

CS 395T - Theory and Practice of Secure Systems  
Fall 2006

Homework #1

Due: 3:30pm CDT (in class), October 10, 2006

**NO LATE SUBMISSIONS WILL BE ACCEPTED**

**YOUR NAME:** \_\_\_\_\_

**Collaboration policy**

**No collaboration** is permitted on this assignment. Any cheating (*e.g.*, submitting another person's work as your own, or permitting your work to be copied) will automatically result in a failing grade. The Computer Sciences department code of conduct can be found at <http://www.cs.utexas.edu/users/ear/CodeOfConduct.html>

## Homework #1 (30 points)

### Problem 1 (4 points)

Give a short snippet of C code which contains a single call to a `libsafe`-protected `strcpy`, and yet is vulnerable to a buffer overflow attack as a result of this call.

### Problem 2 (6 points)

Memory safety mechanisms such as virtual memory and software-based fault isolation prevent one application from corrupting the data or code of another application. Give at least **three** attacks by a malicious application which cannot be prevented by memory safety mechanisms, but can be stopped using system call interposition.

### **Problem 3**

A famous Norwegian hacker OpKødë proposes the following defense against buffer overflow exploits. All code pages, including system library routines, should be mapped to low memory addresses (*e.g.*, from 00000000 to 000FFFFFF on x86 machines), and the rest of the pages should be marked as non-executable.

#### **Problem 3a (4 points)**

Would this technique prevent buffer overflow exploits? What about `return-to-libc` exploits? Explain.

#### **Problem 3b (3 points)**

What are the advantages and disadvantages of protecting memory in this way?

### Problem 4 (4 points)

Give an example of a security property which **cannot** be expressed as a finite-state automaton, and write a snippet of C or pseudo-code which violates this property.

### Problem 5 (4 points)

Will increasing the salt size in UNIX password hashes from 12 bits to 48 bits provide much better security against password cracking? Explain your answer.

## Problem 6 (5 points)

Acme Security has developed an intrusion detection system (IDS), which is designed to detect malicious port scans as well as connections with spoofed IP addresses. It boasts an impressive accuracy rate, reflected in the following table:

Type of connection	How this connection is classified		
	Port scan	IP spoofing	Legitimate
Port scan	85%	5%	10%
IP spoofing	5%	90%	5%
Legitimate	5%	5%	90%

For example, when Acme's IDS observes a malicious port scan, it correctly classifies it as a malicious port scan with probability 85%, misclassifies it as an IP spoofing attack with probability 5%, and misclassifies it as a legitimate connection with probability 10%.

For the purposes of this problem, assume that port scans are 3% of all connections, and that IP spoofing attacks are 1% of all connections, while 96% of traffic consists of legitimate connections. Also assume that a connection cannot be *both* a port scan and an IP spoofing attack at the same time.

When Acme's IDS announces that it detected a port scan, what is the probability that the connection is, in fact, legitimate? Give your calculations.