Name ..... ID ..... Section ..... Seat No .....

Sirindhorn International Institute of Technology
Thammasat University

Final Exam Answers: Semester 2, 2013

Course Title: CSS322 Security and Cryptography

Instructor: Steven Gordon

Date/Time: Thursday 13 March 2014; 13:30–16:30

Instructions:

- This examination paper has 17 pages (including this page).
- Conditions of Examination: Closed book; No dictionary; Non-programmable calculator is allowed
- Students are not allowed to be out of the exam room during examination. Going to the restroom may result in score deduction.
- Turn off all communication devices (mobile phone etc.) and leave them at the front of the examination room.
- The examination paper is not allowed to be taken out of the examination room. A violation may result in score deduction.
- Write your name, student ID, section, and seat number clearly on the front page of the exam, and on any separate sheets (if they exist).

Question 1 [13 marks]

In Diffie-Hellman key exchange, each user chooses a private X, calculates a public Y, and after exchanging Y, calculate K where for user A:

- $Y_a = \alpha^{X_a} \bmod q$
- $K_a = Y_b^{X_a} \bmod q$
- q is a prime number, $X_a < q$, and α is a primitive root of q

You and other users have agreed to use the public values $q = 11$ and $\alpha = 2$. You just received a message from Steve containing his public value, $Y_b = 5$. You choose a private value $X_a = 3$.

(a) What public value will you send to Steve? [2 marks]

Answer.

$$Y_a = \alpha^{X_a} \bmod q$$
$$Y_a = 2^3 \bmod 11 = 8$$

(b) What secret will you share with Steve? [2 marks]

Answer.

$$K_a = Y_b^{X_a} \bmod q$$
$$K_a = 5^3 \bmod 11 = 4$$

Assume Thanaruk has intercepted the packets used in the Diffie-Hellman key exchange between you and Steve.

(c) Show the steps Thanaruk takes to determine Steve's private value. [3 marks]

Answer. *Thanaruk knows the public values: $q = 11$, $\alpha = 2$, $Y_a = 8$, $Y_b = 5$. He will need to find one of the private values. He knows the equations for determining Y:*

$$8 = 2^{X_a} \bmod 11$$

and

$$5 = 2^{X_b} \bmod 11$$

Brute force on either X is possible, but the question is asked for Steve's private value, i.e. X_b. Try $X_b = 0, X_b = 1, X_b = 2, \ldots$ until 5 is obtained. The answer is $X_b = 4$.

(d) What value(s) are recommended to ensure Diffie-Hellman key exchange is secure, i.e. so Thanaruk could *not* determine the secret shared between you and Steve? [2 marks]

Answer. *q must be very large.*

For the following question(s), assume appropriate values were chosen for the Diffie-Hellman key exchange (i.e. those you recommended above).

(e) Explain how a malicious user Thanaruk (user C) can perform a man-in-the-middle attack on two users (A and B) that perform a Diffie-Hellman key exchange. You may use a diagram to illustrate the attack. You should use the variables (e.g. Y_b, α) in your description (don't use the values such as 2 and 5 from the above questions). [4 marks]

Answer. *When Y_a is sent to B, C intercepts and calculates its own Y_c, sending Y_c onto B. When B replies with Y_b, C intercepts and sends Y_c to A. The end result is A uses Y_a and Y_c as public values and determines some secret K_{ac}. C will also be able to calculate K_{ac}. Similar, B determines some secret K_{cb}, which C also has. A and B don't know that C has intercepted, so when they encrypt data with their secrets, C can intercept, decrypt and then encrypt again using the destinations shared secret.*

Question 2 [11 marks]

You have three algorithms that generate random passwords for new users on a system:

Algorithm 1 Select a character from the set of: English lowercase and uppercase characters, as well as the six punctuation characters < > [ ] { }

Algorithm 2 Select a character from the set of: digits and five operators + - / * =

Algorithm 3 Select the first character to be a digit, then all remaining characters are chosen from the set of: English lowercase and uppercase characters

(a) How many passwords are possible if a 5 character long password is generated using Algorithm 1? [2 marks]

Answer. *There are 58 possible characters, so there are* 58^5 *possible passwords*

(b) What is the entropy of a 10 character long password generated using Algorithm 2? [2 marks]

Answer. *There are 15 possible characters, so there are* 15^{10} *possible passwords. Therefore the entropy is* $\log_2(15^{10}) = 39.07$.

