code? Briefly explain why are the checks sufficient to ensure the new policy.

4. System Call Interposition [10 points]

Consider the following program running on a typical UNIX system.

```
1. S0;
2. while (..) {
3.     S1;
4.     if (...) S2;
5.     else S3;
6.     if (... S4 ...) ... ;
7.     else S2;
8.     S5;
9. }
10. S6;
11. S7;
```

For conciseness, the code only show the system calls made by the program as $S_1, S_2 \dots, S_7$. S_0 denotes a network write system, S_1 denotes a network read, S_2 is a file read system call, S_3 is a file write, S_4 is a file open system call, S_5 is file close, S_6 is a `getruid` and S_7 is `setuid`.

(a) [2 points] How can system call interposition prevent an attacker from exploiting this program to execute a shell (via the `exec` system call).

(b) [6 points] Based on the program code, you notice that the program can only execute a certain sequence of system calls under benign operation. For instance, the program never performs a network read or write after the `setuid` system call. Write a regular expression that captures all legitimate sequences of system calls that the program shown above can make. Express your answer using regular expression operators (concatenation, alternation and Kleene's star) over the set of symbols ($S_1, S_2 \dots, S_7$ and ϵ).

(c) [2 points] How does your regular expression model in part (b) help enforce the following policy using system call interposition : *No network writes after file read*. Outline one example of a threat that this policy defends against.

5. Privilege Separation [10 points]

Foomail is a UNIX electronic mail transfer program – it manages sending and receiving mails on a multi-user UNIX system. In this question, you are given a high-level design description of the functions that *foomail* performs. Note that it requires superuser privileges, but many of its operations handle immense amount of untrusted data. Your task is to redesign the *foomail* program with several trusted and untrusted components using the principle of least privilege.

System Resources. The set of resources that are used by *foomail* along with a brief description is outlined below.

- *Mail Queue (Q)* : An operating system resource that stores incoming and outgoing mails, waiting to be processed, for all users on the system.
- *Mail Spool Directory (SD)* : A special directory containing a separate file for each user where the mail transfer program stores the received mail for that user.
- *SMTP Port on local machine (P)* : Port 25 is the SMTP port on which the mail transfer program binds and listens for incoming mail delivery connections.
- *Other high-numbered ports on local ports (N)* : `foomail` needs to bind and connect to high-numbered ports on local machine to access the network for communication with SMTP mail servers on remote hosts.
- *DNS (D)* : Access to a DNS server on a remote host on the network, which resolves network host names to IP addresses.
- */etc/hosts file* : Privileged file (read/write access allowed only to superuser, read-only access to users on the system), that is a local database to lookup IP addresses for network hostnames.
- */etc/passwd file* : Privileged file (read/write access allowed only to superuser, read-only access to users on the system), containing general user information (in addition to other sensitive information) about each user on the system.
- *.forward file* : A user-specific file (owned by a user), that contains information about how to forward email received at one email address to another.

Access to Q, SD, and binding to port P requires superuser permissions. Access permissions to files are as outlined in the above description; other resources can be accessed by processes running as unprivileged (non-superuser) users.

Operations performed by foomail. When the `foomail` process runs, it roughly needs to perform the following tasks, not necessarily in the order defined below:

- Bind and listen on local SMTP port P for receiving mail.
- Upon receiving mail from the network port P, it updates Q with the arrived mail.
- Read Q to fetch incoming mail that has not been processed.
- Access user information from `/etc/passwd` to determine the recipient's information.
- Read `.forward` files of a user to check if mail is to be delivered.
- Consult SD to place received mail for a user in his/her mail file.
- Write outgoing mails from user's mail file in SD to Q for further processing.
- Read Q to get outgoing mail.
- To send outgoing mail to another remote SMTP mail server, `foomail` first needs to get an IP for the hostname of the remote SMTP mail server. It first reads `/etc/hosts` to resolve the network hostname to an IP address.
- To send outgoing mail to another remote SMTP mail server, it first needs to get an IP for the hostname of the remote SMTP mail server. If `/etc/hosts` can not resolve the hostname of target SMTP server, it makes a network request to a DNS server to get the destination IP.
- Send outgoing mail to destination SMTP server by opening a network connection to the IP address on the remote SMTP server.

(a) Split the monolithic design of `foomail` into 2 or more individual components, that minimize the privileges required for each component. Label each of the components as A,B,C,D . . . and so on.

(a.1) [4 points] Describe the functionality provided by each component in no more than 2 sentences per component.

(a.2) [4 points] Write down for each component the kind of access (read, write, read-and-write, or none) for each of the resources described above, as a matrix with rows as the component labels and columns as resource names.

(a.3) **[2 points]** In at most 3 sentences, precisely state what security mechanisms would you use to enforce restrictions you have outlined in part (a.2) on components.

(b) **[5 points]** Consider the component, say `X`, that reads `/etc/passwd` which runs at normal user privileges. `X` is passed the name of the recipient by the component that receives the email from the network, and `X` then retrieves the user's account information from `/etc/passwd`. Suppose a remote attacker discovers a buffer overflow in `X`, which it can exploit remotely by providing a very long recipient user name. Based on the permission restrictions imposed by your solution in part (a.2), can the attacker exploit this bug to leak sensitive information in `/etc/passwd` about all users to the remote attacker directly via the network?

If not, then consider a slightly different setting. Suppose the user has installed a malicious network game application on the same computer as the one running `foomail`. It turns out that the remote attacker who knows about the buffer overflow in `X`, is also the author of the malicious game. The malicious game has been restricted access only to the network, and has no access to files in `/etc`. Can the remote attacker manage to learn the sensitive information in `/etc/passwd` in this setting?