

CS 162 Fall 2003
Discussion Section Quiz 6
TA: Steve Martin

1. What are the primary characteristics of a computer virus? What are some ways a computer virus can spread? How can a virus be detected?

Malicious code that executes on a user's machine, and often replicates itself elsewhere. Viruses can spread over any medium over which files are transferred, from floppy disks (80's) to networks. Usual detection methods revolve around generating a unique signature of the virus code and then scanning for files that match that signature.

2. Suppose we were using a security system using a key server. Assume that you are on machine S, and you want to communicate securely to machine T. All you have is a key that allows you to talk to the key server securely. How would you gain the ability to have secure communications with machine T?

Step 1. Send a request to the server for a key to communicate to T, and encrypt with the key S shares with the server.

Step 2. Server will send back a key S can use to talk to T and a message with the key to send to T, but encrypted with the key it shares with T. This is all encrypted in the key S shares with the server.

Step 3. S then sends the message to T, which T can decode with its own key. T then can communicate with S using this key.

3. Why is a completely random one-time pad (running key cipher) considered 'perfect encryption'?

Because if the key is completely random (i.e. not periodic at all) then the output will be completely randomized as well, leaving no frequency information for attackers to exploit in cracking the code. Note that pseudorandom numbers are not completely random and make one-time pads weaker.

4. What is a Virtual Machine Monitor? Why are they useful?

A VMM is a program that transparently emulates an entire computer to another piece of software running on top of it. These are useful for many reasons, such as maximizing

resource utilization and performance in wide scale non-uniform distributed systems, device driver development, OS development, etc.

5. There are many problems with getting VMMs to work correctly. Many of these issues involve effectively multiplexing system resources in a transparent manner. Memory management we discussed in class; describe how you might manage distributing A) CPU time and B) the hard drive.

There are several problems to be solved for this to work correctly. Here's a short sketch of some of them:

CPU: because VMMs run in a user process, the OS has no concept of what is going on within the virtual process. For this reason, the scheduler can reschedule vm processes at will, since even virtual operating systems that are 'idling' will still generate work for the CPU at the real process level. How to solve this? Well, must somehow allow VMMs to effect CPU scheduling so that only VMMs with useful work can run.

Disk: we need to write blocks to a virtual disk. However, because the blocks are virtual as well, need to somehow translate between virtual blocks and real ones. For this reason we need to intercept all calls to IO on the virtual operating system and do some translation before the job can be scheduled. This holds for other types of asynchronous IO as well.