(c) You know that a random value should be 80 bits in length to withstand a brute-force attack. For each of the three algorithms, give the minimum password length recommended to withstand a brute-force attack. [5 marks]

Answer. *The entropy of each character using Algorithm 1 is* $log_2(58) = 5.858$. *For Algorithm 2 it is 3.907. For Algorithm 3, the entropy of the first character is* $\log_2(10) = 3.322$, *while the entropy of the remaining characters are each* $\log_2(52) = 5.700$. *The entropy of the resulting passwords should be 80. Therefore using Algorithm 1* $80/5.858$ *or 14 characters are needed. For Algorithm 2, 21 characters are needed. For Algorithm 3, the 2nd to last characters must produce an entropy of 76.678, therefore 14 extra characters are needed, i.e. a password length of 15.*

(d) Even if the password generated with any of the algorithms is long enough to withstand a brute-force attack, give two reasons why it may be a poor password? [2 marks]

Answer. *Hard to remember; hard to type (make mistakes)*

Question 3 [13 marks]

Consider the mechanism illustrated in Figure 1.

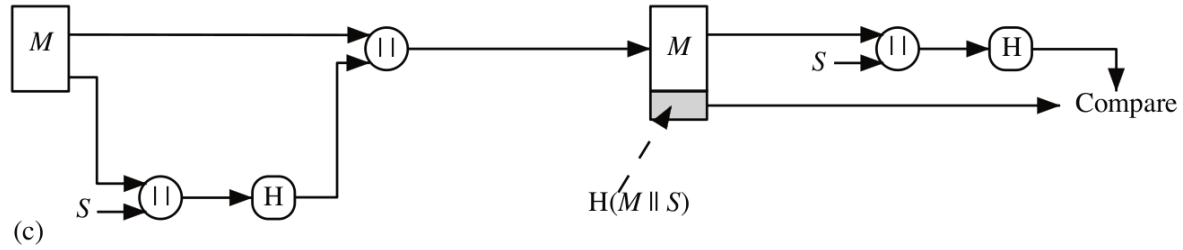

Figure 1: Security mechanism 1

(a) What is a security service that this mechanism provides? [1 mark]

Answer. *Authentication, data integrity*

(b) Explain (or define) the *one-way property* (also called *pre-image resistant property*) of a hash function. [2 marks]

Answer. *Computationally hard to determine the input of a hash function, given only the hash function and the output hash value*

(c) Explain how an attacker can defeat the above security service if the function H() did not have the one-way property. [3 marks]

Answer. *If the one-way property does not hold, then from the hash value,* $\mathrm{H}(M||S)$ *the attacker can find* $M||S$*. Since the attacker also knows* M *they can find* S*, the shared secret. Once they know the secret they could send a message to B, pretending to be A.*

Consider the mechanism illustrated in Figure 2

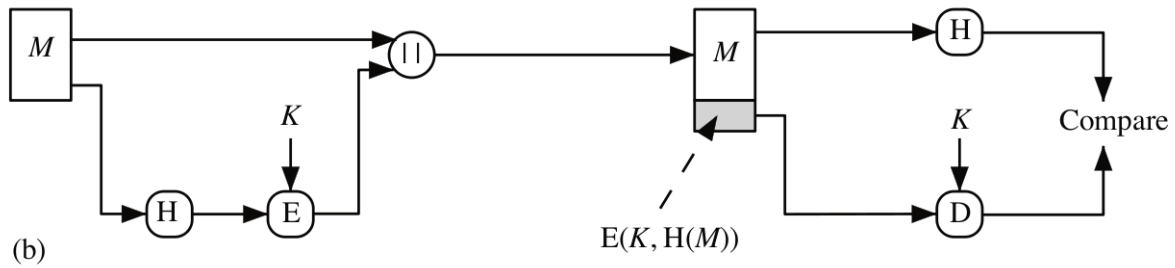

Figure 2: Security mechanism 2

(d) What is a security service that this mechanism provides? [1 mark]

Answer. *Authentication, data integrity*

(e) Explain (or define) the *weak collision resistant property* (also called *second pre-image resistant property*) of a hash function. [2 marks]

Answer. *Computationally hard to find a message y such that $H(x) = H(y)$, given the hash function and x*

(f) Explain how an attacker can defeat the above security service if the function H() did not have the weak collision resistant property. [3 marks]

Answer. *If the weak collision resistant property does not hold, then the attacker can replace the message M sent by A with another message y, where $H(M) = H(y)$, and forward the message to B. B will not detect the change because after decrypting, the received hash value $H(M)$ will match the calculated hash value $H(y)$.*

(g) What is the difference between a hash function and a MAC function? [1 mark]

Answer. *A hash function takes data as input, while a MAC function takes data and a key as input*

Question 4 [13 marks]

A (digital) certificate contains a users public key (PU), identity (ID), a timestamp (T), as well as those three values signed by a trusted entity. Consider a system where the trust relationships between users are shown in a hierarchy as in Figure 3 where a user trusts the user on the next level up in the hierarchy. For example, users A and B trust user X; user X trusts user Z.

