CHALMERS | GÖTEBORGS UNIVERSITET

Network Security

EDA491 & DIT071 7,5 hec
EDA490 & DIT070 6,0 hec

Monday 2009-12-14 14:00–18:00 in building M

No extra material is allowed during the exam except for an English language dictionary in paper form. **No electronic devices allowed.**

Give clear answers. Your thoughts and ways of reasoning must be clearly understood! Questions must be answered in English.

Teacher:	Tomas Olovsson Dept. of Computer Science and Engineering
Questions during exam:	Tomas Olovsson, 772 1688
Answers:	Published on the web page after the exam
Results:	Sent via email from Ladok
Inspection of exam:	See web page for announcement

NOTE: **Indicate clearly on front page if you take an older (6 hec) course!**

CTH Grades:	30-38 → 3	39-47 → 4	48-60 → 5
GU Grades:	30-47 → G		48-60 → VG
Older 6 hec courses:			
CTH Grades:	24-35 → 3	36-47 → 4	48-60 → 5
GU Grades:	24-47 → G		48-60 → VG

1. Attacks, DoS

a) There are several ways to scan a system such as SYN, ACK, and FIN scans. For each of these scans, explain clearly the advantage (or possibly disadvantage) with it and what is *gained* with the scan. It should be obvious from your answer that there is a reason why scanning tools offer different ways to scan systems. (3p)

SYN scan: shows directly what ports are open and accepts communication (SYN/ACK received) and what ports are closed (RST or no answer).
ACK scan: may work through stateless firewalls: RST means someone (a host) is there.
FIN scan: More stealthy than the other but may not work that often (not against Windows for example).

b) In the Ping of Death attack, the attacker creates an oversized ICMP echo message. Explain what the problem with receiving this is, and how it is really possible to create an oversized datagram (give an example)! (3p)

A naïve implementation would assume IP datagrams never exceed 65,535 bytes since the length field is 16 bits long.
An oversized IP datagram can be created that exceeds this size by sending a fragment with an offset and a length extending the datagram beyond this limit, for example by setting offset to 65,000 and length to 1,000 bytes.

c) IP datagrams with overlapping fragments can be problematic, for example for IDS systems and firewalls. Why? Describe the problem by giving an example! (2p)

If datagram 1 = AAAAAA and datagram 2 = BBBBBB and they overlap by 50%, a firewall or IDS system does not know how the receiving system reassembles the datagram. It may become AAABBBBBB or AAAAAABBB.

d) A possible DoS attack is to send lots of SYN packets to a victim. Why can this be a problem for a system? (2p)

e) The allocation of SYN cookies is a possible defense against this type of attack. A SYN cookie may be stored in the initial serial number (ISN) and be:

ISN = hash (source and dest IP and port numbers + secret + client's ISN)

Explain clearly why this may help the situation! Also explain when and in what situations it offers protection? Are there any drawbacks or problems with this method (if not, why is it not the default mechanism)? (4p)

Why: The server does not save any data when a SYN is received. Instead all state information is contained in the ISN which means that it does not have to allocate any resources until it receives the ACK. Then the ISN is checked with what the client sends back.
When: The attacker must be using a valid IP address to receive the SYN/ACK with the ISN to be able to respond with a valid ACK. Only if the ACK is valid, resources are allocated. If one or a few addresses still flood the server, a firewall could easily filter or block this IP address, manually or automatically.

Drawbacks: problems to remember if TCP options (such as window size) were negotiated in the SYN and SYN/ACK packets since they are not saved by the server. Another drawback may be that the computation of the hash may take some time, a fact that may be utilized by the attacker.

