

Operating Systems Design Exam 3 Review: Spring 2011

Paul Krzyzanowski
pxk@cs.rutgers.edu

© 2011 Paul Krzyzanowski

4/18/12

1

1. Why does an IP driver need to use ARP, the address resolution protocol?

IP is a logical network. An IP address is meaningless on a physical network, such as ethernet. ARP is used to find an address for the physical network (link-layer address) that corresponds to the logical (network-layer) address of IP.

© 2011 Paul Krzyzanowski

4/18/12

2

2. Why is the use of the socket buffer (`sk_buff`) crucial for the efficient implementation of a network stack?

Without it, we'd be copying packet data for each layer of the network stack.

© 2011 Paul Krzyzanowski

4/18/12

3

3. Why is file sharing problematic under AFS?

AFS provides session semantics. Each system has its own cached version of a download file. The last system to close a file causes an upload of that version which then overwrites any changes made by other systems.

© 2011 Paul Krzyzanowski

4/18/12

4

Question 4

A loop device is a:

- network driver that echoes every packet sent to it.
- ➔ block device driver that makes any regular file look like a block device.
- file system driver that makes any file look like a file system.
- character driver that echoes any byte written to it (a `read` returns the results of the `write`).

A loop device allows one to have a block driver interface to a file. If desired, a file system can now be placed on that file.

© 2011 Paul Krzyzanowski

4/18/12

5

Question 5

You don't need to know this for the forthcoming exam

The device file system (`devfs`):

- ➔ presents a file system view of all registered block & character devices.
- is a special file system that allows the user to create device files.
- allows network interfaces to be accessed as regular files.
- is a file system that contains all the device drivers for an operating system

`devfs` is a pseudo file system that presents all block & character devices in the system. It does not allow a user to create new entries (via `mknode`). It does not contain drivers.

© 2011 Paul Krzyzanowski

4/18/12

6

Question 6

CSMA/CD (Carrier Sense Multiple Access with Collision Detection) is a technique for:

- a. an Ethernet controller to send multiple packets concurrently.
- b. an Ethernet controller to manage multiple physical Ethernet connections.
- ➔ c. an Ethernet controller to send a packet out onto the network.
- d. multiple processors to share a single Ethernet controller

CSMA/CD is used on shared (not switched) ethernet networks to ensure that a packet was transmitted. The transceiver listens on the network to detect a period of no traffic. Then it transmits the packet while listening to ensure that a transceiver on another system didn't transmit at the same time (causing a collision). If there was a collision, it waits and tries again.

© 2011 Paul Krzyzanowski

4/18/12

7

Question 7

In the Internet Protocol (IP), a port number is present at the?

- a. network layer.
- ➔ b. transport layer.
- c. presentation layer.
- d. application layer.

The transport layer allows applications to talk with other applications as opposed to the network layer, which only handles routing a packet from one machine to another. A port number is used to identify the socket used by an application.

© 2011 Paul Krzyzanowski

4/18/12

8

Question 8

Which socket-related system call is not needed in a connection-oriented client?

- a. socket
- b. bind
- ➔ c. accept
- d. connect

The `accept` system call is used by a server process to wait for an incoming connection.

© 2011 Paul Krzyzanowski

4/18/12

9

Question 9

TCP/IP differs from UDP/IP in that TCP/IP:

- a. requires the use of port numbers.
- b. detects errors in a packet.
- c. provides guaranteed bandwidth.
- ➔ d. provides in-order packet delivery to the application

TCP tries to simulate a circuit-switched connection and provides reliable, in-order delivery.

- (a) Every transport-layer protocol requires this: UDP & TCP
- (b) UDP also detects errors but it just drops the packet. TCP requests a retransmission
- (c) TCP has no control over the available bandwidth on a network.
- (d) Each packet has a sequence number. Out of order packets are buffered in the network stack and delivered to the application in sequence.

© 2011 Paul Krzyzanowski

4/18/12

10

Question 10

In contrast to the older non-NAPI implementation, the Linux NAPI approach to getting data from a network interface card:

- a. coalesces multiple packets into a single socket buffer (`sk_buff`).
- b. generates a single socket buffer (`sk_buff`) for each received packet.
- c. uses interrupts from the network controller for instant response instead of relying on polling.
- ➔ d. uses a combination of interrupts and polling.

