TE WHARE WĀNANGA O TE ŪPOKO O TE IKA A MĀUI
VICTORIA
UNIVERSITY OF WELLINGTON

EXAMINATIONS – 2015

TRIMESTER 2

NWEN 405

Security Engineering

Time Allowed: THREE HOURS

CLOSED BOOK

Permitted materials: No calculators permitted.

Non-electronic Foreign language dictionaries are allowed.

Instructions: There are 180 possible marks on the exam.

Attempt all questions.

Many of the questions require you to discuss an issue, compare options or evaluate a situation. For such questions, the assessment will take into account the evidence you present and any insight you demonstrate.

If additional space is required you may use a separate answer booklet.

Question	Topic	Marks
1.	Security Policies and Models	30
2.	Software Security	30
3.	Malware and Countermeasures	30
4.	Network Security Attacks	30
5.	Firewalls and Intrusion Detection Systems	30
6.	The Bigger Picture	30
	Total	180

Question 1. Security Policies and Models [30 marks]

(a) [10 marks] Summarize the historical reasons leading to the development of the concepts of a Trusted Computing Base and the Bell-LaPadula model of computer security.

Good answer should recognise: need to use commercial time sharing machines, discovery that DAC based OS policies could easily be bypassed using trojan horses. Excellent answer: provides definitions of the terms, need to have a way to formalise information flow.

(b) [10 marks] You have been employed to fix a problem with a secret government airline's travel booking system.

The system implements multilevel security and uses *polyinstantiation* to avoid the identify of High users being revealed to Low users. However, this is leading to situations where users find their assigned seat has been double booked by another user.

Explain the cause of this double booking, propose and evaluate *TWO* alternative solutions.

Good answer – explains that a low user cannot see a high user's booking to avoid a covert channel constructed by using seat bookings, suggests either autmatically reply classified to low user (uk) or high users must use a cover identity when booking a flight (usa).

Excellent answer – recognises booking a cover identity (usa) is harder to implement and still opens up a covert channel whereas uk also has covert channel but has the advantage of simplicity.

(c) [10 marks] Describe the implementation of a Chinese Wall security model using an operating system that allows access control policies to be expressed using Role Based Access Control.

Good answer would suggest that users would organised into roles based upon which client they represented, access permission to client files would associated with the role associated with representing that client. This ensures separation of duty. Excellent answer includes discussion of the need to dynamically associate a user with a given role using the concept of a session

Question 2. Software Security [30 marks]

(a) [10 marks] A major challenge in crafting a successful stack-based buffer overflow attack is the address space of a given process is unpredictable making this process difficult and subject to guesswork. Compare and contrast the following buffer overflow attacks with respect to the techniques developed by attackers to overcome these challenges.

(i) NOP sledding.

Aim is increase the chances of guessing the right return address by increasing the size of the target by including null instructions that if executed simply cause control to fall through to the malicious program.

(ii) Trampolining.

Reduces the uncertainty compared to NOP sledding. Unlike NOP sledding we don't directly use the guessed address of the malcious code, instad Overwrite the return address of current function with the location of code contained within external libraries that are always loaded into predictable locations in memory. The code targeted for execution are usually instructions that branch to a memory location stored in a register. The aim is to load the malware at the location in the register. Increased liklihood of success because of increased predictability.

(iii) Return-to-libc.

Similar to trampolining in that we use an extenal library as our target – in this case libc. Idea is to determine address of a function that will allow execution of arbitrary instructions such as system() and to overwrite the return address with this along with the arguments for the function that contain the shellcode.

Unlike NOP and trampolining in that we do not directly try to execute code on the stack itself.

(b) [10 marks] Explain why programmers should still follow safe coding practices depspite almost all recent operating systems and compilers have features to reduce the risks of buffer overflows. Provide relevant supporting reasons and examples.

