Practice Questions for Final Exam, Fall 2008

Questions to help you prepare for final exam. These are not representative of questions on final exams -- merely to help you think about the topics in scope. The scope of questions/knowledge is not limited to those suggested by this questions -- you must use the lecture notes, slides, textbooks, homeworks, projects and other mandatory readings posted on the web page for thorough preparation.

1. What is the principle of Least Priv.?

2. An executable, say "passwd" program that changes user passwords, has setuid bit set and is owned by root. It is invoked by a user with ID 18. What will be the process's real and effective UIDs at start of the process? Suppose it calls exec on a executable that does not have the setuid bit set. What will the real and effective UIDs be after the call? What happens if the process calls "setuid(geteuid())" -- what access permissions to files will the process have after this.

3. What is the difference between "setuid" and "seteuid" on Linux, when dealing with root and non-root real UIDs?

4. How do we achieve automatic privilege separation? Describe one safe way of handling variables that are inferred both as "privileged" and "unprivileged"? Does privilege separation remove all possible attacks that caused by buffer overflows -- give on example that is not stopped?

5. Suppose you have a buggy device driver? How would you use segmentation registers on Intel x86 to apply hardware-fault isolation.

6. What are the advantages and disadvantages of SFI over hardware fault isolation?

7. How does SFI technique significantly reduce the trusted computing base (TCB) by using a verifier?

8. On the RISC architecture we studied in class, all instructions were

of fixed size (say 4 bytes). On x86, instruction sizes may vary between 1-16 bytes, and instructions can begin at any byte in the code memory. Does the classical SFI technique prevent against attacks where the sandboxed code tries to jump at in middle of an instruction stream? If not, suggest changes to fix it.

9. System call interposition -- how do we extract a policy for legitimate sequences of system calls, given the control flow graph (CFG) (code) of the program? Suppose the CFG or code is not given, you just have access to the program -- suggest one way to develop the policy to enforce with system call interposition on a application?

10. Can you use inline reference monitors for confining actions of browser plugins? Explain what properties should you enforce on the "mplayer" MPEG codec plugin for Firefox?

11. System Call interposition based monitors, SFI inline checks, Virtual machine monitors are all examples of the general concept of ..... monitors.

12. How can a infected VM communicate with a listener in another VM using CPU loading/unloading?

13. Specify two other covert channels in virtual machines (assume that the network device and hard disk controller are shared)?

14. (a) Can the TPM be used to prevent a virus from modifying the machines Master Boot Record (MBR), used for bootstrapping the OS, without being detected? If so, explain why. If not, explain why not.

    (b) Can the TPM be used to prevent a virus from modifying the machines BIOS boot block without being detected? If so, explain why. If not, explain why not.

    (c) Suppose user A is able to extract the secret AIK signing key from the tamper resistant chip in his machine. Explain the implications of this for the va- lidity of the attestation process. How could A use this key to fool a remote server about the software running on As machine?

    (d) How would you defend against this problem? You may assume that the private key extracted from the chip is published on the web (anonymously) so that anyone can mount the attack from part (c).

15. Suppose a music player vendor wishes to allow only CDs sold by that vendor to be played on its player. How can it use a special

hardware, like TPM, to achieve this? Can it use purely software techniques?

16. List atleast one other application of specialized cryptographic hardware other than those mentioned in previous two questions.

17. SQL injection can be prevented by using PREPARED statements, as seen in the homework. Explain what are "?" (bind or placeholder parameters) used in prepared statement?

18. Suggest one-way to prevent HTTP response splitting.

19. Distinguish between reflected XSS and stored XSS attacks.

20. What is the difference between a XSRF vulnerability and a XSS vulnerability.

21. You notice that a pizza purchase web site is using a MAC for some part of the data in its cookie. What could be this data, and why is the MAC being used?

22. (a) State the same origin policy as it applied to the DOM, as clearly and precisely as you can, in one or two sentences. Do the same for the same origin policy as it applies to cookies.

    (b) Why is it consistent with the same-origin policy for content from site A to include an image (such as <img src="http://anothersite.com/picture.jpg" >) from another site B?

    (c) Suppose that web pages from several sites request images from TripleClick.com. Explain how each site can pass TripleClick some information about the content of the page that will contain the image. Write a variant of the HTML <img src="http://tripleclick.com/picture.jpg" > that passes information to TripleClick as part of the request for a picture.