NAPI was designed to avoid high rates of interrupts from network transceivers by shutting off interrupts from the network card after the first interrupt and resort to polling. If polling yields no incoming packets, interrupts are then re-enabled.

© 2011 Paul Krzyzanowski

4/18/12

11

Question 11

The network device driver is responsible for which layer of the OSI stack?

- ➔ a. Data link.
- b. Network.
- c. Transport.
- d. Presentation.

The device driver itself interacts with the physical network transceiver and hence implements the data link layer: the mechanics of creating ethernet packets, getting a packet, and sending it out.

The network & transport layers are not handled by the device driver but by the network stack.

The presentation layer is handled at the application layer.

© 2011 Paul Krzyzanowski

4/18/12

12

Question 12

The interface definition language (IDL) was created to:

- ➔ a. enable the use of remote procedure calls with languages such as C, C++, Java, etc.
- b. extend the syntax of conventional languages (such as C, C++, Java) to support remote procedure calls.
- c. serve as a high-level programming language that understands remote procedure calls.
- d. serve as a system call interface between user processes and kernel RPC facilities.

The IDL was created to allow a *precompiler* to create client stub functions and a server program for channeling remote procedure requests. It is not a programming language; it's essentially a series of function declarations similar to function prototypes.

© 2011 Paul Krzyzanowski

4/18/12

13

Question 13

Remote procedure calls rely on marshaling parameters. Marshaling means:

- ➔ a. transforming the parameters into a form suitable for network transmission.
- b. sending the parameters over the network to a remote server.
- c. having the remote server extract the parameters from the network message.
- d. having the remote server place the parameters on the stack to make a procedure call.

Marshaling is placing all the parameters into a stream of bytes in an agreed-upon format so that they can be transmitted. Unmarshaling is the opposite process (c).

© 2011 Paul Krzyzanowski

4/18/12

14

Question 14

What facilities does the operating system kernel provide to allow user processes to use remote procedure calls?

- a. Interface definition language.
- b. Remote procedure system call.
- ➔ c. Sockets.
- d. All of the above.

Remote procedure calls are not provided by the operating system; they're a function of the language (or supporting tools for the language). The operating system provides *sockets* as a mechanism for accessing communication networks. Everything else is built on top of that.

© 2011 Paul Krzyzanowski

4/18/12

15

Question 15

An NFS client invalidates locally-cached data via:

- a. callbacks: receiving a message from the server informing it to invalidate the data.
- b. tokens: getting tokens from a server telling the client how it may cache data.
- c. disabling caching for all remote data.
- ➔ d. validation: checking the timestamp on a remote file.

NFS checks the timestamp on a remote file IF it happens to be making a remote request. Otherwise, it uses the cached version but gives it an expiration time of several seconds.

© 2011 Paul Krzyzanowski

4/18/12

16

Question 16

Privilege separation is when:

- a. users on a system get unique user IDs and privilege levels.
- ➔ b. a process is split into components with limited privileges.
- c. a process cannot create data that processes running at lower privilege levels can access.
- d. a process cannot interact with any other process or data at another privilege level.

Privilege separation is where each functional component is provided with only the privileges it needs to do its job. It's often implemented by running multiple communicating processes, with different privileges in each process.

© 2011 Paul Krzyzanowski

4/18/12

17

Question 17

A capability list is a list of:

- a. per-user permissions associated with an object.
- ➔ b. permissible operations on objects associated with a user.
- c. system capabilities associated with a process.
- d. permissible operations for a specific user and object.

An access control list associated a list of per-domain permissions with objects (domain = users or groups of users)

A capability list is the opposite way of implementing an access control matrix: a list of per-object permissions is associated with each domain (user or group of users)

© 2011 Paul Krzyzanowski

4/18/12

18

Question 18

A Multi-Level Secure model:

- a. uses multi-factor authentication to access system objects.
- b. allows a process to assign a classification level to objects that it creates.
- c. allows a process to assign multiple classification levels to objects it creates.
- ➔ d. limits the ability of a process to create objects that another process can access.

MLS is an example of Mandatory Access Control. System-wide policies restrict how a user can communicate and the user cannot override these restrictions.