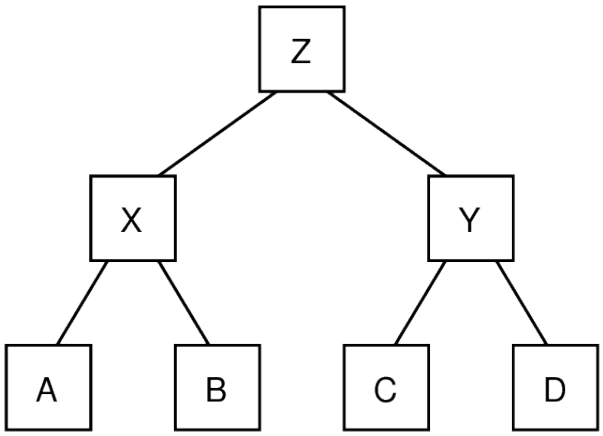

Figure 3: Trust relationship

In the following questions you may use the operators: $||$ for concatenation, $E(k, p)$ and $D(k, c)$ for encryption and decryption, and $H()$ for a hash function.

(a) Write an equation that shows the certificate of user A, i.e. $C_A = \ldots$ [3 marks]

Answer.

$$C_A = ID_A||PU_A||T||E(PR_X, H(ID_A||PU_A||T))$$

(b) Who signs C_Z? [1 mark]

Answer. *C_Z is a self-signed certificate; Z signs its own certificate*

Assume all users except user B already have their own certificate and the certificate of the other user they trust. For example, A has C_A and C_X. Consider user B.

(c) What algorithm can B use to generate its own key pair? [1 mark]

Answer. *RSA*

(d) Which user does B send a Certificate Signing Request to? [1 mark]

Answer. *X*

All users, including B, now have their own certificate and the certificate of the other user they trust. Consider A wanting to communicate confidentially with B.

(e) Draw a diagram that shows the steps that A and B (and others if necessary) take to exchange public keys. The diagram should be similar to those seen in lecture (e.g. like Figure 4 or 5), labelling the messages with numbers to indicate the order and showing the contents of messages. [2 marks]

Answer. *A sends C_A to B and B sends C_B to A*

(f) Explain how A verifies the information it receives from B, including what information A must know to perform the verification. [2 marks]

Answer. *A needs C_X. A uses PU_X from C_X to decrypt the signature of C_B. If the decrypted value matches the hash of the received public key then it is successfully verified.*

Now consider B wanting to communicate confidentially with C. They exchange their public keys.

(g) Explain how C verifies the information it receives from B, including what information C must know to perform the verification and what other communications may need to take place to complete the verification. [3 marks]

Answer. *This question had an error in it—it referred to A instead of C. Therefore I gave everyone full marks for this question.*

Question 5 [6 marks]

The following are a selection of Linux commands used for cryptographic operations. Some commands have selected parts hidden with XXX.

(a) `openssl pkeyutl -verify -in file.txt -sigfile sign.bin -certin -inkey cert.pem`

(b) `openssl rand 16 -hex`

(c) `openssl req -new -key privkey.pem -out req.csr`

(d) `openssl pkeyutl -encrypt -in file.txt -certin -inkey cert-XXX.pem -out file.bin`

(e) `openssl genpkey -algorithm XXX -out privkey.pem`

(f) `openssl pkeyutl -sign -in file.txt -inkey privkey.pem -out sign.bin`

(g) `openssl verify -CAfile cert-CA.pem cert.pem`

Select the most appropriate command from above that is used to perform each of the following operations. To answer, in the space for each operation, give the letter, from between *a* and *g*, of the command. [1.5 marks each]

(a) Create a request message to be sent to a CA so that CA can generate a X.509 certificate *c*

(b) Create a public and private key pair *e*

(c) Check that a received file has not been modified and came from a particular person *a*

(d) Create a secret key that can be used for encryption *b*

Question 6 [6 marks]

