CS4410 Homework 5

HW5 WILL NOT BE GRADED. SOLUTIONS WILL BE POSTED ON TUESDAY APRIL 1

1. (If we graded it, would have been worth 10 points). Describe a situation in which your computer might have multiple IP addresses over the course of a single day.

At home, my laptop was connected via a home wireless network; at work I connected to an Ethernet, then later took it to class and used Red Rover. Each use had its own IP address.

2. (If we graded it, would have been worth 10 points). Suppose that you and your friend both check your respective IP addresses at the exact same moment in time. Could both claim to have IP address 196.168.1.27? Explain.

Absolutely. Perhaps we both own the same brand of wireless router and the devices are using the same numbering system for IP addresses they hand out. These IP addresses work only in very limited environments: my IP address might be useable between my laptop and the printer, or between it and the wireless router, but the "outside world" would see my machine under the address (and a port number) associated with the wireless router. The router itself does network address translation. Thus from the outside world's perspective, there these internal addresses might not be meaningful or even accessible.

3. (If we graded it, would have been worth 10 points). Suppose that someone is using a packet snooper technology such as "ethereal" on the Ethernet while you read or write a file called "homework 5 solutions" which is stored on the NFS server in your department. Furthermore, assume that Sun's proprietary NFS cryptographic security features are not in use. What could the intruder learn, if anything?

They can see the contents of any file blocks that you access via NFS read or write.

4. (If we graded it, would have been worth 10 points). Summarize the security properties of NFS when Sun's proprietary NFS cryptographic security is turned off, as is the "default" for this file system. (Recall that in many settings, NFS is used with a mixture of other technologies, and that this can preclude use of the NFS cryptographic security. Also, NFS security relies on a cryptographic protocol that cannot be exported from the United States. This is why NFS cryptographic security is often disabled.)

NFS doesn't really have any security properties at all in this mode. Anyone can really do anything. Obviously, some pretense of respecting file access permissions occurs, but in fact you can construct NFS requests with fake IDs in the fields associated with file access and having done so, you can access files under IDs that you really shouldn't be allowed to touch. For example, you could build a message in the official NFS format claiming to be Bill Gates and send it to an NFS with his bank account data, and if your message can reach the NFS, it will accept that you are Bill Gates without doing any extra checking. Of course the normal NFS software only puts sensible things into user ID and group ID fields of messages, but nothing stops you from building this sort of "broken client" except the hassle of doing it. (And with FUSE available for free download, the hassle isn't very much!)

5. (If we graded it, would have been worth 10 points). Suppose that two different users, on different computers, nearly simultaneously open a file on an NFS server. The open from user A occurs just before the open from user B. They use the same file name and both opens succeed. Now both users write 10 blocks into the file, concurrently. Later you print the file. What would you expect to see in it?

Any mixture of the blocks would be possible here. You could see A's versions, B's versions, or some blocks from each (this last seems more likely given the description in the problem statement).

6. (If we graded it, would have been worth 10 points). Suppose that you had a private, dedicated, 10 Gbit link from New York to Bangalore, round-trip-time 200ms. Would you expect that a TCP connection from a machine in New York to a machine in Bangalore could run at the full speed of the link? What factors would limit performance?

TCP will be limited by the size of the TCP Window, which by default is 128Kbytes. A 10GBit link could need 500 Mbytes of windowed data to run continuously. So the performance will be far lower than the theoretical speed of the link. In fact, though, your computer has a network interface card and that card certainly won't run at 10 Gbits either. So: your NIC will limit you, and the TCP window size. If any loss is occurring on that long link, the need to send a NAK from the remote side back to New York, and to wait for the retransmission to get through, could also cause problems: each time this happens the link will pause for 400ms.

7. (If we graded it, would have been worth 10 points). Just as you start to download a video to your computer over a wireless link, your roommate decides to make a fruit frappe and runs the blender. The blender generates a lot of electronic noise and 50% of the wireless packets are corrupted and dropped. Assume that she runs the blender until your file transfer finishes. If it normally takes T seconds to download the video, would you expect it to take more, less, or exactly 2T seconds under these conditions? Why?

Probably much more than 2T seconds. TCP will assume that the loss is caused by congestion (TCP has no other behavior) and so each lost packet causes TCP to slow the transmission rate down. It could take 100T for all we know: when TCP slows down it can slow to a crawl! ("Multiplicative rate descrease").

8. (If we graded it, would have been worth 10 points). An application uses TCP to talk to a remote machine but it sends data in bursts with long delays between the bursts. Sketch a graph of the likely throughput for this application.

We actually discussed this in class: assuming the interval between bursts is long enough, TCP will keep falling back into slow-start mode and you get a graph like the one below. Conversely, if the delays are not quite long enough to trigger slow start, then TCP will just alternate between the basic sawtooth and periods with zero bandwidth (because nothing is being sent). The use of the word "long" in the question makes it more likely that the graph from class was the intended one:

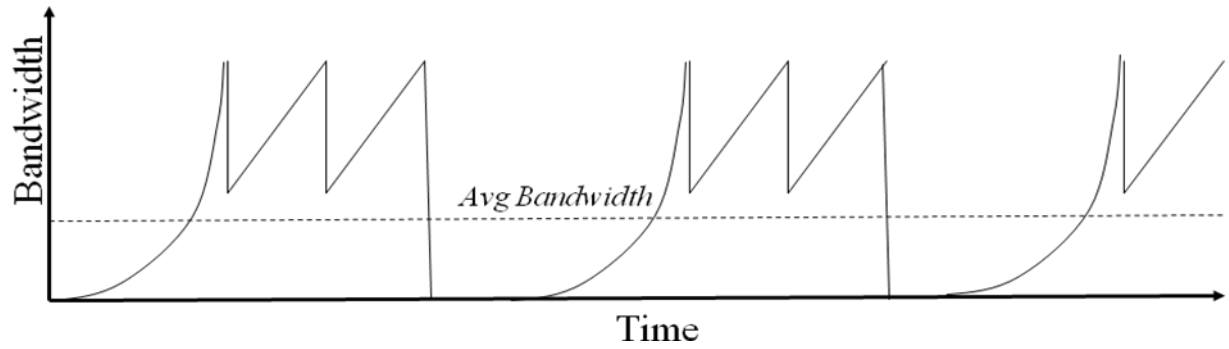

9. (If we graded it, would have been worth 10 points). Suppose that you are viewing a web page from your bank. Describe some reasons that some of the things shown on the web page might actually not have come from your bank.

When you download data from a site like mybank.com, unless you use the HTTPS secure protocol, the DNS resolves the address and could actually misdirect you to a site like fakebank.ru (web site spoofing). In addition, the actual page may have frames that were really downloaded from other sources such as advertising web sites, offline image hosting systems, etc. Finally, some content is rendered on your machine by downloaded Javascript; this is code and could decide to show a picture other than what the bank itself was intending to send.

10. (If we graded it, would have been worth 10 points). Read about the new services associated with the Google internet telephone platform, Google Voice. List some things that you might normally think of as private that Google might be in a position to learn about if you adopt Google Voice as your standard telephone service.

They would potentially be able to create logs of your location (via GPS), who you were talking to, when, who you spend time with (by noticing other people who had the same GPS data at the same time), and even would have access to your voice mail or to transcripts of your calls, if you enable those features of Google Voice (Google Voice has a feature to automate creation of transcripts). Same goes for pictures you might share while talking with your friends or other kinds of media content. All of this is stored on the Google site and hence indexed by them.