2. Firewalls

a) Give six examples of different types of rules that a screening router likely would have to filter traffic! You may use text or your own "pseudo-language" for the filter rules as long as they are possible to understand (comments to the rules are necessary). (3p)

See slides. Examples could be private (reserved RFC 1918) addresses:
(block inout 127.0.0.1), incoming traffic to broadcast addresses (block in 1.2.3.255), outgoing ICMP packets such as ICMP echo reply, clearly erroneous incoming packets (SYN=1, FIN=1), outgoing packets not belonging to our IP address range, etc.

b) Why is it problematic to send protocols like FTP through non application-level firewalls? What is the fundamental problem with this protocol? Also give one example of how a firewall working on transport and network level (such as a router) could handle such protocols! Use your own "pseudo-language" again if needed. (3p)

c) Explain briefly how NAT can be used as a firewall feature. Do you think it is a reasonably good and secure enough for a home network? Explain! (2p)

d) Stateful firewalls maintain both a state table (connection table) and a table of rules. What is the difference between these tables, why are two needed? Which table is consulted first when an incoming packet is received? Show with an example what happens when some packets go through the firewall and how these two tables are used and what they may contain! (4p)

State table is used for established connections.
The ACL contains the rules for the firewall and is consulted when new connections are estabished.

Example:

ACL:

- Pass in from any to 1.2.3.4 port 80 # web traffic
- Block in all

When an incoming SYN packet is received from host 22.33.44.55 to port 80 on 1.2.3.4, the ACL list is consulted and a state table entry is created:

```
  22.33.44.55  999   1.2.3.4   80   SYN-recvd
```

When an SYN/ACK is seen (from 1.2.3.4), the State table is directly consulted and since it contains an entry for this connection, the packet is passed on. It is now updated to:

```
  22.33.44.55  999   1.2.3.4   80   SYN/ACK-sent
```

3. SSL and VPN systems

In SSL, the client and the server negotiate a master secret. This master secret ("master") is then used to create key material for the current session:

key_material = MD5(master || SHA('A' || master || r1 || r2))
|| MD5(master || SHA('BB' || master || r1 || r2))
|| MD5(master || SHA('CCC' || master || r1 || r2))
|| …

a) Six keys are normally created when a new session is initiated. Which are these keys and what are they used for? (2p)

The six keys are Client and Server write keys, Client and server MAC keys, Client and Server IV. Write keys are used to encrypt messages (confidentiality), MAC key to create the HMAC (message integrity), and the IV is used if the encryption algorithm requires an IV. (The client keys are used in one direction, the server keys in the other.)

b) Why six keys? Why not just three or even just one key such as using the master secret as the key? It is the only unknown parameter in the generation of the key_material above anyway, so why bother creating more keys by hashing it a couple of times? (3p)

The reason why six keys are used and not three, is that to use different keys in each direction in order to limit the damage if a key is compromised.

We definitively don't want to use the master secret directly since the more we use a key (or master secret), the more information is given to potential attackers and it makes it easier to break it. It is therefore used sparsely and just to create new keys that are regularly changed.

c) At the end of the SSL handshake protocol, both sides sends a "Finished" message to the other side:

hash(master || *opad* || *hash(msgs* || *0x53525652* || *master* | *ipad))*

Where *msgs* here is a list of all exchanged messages. This message was not present in the early versions of the protocol. What type of attack does this protect against? (2p)

d) Suppose you are the person in charge of selecting a VPN system to be used for remote access for a medium size company (500 employees) with a mixed client environment (Windows, MAC, …). You have to choose between an SSL (or SSH) based VPN system and a system based on IPsec. Your manager will purchase the system you recommend. What would you chose? What would your report to the manager say? What are the strong and weak points with the solutions? (5p)

(There may be several answers to this question, but you need to come with valid arguments and show that you have understood the technology and therefore can make a comparison.)

