CS 5434 HW 3 – Fall 2015

Due Monday Nov 9th @ midnight.

In this homework you will build on your homework 2 code in order to develop a portscan detector. Work in the same groups as for hw2, but rename the new version "hw3". Use the same VM as provided for HW2.

Your program should be able to read a series of pcap files (specified on it's command line via successive "-r my.pcap" options) using libpcap (see "man 3 pcap" for details). In addition to the flow table of hw2, you should now add a source table keyed on the source-ip of the client that initiates each flow. Use a hash table with chaining with the hash key based on the src-ip. You can review the algorithms discussed in class for ideas of what state to keep in your source table:

http://www.cs.cornell.edu/courses/CS5434/2015fa/Lecture7-port-scan.pdf
http://www.cs.cornell.edu/courses/CS5434/2015fa/Lecture8-port-scan-worms.pdf

For this exercise, we have provided a test pcap that contains background traffic:

http://www.cs.cornell.edu/courses/CS5434/2015fa/port_scan.pcapng

The background traffic is a sample of about 20 minutes of real traffic from a large network where the packet bodies have been removed and the IPs rewritten.

There are twelve port scans embedded in the traffic, each from a different source IP. Your goal is to find them all. A number of them are large, fast and noisy, but several are relatively small and subtle. Not all of them are syn or connect scans. Your goal is to find all the scans. Your code, when run on the file, should report the source-ip, the time of the first and last scan attempts, and the number of ports scanned. Most of the credit will be given for finding the correct source-ip and approximate time (ie if you miss a few packets at the beginning or end of the scan, you'll still get most of the credit).

If your code emits false positives (ie reports on activity that is not a portscan) credit will be deducted.

Note, it's possible that there will turn out to be extra port scans in the background traffic. If so, reporting those will be neutral (neither credit given, nor taken away). If any such extra scans come to our attention during the homework period, we will announce them on Piazza.