Many examples of weaknesses in protection mechanisms at the operating system and compiler level. Should give at least one example of an operating system protection and one for a compiler level protection. Stack canaries and similar approaches can be defeated by attacks that provide the expected canary value at runtime while still overwriting the return value. ASLR implementations have used insufficient randomness. DEP/NX is sometimes switched off or the attackers change tactics – for example malloc/heap/libc attacks.

Attackers will use these weaknesses to bypass the protection mechanisms and exploit the vulnerabilities despite the presence of the these protection mechanisms so the only way to guarantee safety to remove them entirely.

(c) Consider operating system architecture.

(i) [5 marks] Outline the historical reasons for the reduction in protection levels over the last 50 years.

Enforcement of protection rings rely upon CPU support for different modes of operation. Except for Apple, most operating systems are designed to run on many kinds of hardware. Means that has to be designed for lowest common denominator in terms of hardware support for multiple states and this currently means two states.

(ii) [5 marks] Explain why the reduction in protection levels combined with a trend to a more modular architecture led to less secure operationg systems?

Third party code – risky, excutes with high privilege. Means can affect the whole system.

Question 3. Malware [30 marks]

(a) [10 marks] Describe how polymorphic and metamorphic viruses differ with respect to concealment and mutation strategies.

Concealment strategy – polymorphic virus behaviour always stays the same where a metamorphic virus changes it behavour with every infection. This means that metamorphic viruses can evade detection by behaviour-based anti-virus tools.

Mutation strategies – The method of mutation also is different between the two types of virus. Polymorphic viruses use encryption to hide the signature of the virus, the virus has an encrypted body that is decrypted during execution. The actual plaintext virus code that eventually executes is static (the decryptor itself may be static or metamorphic). By contrast, a metamorphic virus, changes the particular instructions used by the virus each time it spreads.

(b) Consider the use of Tumblr for botnet command-and-control. Assume a simplified version of Tumblr that works as follows: (1) users register accounts, which requires solving CAPTCHAs; (2) once registered, users can create text and image blog posts; (3) user A can subscribe to user B so that A receives copies of Bs blog posts; (4) user B can tell when user A has decided to follow user B.

(i) [6 marks] Sketch how a botmaster could make use of Tumblr to create a centralized C&C structure. Clearly identify what actions the different parties (individual bots, botmaster) take. You should consider at least two different approaches that differ in terms of the communication pattern between the botmaster and individual bots. Assume that there is no worry of defensive countermeasures.

A simple approach is that the botmaster registers two Tumblr accounts, A and B. A is used to send commands, and B for receiving commands. The bot malware includes within it the credentials for the B account. New bots then access the B account to read the tweets sent by the A account, which encode the instructions to the bots.

An alternative approach would be to create a new account per bot. This requires some discussion of how to solve all the CAPTCHAs; for example, by purchasing solutions from a solving service available from the underground economy.

(ii) [4 marks] Briefly describe a method that Tumblr could use to detect botnets using this C&C scheme.

For the first scheme, Tumblr could look for access to the same account from many different IP addresses.

For the second scheme, Tumblr could look for accounts whose followers all only follow that account.

(c) Evaluate the effectiveness of filter-based and rate halting worm containment mechanisms with respect to the following scanning techniques.

(i) [4 marks] Random.

Each worm probes random address in the IP address space, produces a high volume of scanning traffic.

Filter-based: worm will be detected at the end-hosts and traffic from the worm processes is filtered out. Rate halting: all traffic from an infected host is blocked once a threshold is exceeded.

(ii) [4 marks] Hit list.

Attacker performs slow scans for vulnerable machines over a long period and seeds the worms with these lists and resulting in a very short scanning period.

Filter-based technique will detect the presence of the worm based upon monitoring activity on the end-host whereas rate halting will fail to detect the worm because the volume of scans would be too low be suspicious.

(iii) [2 marks] Local subnet.

Host looks for targets in its own network using the subnet address structure.