© 2011 Paul Krzyzanowski

4/18/12

19

Question 19

The Bell-LaPadula model is an example of:

- a. applications using cryptographically secure communications.
- b. an access control matrix.
- ➔ c. mandatory access control.
- d. discretionary access control.

The Bell-LaPadula model organizes users in a hierarchy of increasingly higher classification levels. At each level, a process can read data from the same or lower levels but not from higher levels. Conversely, a process can write data to the same or higher levels but not to lower levels.

This is a system-wide policy that the user cannot override.

© 2011 Paul Krzyzanowski

4/18/12

20

Question 20

To authenticate himself to Alice, Bob would encrypt a nonce sent by Alice with:

- ➔ a. his private key.
- b. his public key.
- c. Alice's private key.
- d. Alice's public key.

To authenticate, Bob has to do something that nobody else can do.

(b) and (d) are not valid choices because anyone can get access to public keys. (c) is not a valid choice because only Alice has access to her private key.

By using his private key to encrypt a nonce (a random bunch of bits) that Alice sent him, Bob can do something nobody else can. Alice can then verify that Bob's private key was used by decrypting the result using Bob's public key.

© 2011 Paul Krzyzanowski

4/18/12

21

Question 21

To send a secure message to Alice, Bob would encrypt the message with:

- a. his private key.
- b. his public key.
- c. Alice's private key.
- ➔ d. Alice's public key

To communicate securely, Bob has to make sure that nobody but Alice can decrypt the message. With public key cryptography, each user has a pair of keys (one called a public key and the other called a private key). If you encrypt a message with one key, you have to use the other key in the pair to decrypt. If he encrypts with Alice's public key, only Alice has the corresponding private key to decrypt the message.

If Bob encrypts with his private key, anyone can decrypt the message using his public key. If he encrypts it using his public key, only he can decrypt the message using his private key.

© 2011 Paul Krzyzanowski

4/18/12

22

Question 22

Which is NOT a technique that the operating system can use to guard against buffer overflow attacks?

- a. Address space layout randomization (ASLR).
- b. Non-executable stacks.
- ➔ c. Stack canaries.
- d. All of the above are viable operating system techniques to guard against buffer overflow attacks.

This is a trick question because it's asking about a technique that the operating system cannot use. ASLR and non-executable stacks can be implemented by the operating system; the program loader can assign random offsets for loading libraries and the memory management unit can have pages holding stack and heap data set to no-execute.

Stack canaries, however, are not implemented by the operating system but are code generated by the compiler at the entry and exit of each function.

© 2011 Paul Krzyzanowski

4/18/12

23

Question 23

A DES key is 56 bits. A triple-DES key is:

- a. 56 bits.
- b. 112 bits.
- c. 168 bits.
- ➔ d. Any of the above.

The standard DES algorithm uses a single 56-bit key. Triple DES uses three 56-bit keys and runs the algorithm three times:

Encrypt with K1 → Decrypt with K2 → Encrypt with K3

If K1=K2=K3, we have a 56-bit key.

If K1=K3, we have a 112-bit key.

If K1, K2, and K3 are all different, we have a 168-bit key.

© 2011 Paul Krzyzanowski

4/18/12

24

Question 24

The Diffie-Hellman algorithm is:

- a. a public key encryption algorithm.
- b. a combined key exchange and public key encryption algorithm.
- c. a symmetric encryption algorithm.
- ➔ d. not an encryption algorithm but allows two parties to compute a common key.

Diffie-Hellman is not an encryption algorithm.

© 2011 Paul Krzyzanowski

4/18/12

25

Question 25

You don't need to know this for the forthcoming exam

Which of the following is FALSE about the Secure Sockets Layer (or Transport Layer Security)?

- a. Public key cryptography is used for authentication.
- b. Symmetric cryptography is used for data communication.
- ➔ c. A one-way hash function is used to store keys.
- d. A session key is established for data exchange

© 2011 Paul Krzyzanowski

4/18/12

26

Question 26

What best describes Kerberos?

- a. A database that keeps track of malicious IP addresses/machines.
- ➔ b. An authentication service that uses a trusted third party that holds all the passwords.
- c. A program developed in 1989 that used SYN flooding to attack specific machines.
- d. The first anti-virus program that used virus-scanning to find a signature of a virus.