(a) Steve wants to send a MAC authenticated message M to Thanaruk. Give an equation that describes what Steve sends to Thanaruk, i.e. $Sent = \ldots$. You must also describe all variables used. [2 marks]

Answer. *$Sent = M||MAC(S, M)$ where S is a shared secret key with Thanaruk.*

(b) Assume a malicious user, Pakinee, intercepts the message sent by Steve. Pakinee modifies the message M. Can Thanaruk detect this modification? Explain your answer (i.e. what Thanaruk does to detect or why Thanaruk cannot detect). [2 marks]

Answer. *Yes, Thanaruk can detect the modification. If M is modified to M' then upon reception Thanaruk uses the shared secret key S to calculate the MAC of the receive message, M', i.e. $MAC(S, M')$. Thanaruk finds $MAC(S, M') \neq MAC(S, M)$ and hence has detected a modification.*

(c) Explain why MAC-based authentication cannot be used as a digital signature. [2 marks]

Answer. *A MAC function uses a shared secret key. A message authenticated with a MAC function confirms that the message was generated by either the of the parties that has the secret. It does not confirm which of the two parties generated the message (which is the purpose of a digital signature).*

Question 7 [10 marks]

Consider a system with 100 users. Confidentiality of communications between users must be provided using AES. Any pair of users may potentially communicate. Figure 4 shows the key distribution protocol used in the system, illustrating the steps that two example users, A and B, use.

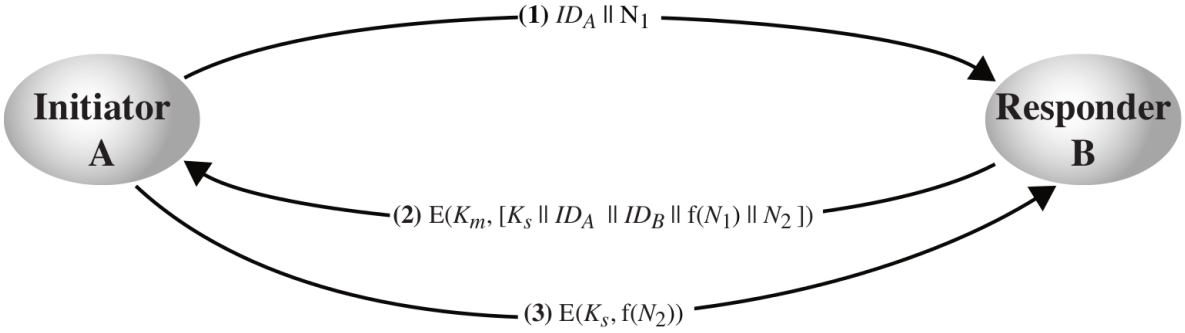

Figure 4: Key distribution protocol

(a) Prior to the key distribution protocol in Figure 4 being used, master keys must be exchanged manually. How many masters keys must be manually exchanged in the entire system to allow any pair of users to communicate? [2 marks]

Answer. *Each user must manually exchange a master key with 99 other users. Therefore a total of* $\frac{100 \times 99}{2}$ *or 4,950 keys must be exchanged*

(b) What is the benefit of performing the key distribution in the figure, as opposed to just using the manually exchanged master keys? [3 marks]

Answer. *This allows the session key* K_s *to change on a regular basis, thereby not re-using keys for too long. If only the master key was used for data encryption, then either the key would be used all the time (insecure) or manually exchange would need to be repeatedly performed (inconvenient).*

(c) Explain one way in which N_1 may be chosen or generated. [1 mark]

Answer. *Random number, timestamp, counter*

(d) If after the steps in Figure 4 are successfully completed, user A has data D to send confidentially to B, write a statement showing what A sends to B. Use the same notation as in the figure. [2 marks]

Answer. $E(K_s, D)$

(e) Explain an advantage of the protocol in Figure 4 compared to using a Key Distribution Centre (KDC)? [1 mark]

Answer. *No need to trust KDC, no performance bottleneck at KDC*

(f) Explain a disadvantage of the protocol in Figure 4 compared to using a Key Distribution Centre (KDC)? [1 mark]

Answer. *Fewer keys to be manually distributed before the protocol operation when using KDC.*

Question 8 [12 marks]

Consider the key distribution protocol in Figure 5.

