System Administration for the Web: Week 3 Lab

13th February 2006

1 Lab for Week 3

Due to security concerns, most of the commands you'll need to perform this laboratory have been removed from the servers in Soda Hall. Thankfully, the OCF provides access to all the commands you need. The homework assignment for the first week was to create an OCF account, so we'll be using your OCF account to perform this lab.

NOTE: Please refer to Section **??**: Submitting the Lab before you begin.

1.1 Logging Into the OCF

We're going to use SSH to log into your OCF account. SSH is a nifty tool that allows you to log into a remote server over an encrypted connection, among other things.

[1] Open a terminal window. Enter "ssh [your OCF login]@ocf.berkeley.edu".

[2] You'll be asked for your current OCF account password. Please enter it. If you get a message asking you if you want to continue connecting, please enter "yes".

From now on, you'll do everything in this terminal window. If for some reason you close the terminal window or need to use another one, you'll have to log into your OCF account in the new window.

1.2 IP Addresses

Every computer connected to the Internet has a unique address called an *IP address*. IP addresses are a set of 4 numbers separated by periods, like `128.32.42.39`. Each number can be between 0 and 255. Most IP addresses owned by Berkeley have 169 and 229 as their first two numbers; that is, most IP addresses owned by Berkeley are in the form `169.229.x.x`, where x is a number between 0 and 255.

[1] Given the above information, how many IP addresses can Berkeley have in that *address space* – that is, theoretically, how many unique IP addresses can Berkeley have in the form `169.229.x.x`?

[2] In total, how many unique IP addresses are possible?

[3] The current world population is just under 6.5 billion. Are there enough IP addresses for every single person in the world? While this may sound like a stupid question, please note that corporations and organizations (like Berkeley) often own large blocks of IP addresses, reducing the number of publicly available IP addresses.

1.3 DNS

[1] *host* is a command to translate a hostname into the IP address to which it is linked. Read the man page for *host* and determine the syntax for translating `solar.cs.berkeley.edu` into its IP address. Please note that when man pages specify a command's syntax, anything in []'s are optional parameters that you probably don't need.

[2] As explained above, IP address translation can also work in the reverse direction. Run *host* on the IP address you obtained for `solar.cs.berkeley.edu`. This is called looking up the *reverse DNS*. Please note that the reverse DNS does not necessarily have to match the *forward DNS* or even exist.

[3] Now look up the IP address for `www.berkeley.edu`.

The output should look different from what you got when you looked up the IP address for `solar.cs.berkeley.edu`. This is because hostnames can be actually linked to other hostnames. There are different classes of hostnames (more generally called *records*), and they do not necessarily have to be linked directly to an IP address. Hostnames that are linked directly to IP addresses are called *A records*. Hostnames that are linked to another host name are called *CNAME records*.

[4] Open a web browser and use a search engine to look up two more types of DNS records. Write them down.

[5] Let's find more about `www.berkeley.edu`. *dig* is a command to query a DNS server to find out more about its DNS records. Please *dig* `www.berkeley.edu` and see if you can locate the A and CNAME records, as well as any records of the type you looked up in the previous exercise.

[6] *dig* a few more of your favorite websites. If you notice a strange record, ask an instructor to explain it to you.

1.4 Ping

As a systems administrator, you'll want to make sure your servers are always up and running. Unfortunately, you may not always be physically near you servers to check on them, and even if you were, just seeing the on light wouldn't tell you if your server was still connected to the Internet. A tool to check the connectivity of a server is *ping*.

[1] *ping* `solar.cs.berkeley.edu` to make sure if its connected to the Internet.

Sometimes systems administrators block ping requests with a firewall. This is because *ping* can be used by malicious people to attack your server and flood it with requests.

[2] *ping* `www.cnn.com`. Usually, when you get no response from the 'ping' command, it means that the server is either down or has blocked your request. Recall that you can press Ctrl-C to abort a command.

1.5 Traceroute

traceroute is a tool to follow the path data will take between you and a recipient server. It will list the routers it passes through along the way, as well as the round-trip time it takes to reach a certain router. By using *traceroute*, you can determine if a particular Internet link is responsible for any slowness in your connection.

[1] *traceroute* the path to `blizzard.ocf.berkeley.edu`.

[2] *traceroute* the path to `www.stanford.edu`.

[2] *traceroute* the path to `www.mit.edu`.

[3] Besides the different names of routers, what is the difference between the outputs you received for the last three exercises? Why is this so?

[4] *traceroute* the path to your favorite websites. Occasionally *traceroute* will fail after a number of *hops*. You'll have to abort the command or wait for it to finish. Do you remember how to abort a command?

1.6 Whois

[1] To register a top-level domain, ICANN regulations state that you must provide contact information for your domain to the public. Usually this involves your name, address, and phone number. *whois* is a command to look-up this information. Run *whois* on `berkeley.edu`.

[2] While corporations and organizations can provide a business contact, most individuals can not. If you own a domain or know somebody that does, run *whois* on the domain.

2 Submitting the Lab

In order to ensure that you've completed the laboratory, please send the results of your laboratory to dima+decal@ocf.berkeley.edu. In the case where an exercise produces output, please paste the relevant information into the email. On a UNIX system, you copy by simply highlighting the text you wish to copy, and you may paste by pressing the middle mouse button (or holding down both mouse buttons if no middle mouse button exists).

Along with your laboratory, submit your name, OCF login, and inst login. We need this information to associate your name with your work.