Kerberos is a trusted third party that knows everyone's passwords. If one service (or user) wants to talk to another service, it has to contact Kerberos for authorization and to get a session key to use in communicating with that service.

© 2011 Paul Krzyzanowski

4/18/12

27

Question 27

A digital certificate is:

- a. a hash of a digital document that is encrypted with your private key.
- b. a digital document that is encrypted with your public key.
- c. your identity and your private key, digitally signed.
- ➔ d. your identity and your public key, digitally signed.

A digital certificate is a way of distributing public keys where the recipient can validate whose key it is and that it was not modified. The key elements of a digital certificate are the public key, information about the owner of the public key, the certificate issuer, and the signature. The signature is a hash of all the data that is encrypted with the certificate issuer's private key.

© 2011 Paul Krzyzanowski

4/18/12

28

Question 28

Ethernet is an example of a broadband network.

True

False

False.

A broadband network is one where the network bandwidth is broken up into frequency bands (channels) so that everyone can talk at the same time but only on their band. Ethernet is an example of a baseband network, where you share time with other users on the network but, for your time of access, you have access to the full bandwidth of the network.

© 2011 Paul Krzyzanowski

4/18/12

29

Question 29

FUSE, the file system in user space, uses a loop device to bridge the VFS layer to a user process.

True

False

False.

FUSE is actually implemented via a kernel module that appears as a VFS file system. This module interacts with the user-level FUSE library via a special file descriptor obtained by opening /dev/fuse.

© 2011 Paul Krzyzanowski

4/18/12

30

Question 30

Sockets support the ability to use file system *read* and *write* system calls because sockets are implemented under the VFS layer.

True

 False

False.

A socket is designed to be compatible with files to some degree and sits in the same table as file descriptors (open files). To maintain this compatibility, it has to implement a set of file operations.

However, sockets are not implemented as a file system and hence is not in the VFS layer.

© 2011 Paul Krzyzanowski

4/18/12

31

Question 31

NFS uses the download/upload model for remote file access.

True

 False

False.

NFS uses a remote-access model, where operations such as read *n* bytes, write *n* bytes, get attributes, etc. are sent as remote procedure calls.

© 2011 Paul Krzyzanowski

4/18/12

32

Question 32

The Password Authentication Protocol with stored hashed passwords provides network-safe authentication.

True

 False

False.

No. With PAP, the password is sent in clear text and can be sniffed on the network.

Hashed password storage prevents somebody from knowing your password even if they see it in a file. To validate, one needs to hash the password and compare the result with the stored hash.

© 2011 Paul Krzyzanowski

4/18/12

33

Question 33

A prominent feature of the original NFS design is that the servers keep no state.

 True

False

True.

The original design of NFS was stateless.

As a result, file locking could not be implemented. Also, a file could be deleted on the server even if it was still open and in use by a remote client.

© 2011 Paul Krzyzanowski

4/18/12

34

Question 34

You don't need to know this for the forthcoming exam

Secure Sockets Layer (SSL, or Transport Layer Security) is an example of a hybrid cryptosystem.

 True

False

True.

SSL uses public key cryptography for authentication and key exchange followed by symmetric cryptography for communication.

© 2011 Paul Krzyzanowski

4/18/12

35

Question 35

A Kerberos ticket (sealed envelope) contains the session key.

 True

False

True.

Kerberos returns two messages to someone who requested access to a service: (1) a message encrypted for you that contains the session key you will use to communicate with that service and (2) a message encrypted for that service which contains the same session key. This second message is called a ticket and you cannot decrypt it but the remote service can once you send it.

© 2011 Paul Krzyzanowski

4/18/12

36

Question 36

Kerberos uses public key cryptography to provide secure communication on insecure networks.

True

False

False.

Kerberos is implemented entirely with symmetric cryptography. Because Kerberos is a trusted third party, everyone's key is shared with Kerberos. There is no problem of key explosion because the key that you use to communicate with a service is a session key that is generated on demand by Kerberos and then thrown away.

© 2011 Paul Krzyzanowski

4/18/12

37

Question 37

A socket is simply a number included at the start of each IP packet.

True

False

False.

A socket is a communication endpoint used by an application. For sockets over IP, there will be a port number that is associated with the socket that ends up in the TCP or UDP header.

© 2011 Paul Krzyzanowski

4/18/12

38