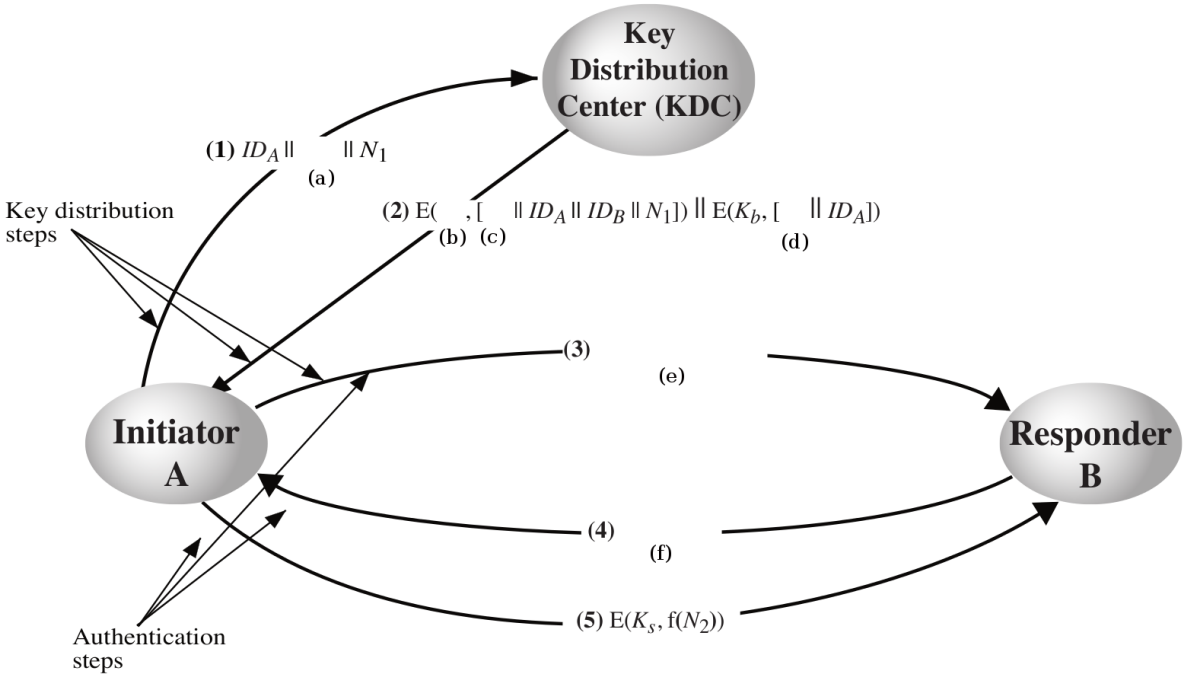

Figure 5: Key distribution protocol

There are six missing values in Figure 5, labelled (a) through to (f). Write the missing values below. [1.5 marks each]

(a)

Answer. ID_B

(b)

Answer. K_a

(c)

Answer. K_s

(d)

Answer. K_s

(e)

Answer. $E(K_B, K_s||ID_A)$

(f)

Answer. $E(K_s, N_2)$

Consider the keys known before any of the steps are performed in Figure 5. What keys are known by: [1 mark each]

(g) KDC:

Answer. K_A, K_b

(h) A:

Answer. K_A

Consider the keys known after all of the steps are performed in Figure 5. What keys are known by: [1 mark]

(i) A:

Answer. K_a, K_s

Question 9 [7 marks]

Consider a computer system that uses passwords for authentication.

(a) Assuming the developers of the system (i.e. insiders) can be trusted, explain why passwords should not be stored in the clear on the system. [2 marks]

Answer. *If an outside gains access to the system then passwords of all users are immediately available.*

(b) Explain the recommended method for storing password information on the system (i.e. what should be stored). [3 marks]

Answer. *Store the username, salt, and hash of the salted password, such as* $H(P||S)$

(c) Assuming you do not limit how users select passwords, give two countermeasures for online password guessing by an attacker. [2 marks]

Answer. *Limit the number of attempts; introduce a delay between attempts.*

Question 10 [9 marks]

Consider the internet in Figure 6. On the 4 subnets assume there are many hosts (although only two hosts are shown for each subnet due to space). The host IP addresses are obtained from the subnet address and the host number, e.g. host 2 has IP 1.1.1.2. The two routers have three interfaces.