Neither technique can prevent the spreading of the worm because filtering is applied at the network perimeter rather than within the local area network. Worms will be able to infect other vulnerable hosts on the same subnet.

Question 4. Network Security Attacks [30 marks]

(a) Network vulnerabilities and erronerous assumptions.

(i) [5 marks] Outline the *THREE* main vulnerabilities in the ARP protocol and show how these can be exploited by an attacker to conduct a successful cache poisoning attack.

(1) no authentication of requests or replies allowing ARP requests and replies to be forged; (2) ARP is stateless so forged replies will not be rejected because there is no marching request; (3) ARP requests or replies are blindly applied by nodes so a forged reply will be accepted;

(ii) [5 marks] Outline *THREE* main assumptions that were incorrectly assumed would prevent TCP session hijacking.

(1) IP addresses cannot be easily spoofed; (2) Even if spoofed cannot easily take over existing one because a valid packet must include the sequence number expected by the received and because sequence numbers are random it is unlikely that they could be correctly guessed.

(b) Consider Kaminsky's DNS cache poisoning attack.

(i) [5 marks] Outline the main steps involved in carrying out the attack.

Attacl: (1) Force target nameserver to initiate query to authoritative server of his choice. (2) Generate many multiple requests and replies (assuming enough bandwidth). (3) Place the malicious mapping in the answer section (as in the old pre-balwick attack).

(ii) [5 marks] Describe and briefly evaluate the weaknesses of *THREE* proposed countermeasures against this attack.

Countermeasures: (1) Add source port randomisation and distinguish requests using combination of transaction ID and source port. Increases work that needs to be done by the attacker. (2) Do not allow untrusted clients to ask a nameserver to do a recursive lookup. What about malware in your organisation? (3) Adopt DNSSec so you can authenticate who you are communicating with. How widespread is DNSSec?

(c) Denial-of-service attacks.

(i) [5 marks] Why are Reflection attacks considered easier to deploy than traditional DDoS attacks?

DDoS use compromised systems under control of the attacker whereas Reflection attacks use network systems functioning normally. Essentially the hosts involved in the reflection attacks are free as opposed to the compromised systems which must be either obtained by the attacker directly or bought on the black market.

(ii) [5 marks] Imagine that you wish to launch an Amplification attack but don't want to target DNS servers. What criteria would you use to choose between an alternative set of servers to target?

Uses UDP protocol instead of TCP, this allows the use of a spoofed IP address in the request to cause the reply to sent to the target system. The size of the reply packet must be larger than the request. Must accept request from anyone. Ideally many of them and they are are have good Internet connectivity. Should be high capactity hosts to avoid intermediate systems being overwhelmed.

Question 5. Firewalls and Intrusion Detection Systems [30 marks]

(a) Imagine that you want to enforce a firewall security policy that allows any inside host to send mail to any outside host (port 25).

(i) [5 marks] How would you express this security policy for a packet filter firewall? What are the limits of the ability of the packet filter firewall to enforce this policy.

Default deny all. Allow traffic to any external machine from any internal machine to port 25. Only allow traffic coming back in if it has an ACK flag set and target is port 25. Problem is that this doesn't capture the relationship between a previous SYN packet sent from the network, attacker can just forge any packet with ACK and have it accepted

(ii) [5 marks] How would using a stateful inspection firewall overcome the limitations of a packet filter firewall for enforcing the security policy?

Can specify that only accept incoming packets for a TCP connection initiated from an internal host. Active outgoing connections are recorded in a connection table. Only incoming accept packets that match one of the existing connections

(iii) [5 marks] Explain why packet filters typically reject any packets with IP fragments containing less than a certain minimum amount of the transport header.

Attackers deliberately fragment the packets so that the first fragment does not contain enough TCP header information for a packet filtering rule to be applied. They hope in thts case the the first packet will be allowed the firewall and subsequent packets are also passed through to be reassembled on the other side.