You are allowed to select either type for full score (both types exist on the market and are used in such environments even if SSL-based solutions are most popular). You will get one point for each valid argument. Typical things to mention are:
- NAT, application transparency, administrative privileges needed, simplicity of web based applications and applets, multi-user system support, ease to deploy on different client platforms, ease of administration, ease to traverse firewalls, …

4. Authentication, link level security

a) Identity management (IDM) is a relatively new field. What is the fundamental problem IDM tries to solve? (2p)

Central coordination of user accounts and user identities in an organization.

b) SAML (Security Assertion Markup Language) can be a part of an identity management system. What is its purpose? What problem does it address? (2p)

SAML is a language developed to allow one server (or organization) to tell another that a user is authenticated and authorized for a certain service, something which is useful if they don't have common account databases. Instead a message is sent between two parties that trust each other telling the other about the user.

c) A self learning switch learns where hosts are located and will normally not forward private traffic between two ports to any other ports. How can an attacker see this supposedly private traffic between two other communicating parties? What could an administrator do to decrease this problem? (3p)

d) Assume you are the systems administrator at a school and you need to create a link level solution using only switches. There are three networks available: a student network, a teacher network and a security lab connected as shown in the picture. There are some security requirements you need to consider:

- The traffic between the networks should never be mixed (students should not be able to access the teachers' computers or security lab computers. Same for teachers; no access to student computers or the security lab.)
- The traffic from the lab should never leave the lab network except to a printer located in the students' network (students in the lab should be able to print to this printer). This printer should only be available for people in the security lab.
- Both teachers and students should have Internet access but not the security lab.
- You are the trusted administrator and only you can access and configure switches.

Assume there are around 100 computers on each network. You are not allowed to work with routers. The users are connected to the large switches; the smaller switches are used to connect networks and other devices. You can also assume that the switches have the necessary functionality. How would you implement your solution? Your answer should be instructions for the person who configures the devices – he/she can read their documentation but does not know much about security. (5p)

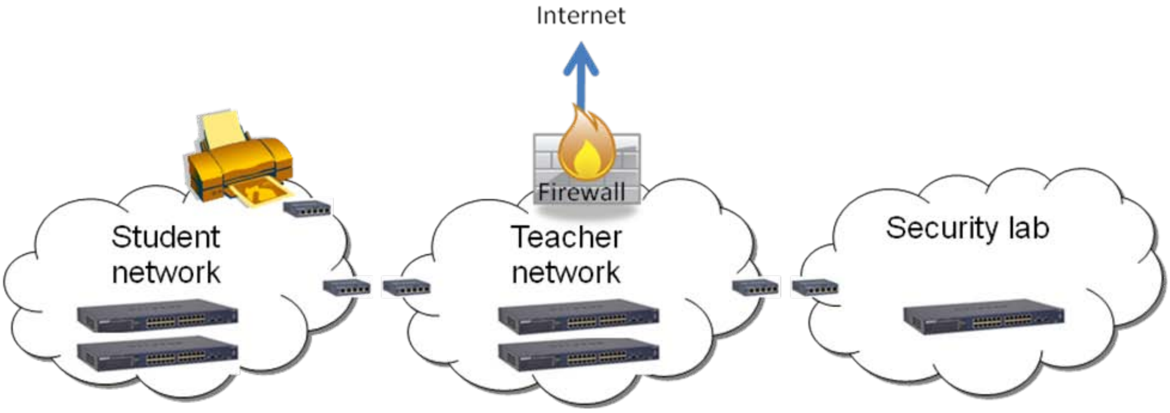

The technique to base the solution is VLAN separation where the switches set the tags to make sure that traffic passing through the teacher network is kept isolated. Explain where tags are set, what VLANs are assigned, etc.

5. Misc short questions

Answer the following questions with one or maximum two sentences, not more. The answer must be detailed enough to show that you understand the topic!

a) WEP is not known for its good security. Describe one vulnerability or possible attack against it! (2p)

b) 802.1x implements port-based network access control. Describe with one sentence what this is! (2p)

c) The Kerberos server is stateless which means that it does not keep track of the users it has authenticated. It is still possible for a user to come back to it and request tickets to services without repeating the authentication process. How is this possible? (2p)

d) In the paper "*A security analysis of Windows Vista's network implementation*" Symantec discusses several weaknesses that they have found in the beta version of Windows Vista. Describe briefly one problem found! (2p)

e) Bonus for answering course evaluation (see web page for details, note the deadline) (2p)