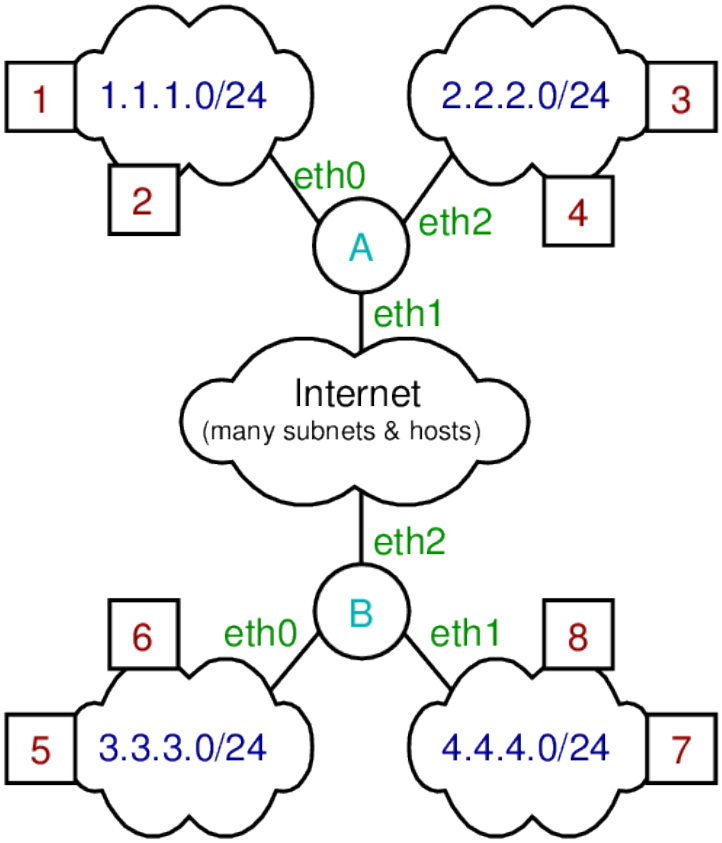

Figure 6: Firewall network

You are the IT administrator for the two subnets attached to router A and need to add a rule to the firewall running on router A. The default policy for the firewall is accept. Stateful Packet Inspection is enabled on the firewall.

For each of the following policies, write a rule that implements it. The format/syntax of your rule is not important, as long as it is clear as to what conditions must be met and what action to take. (For each part, assume initially there are no firewall rules; i.e. your answer in part (b) is independent of your answer in part (a)). [3 marks each]

(a) Block all hosts on network 1.1.1.0/24 from accessing the web server on host 8

Answer. *Source IP is 1.1.1.0/24, destination IP is 4.4.4.8, protocol is TCP, destination port is 80, action is drop*

(b) Block host 7 from communicating with any internal hosts.

Answer. *Source IP is 4.4.4.7, action is drop*

Consider the same network as in Figure 6. Below is a single packet that is received by the firewall, and the current firewall table and SPI table.

```
+----+-----+------+
| IP | TCP | Data |
+----+-----+------+
  |     |
  |     |-SrcPort=25
  |     |-DstPort=47231
  |
  |- SrcIP=1.1.1.1
  |- DstIP=3.3.3.5
```

Firewall Table

Rule	SrcIP	SrcPort	DstIP	DstPort	Protocol	Action
1	3.3.3.0/24	ANY	1.1.1.1	25	TCP	ACCEPT
2	3.3.3.6	ANY	1.1.1.2	ANY	TCP	ACCEPT
3	4.4.4.0/24	ANY	2.2.2.4	23	TCP	ACCEPT
4	2.2.2.0/24	ANY	4.4.4.7	23	TCP	ACCEPT
5	2.2.2.3	ANY	4.4.4.7	23	TCP	ACCEPT

Default policy: DROP

Figure 7: Firewall Table

Stateful Packet Inspection Table

SPI	SrcIP	SrcPort	DstIP	DstPort
1	3.3.3.6	40327	1.1.1.1	25
2	4.4.4.7	47231	2.2.2.4	23
3	3.3.3.5	47231	1.1.1.1	25
4	2.2.2.3	40327	4.4.4.7	23
5	2.2.2.3	44981	4.4.4.7	23
6	4.4.4.7	40327	2.2.2.4	23
7	2.2.2.4	40327	4.4.4.7	23
8	2.2.2.3	44981	4.4.4.7	23

All entries for TCP

Figure 8: SPI Table

(c) What action does the firewall take for the received packet? Explain why it takes that action (i.e. refer to the (non-)matching table entries in your answer). [3 marks]

Answer. *The packet matches the 3rd row in the SPI table. Note that the SPI table checks packets in either direction.*