(iv) [5 marks] Compare the packet filter approach and stateful inspection filter approaches to dealing with IP fragmentation with respect to number of packets falsely discarded and resource usage.

Stateful implements TCP/IP stack and reassembles packets whereas the packet filter will just reject if not enough TCP header information stored in IP fragment. Can lead to large number of packets discarded that are fine if packet filter but stateful uses more resources to implement the reassembly.

(b) Consider Intrusion detection systems (IDS).

(i) [5 marks] What metrics are used to assess the performance of a detection system and what is considered the ideal tradeoff? Justify your answer.

False positives, false negatives. Want to have high detection rate – ratio of detected to total attacks while minimising the number of false alarm. A low false alarm rate is important for reducing operator fatigue which can lead to real attacks being missed.

(ii) [5 marks] Discuss the fundamental reasons for a tradeoff between how tightly intruder behaviour is defined and the performance of a detection system.

Too tight a profile and you miss some attacks. Too loose and you diagnose some non-attacks as attacks. The problem is that there is usually an overlap between the profile of an attacker and a normal user.

(iii) [5 marks] What are the main differences between inline and passive Network Intrusion Detection System sensors? Explain the why these differences would lead you to perhaps choose one type for a small organisation and another for a large organisation?

Inline sensor can be implemented in an existing network device so no extra hardware is required compared to a passive sensor that requires special hardware. This would be attractive for small businesses wanting to save money. On the other hand passive sensors monitors a copy of the network traffic and so avoids introducing packet delay into the system. This would advantageous when monitoring backbone high speed links.

Question 6. Individuals, Society and the Physical World [30 marks]

(a) Consider a scenario where you receive an unsolicited phone call at work from a person claiming to work for Visa. They inform you that they are checking some odd activity on your account that is happening right now. They quote your cedit card number to you and ask if you could read the three security digits on the signature strip to prove you still are in possession of the card.

(i) [5 marks] Explain why this is an example of pretexting attack.

Has all the main features – attempt to gain private information, attacker claims to be someone authorised to have access and create a sense of false urgency. Answer should relate the example to these general features

(ii) [5 marks] What do the results of Milgram's experiments suggest about the probability of someone falling for this attack?

Miligram investigated this question. He wanted to see if people would obey authority than their own conscience. He found the majority would (in his case administer electric shocks) and this suggests that many people might comply with the authority figure in this case.

(b) You have been employed by a shadowy organisation to craft a targeted attack on a Government minister. After a review of his Facebook profile and Trademe account you discover the following about him: (1) he greatly admires Donald Trump; (2) he is currently in a bidding war for a ticket to see Justin Beiber perform an acoustic set in Auckland.

Consider for each of the following how you could use an understanding of psychology to maximise the chances that he will click on a link provided by you in an email.

(i) [5 marks] Affect heuristic.

Idea of affect heuristic is to get people to think with heart rather than head. Could ask a personal question but a more common approach is to send them an email featuring or being sent you one of their heroes – in this case Donald Trump.

(ii) [5 marks] Prospect theory.

Prospect theory: models risk adversion, people are likely to take an action framed as a gain rather than a loss. Target their current bidding war for Justin and send email saying that unless the follow the link they are likely to lose the chance of gaining a ticket.

(c) [10 marks] You have recently moved your datacentre into an area of the city where many other datacentres are also located. You are worried about a recent spate of robberies where attackers have broken into datacentres to steal hardware. Your neighbouring datacentres pay for a very expensive but effective security service featuring silent invisible alarms and a security guards who respond to the alarms.

Evaluate whether you need to also join this security service. When answering this question you should consider the lessons from the Lojack study.

Want them to spell out similarities in terms of invisible protection. How the response rate created a deterrence for attackers or removed them from the area. How it does depend upon effectiveness, without it wont work. Also important that attackers cannot tell that it is NOT being used. Also would like them to talk about freerider effect.

* * * * * * * * * * * * * * *