

Solutions

Each answer shows the steps, then the **result**.

2.1

- **acceptable use policy** - the range of activities permissible, prohibited or required by users on the organisation's devices; **password policy** - the requirements for user passwords; **software installation policy** - what software, if any, users may install, and the process for requesting it

2.2

- any three of: a minimum or maximum password length; a minimum or maximum time a password may be kept; a prohibition of password reuse; construction rules such as no dictionary words or character-set requirements; a suggestion to use a secure password manager rather than writing passwords down

2.3

- users installing software freely can install malware disguised as a legitimate tool, or software with vulnerabilities the organisation has not assessed; the policy therefore provides a **process for requesting** software so it can be checked and installed by someone authorised

2.4

- it holds a database of malware **signatures** and periodically scans the files on the device, checking whether any file matches one of those signatures

2.5

- it **quarantines** the malware and **removes** it

2.6

- detection depends on a matching signature already being in the database; brand-new malware has no signature yet, so nothing matches and it is not recognised - and the same applies to fileless malware, which is not a file to scan at all

2.7

- publishing a fix also publishes the existence of the flaw, so adversaries know exactly what to exploit and exploit code often follows quickly - the vulnerability becomes easy to use and the likelihood of attack rises sharply while the device remains unpatched

2.8

- any two of: it protects a single device even on a **compromised network** such as public Wi-Fi; it can block **outbound** traffic, stopping a compromised device calling out to a command-and-control server; it can block ports and services not needed for that specific device's role, which a shared network rule set cannot do per device

2.9

(a)

- **no** - it is allowed. Rule 1 matches all outbound traffic first, and the first matching rule executes, so rule 2 is never reached

(b)

- put the specific deny above the general allow: 1 DENY outbound any -> port 4444
2 ALLOW outbound any -> any

3.1 B

3.2 B

- It can block outbound traffic, and it does use an ACL

3.3 B

3.4 B

3.5

(a)

- any four, each with a compromise: **no software installation policy** - staff install malware disguised as a font or plug-in, giving an adversary control of the device; **a signature database fourteen months stale** - any malware released since then is unrecognised, so infection is undetected; **dismissed operating-system updates** - publicly documented vulnerabilities remain exploitable, allowing an adversary to take over the device with little skill; **no host-based firewall** - nothing filters traffic on an untrusted cafe network, and a compromised laptop can call out to a command-and-control server unimpeded

(b)

- the office firewall protects traffic crossing the office boundary, and these laptops spend their working time on client and cafe networks, where that firewall is not in the path at all. A host-based firewall is exactly the extra layer for a host connected to a network that may already be compromised - and it also blocks outbound traffic, which the office firewall cannot do once the laptop is elsewhere

(c)

- the malware was first seen two months ago, so its signature exists, but the studio's database is fourteen months old and therefore does not contain it - nothing matched during the scan. After an update the signature would be present, so a scan would then detect and quarantine it - though only after the infection, which is why patching and prevention matter more than detection alone

(d)

- three sensible rules with a defensible order, for example: 1 ALLOW outbound any -> ports 80, 443 (web) 2 DENY inbound any -> any 3 DENY outbound any -> any
Justification: rules are checked in order and the first match executes, so the specific

allows must come above the broad denies; putting the deny first would block the studio's own web traffic and the laptop would be unusable. Blocking all inbound and all other outbound follows the rule that ports and services not needed for the device's role should always be blocked

(e)

- any two policies, each tied to a weakness: a **software installation policy** prohibiting staff installations and providing a request process, which closes the free-fonts-and-plugins route; an **acceptable use policy** stating what staff must do on studio devices - including not dismissing updates and not working on untrusted networks without the firewall enabled; a **password policy**, if the answer identified weak authentication among the weaknesses