    (d) How can TripleClick use the requests you described in part (c) to build up a database of interests of each web user? Explain the browser mechanism that will let TripleClick tell if two requests for images come from the same user and machine, even if the user changes IP addresses.

23. In Microsoft Internet Explorer 6. This feature is a new attribute for cookies which prevents them from being accessed through client-side script. A cookie with this attribute is called an httpOnly cookie.

    (a) What attack are httpOnly cookies intended to prevent? Give an

example attack that does not work if the site uses httpOnly cookies, but works with normal cookies.

(b) Show that httpOnly cookies do not eliminate the class of attacks from part (b). Give an example where httpOnly cookies do not improve security.

24. Suppose you have a very old "rsh" server that only wishes to communicate with a client machine of a known IP address T. It naively trusts the TCP protocol design to setup a connection with T; when the client connects to the server using TCP, the server relies on the TCP 3-way handshake to guarantee that IP address of the client is indeed T and uses no other authentication mechanism.

Recall from class that the TCP/IP 3-way handshake works as follows, with C denoting the client and S denoting the rsh server:

```
C -> S : SYN (X)            % [SRC IP = T]

S -> C : SYN (Y), ACK (X)   % [DST IP = T]

C -> S : ACK (Y)            % [SRC IP = T]

C -> S
or
S -> C : Data
```

(The part in square brackets denotes the IP src/dst fields, if relevant, in the IP packets sent on the link)

(a) How do TCP sequence and ACK numbers provide some level of security, if they can not be guessed by the attacker.

(b) Consider a network attacker Eve who has _no_ access to the physical network between the client with IP address T and the rsh server, i.e Eve can not physically eavsdrop or inject or change packets in transit between them. Suppose, however, that Eve has broken the psuedo-random number generator of the rsh server and can reliably predict the next initial sequence number that the rsh server will generate.

Describe an attack that allows Eve to connect to rsh server, impersonating as the legitimate client with IP address T, without ever recieving any packet from the server during the TCP connection setup. Show your attack as sequence of messages with the source and destination fields of the IP packets sent. The last line of your attack transaction should be "Eve -> S : Malicious Data"

25. Simple Syndication Ads (ssads) is an ad syndication service that publishers can use to run advertisements on their web sites. Ssads generates banners that link to URLs of the following form: http://ssads.com/pub=3942&adv=964&landing=http://www.yoursite.com/

    Here 3942 is the publishers account number, 964 is the advertisers account number, and www.yoursite.com is the web site to which the advertiser wishes to send peo- ple who click on the banner. When ssads.com receives a request for such a URL, it charges the advertisers account for a click, credits the publishers account for a click, and generates an HTTP redirect to www.yoursite.com.

    (a) Suppose that an attacker breaks the pseudo-random generator used by ssads.coms TCP implementation for initial sequence numbers, and thus can guess the servers TCP sequence numbers with relatively good probability. Explain how a malicious publisher could exploit this fact to perpetrate click fraud.

    (b) Give one technique that ssads.com could use to thwart such attacks, even if its TCP sequence numbers are easily guessed.

26. What is a SMurF attack? How does ingress filtering in the network try to protect against it.

27. What is a SYN flooding? How do SYN cookies prevent against this attack?

28. What is a network telescope? How would you estimate the number of DOS attacks prevelant on the Internet if you had a large network telescope?

29. What is the fundamental difference, in terms of the resource being targeted, between a network bandwidth exhaustion attack and a SYN flood?

30. Suppose we had the ability to identify the IP addresses responsible for causing a DDoS attack that happened 2 hours back? Today, with the prevalence of the massive botnets, why is this capability not as useful for defense as imagined 15 years ago.

31. Is DNS cache poisoning easier when the attacker is "in-path" between the victim and the name server, that if the attacker is not in path?

32. Precisely which guarantees does DNS-SEC provide about a response from a DNS server?

33. Suppose you get a SSL certificate error "Domain Mismatch" when you visit https://foo.com. Is this likely a phishing attempt?

34. You wish to block all incoming emails with the phrase "Rush for Gold here!". Can you use a packet filtering firewall for this? If yes, how. What is an application gateway firewall, and how would you use it for this defense.

35. Suppose a company being targeted deploys a NIDS with signature like those used in Snort. Suggest two ways to send your attack payload to the companies SMTP server, bypassing the NIDS using (a) a short TTL (b) IP fragments. What assumptions did you have to make about the NIDS?