Database Security & Forensics Practical Assessment ½

1. Why is Bridged networking mode preferred over NAT when hosting a database server in a VM?

- ☐ **A.** Bridged mode assigns the VM an IP address directly accessible on the local network, which is essential for testing connections from the host machine. ← Correct answer
- ☐ **B.** NAT mode consumes more RAM than Bridged mode.
- ☐ **C.** NAT mode prevents the installation of PostgreSQL.
- ☐ **D.** Bridged mode is more secure because it completely isolates the VM from the external network.

2. Why is Debian installed in minimal mode (no graphical interface) on a database server?

- ☐ **A.** The graphical interface reduces PostgreSQL performance by 80%.
- ☐ **B.** A server without a GUI reduces the attack surface, limits installed packages, decreases memory consumption and simplifies remote administration via SSH. ← Correct answer
- ☐ **C.** VirtualBox does not support graphical interfaces in Debian VMs.
- ☐ **D.** The graphical interface prevents PostgreSQL from starting automatically.

3. What is the direct consequence of a single-partition layout (one root partition /) on a production PostgreSQL server?

- ☐ **A.** PostgreSQL refuses to start if /var is not on a dedicated partition.
- ☐ **B.** Only one user can connect to the database at a time.
- ☐ **C.** Log or data bloat in /var can fill the root partition entirely, causing a complete system halt and potential data corruption. ← Correct answer
- ☐ **D.** Debian does not support single partitions on virtual disks.

4. Why does the lab require installing only SSH server and standard utilities, with no additional Debian components?

- ☐ **A.** Debian only allows one server role per installation.
- ☐ **B.** PostgreSQL is incompatible with Debian environments that have more than 200 packages installed.
- ☐ **C.** Extra packages conflict with the separate partitioning scheme.
- ☐ **D.** Every additional package widens the attack surface, increases dependencies to monitor and raises resource consumption; a minimal server is easier to audit and maintain. ← Correct answer

5. Why is the SSH port changed from 22 to a non-standard port? What is the main limitation of this measure?

- ☐ **A.** Changing the port improves SSH connection speed; the limitation is that some clients do not support ports above 1024.
- ☐ **B.** Changing the port reduces noise from automated bots that systematically target port 22; the limitation is that a full scan (nmap -p-) will reveal the new port to a determined attacker. ← Correct answer
- ☐ **C.** Changing the port strengthens SSH encryption with a stronger algorithm; the limitation is compatibility with older clients.
- ☐ **D.** Changing the port permanently prevents brute-force attacks; there is no limitation if the chosen port is above 10 000.

6. An administrator sets PermitRootLogin no in sshd_config but forgets the AllowUsers dbadmin directive. What security risk remains?

- ☐ **A.** No risk: the PermitRootLogin no directive alone is sufficient to secure SSH access.
- ☐ **B.** Without AllowUsers, any valid system account (www-data, postgres, nobody…) can attempt SSH authentication if an attacker obtains its credentials, significantly broadening the attack surface. ← Correct answer
- ☐ **C.** SSH refuses to start if the AllowUsers directive is absent.
- ☐ **D.** The postgres account will be able to connect but without any privileges, which is acceptable in a test environment.

7. Why is a 'deny all incoming' default policy applied in UFW rather than only blocking known dangerous ports?

- ☐ **A.** UFW can only block 10 ports simultaneously; blocking everything by default works around this limitation.
- ☐ **B.** Blocking only dangerous ports significantly slows packet processing by the kernel.
- ☐ **C.** A blacklist is inherently incomplete: a new vulnerable service on any port would remain exposed. A whitelist (allow only what is necessary) is exhaustive and covers all cases without exception. ← Correct answer
- ☐ **D.** Debian 12 requires the deny all incoming policy to allow package installation via apt.

8. An administrator runs 'ufw enable' before adding the rule allowing SSH. What happens immediately?

- ☐ **A.** UFW asks for additional confirmation and does not activate the policy until at least one SSH rule is defined.
- ☐ **B.** The 'deny all incoming' policy takes effect immediately and cuts the active SSH session. The administrator loses all remote access and can only restore it via physical console or KVM. ← Correct answer
- ☐ **C.** UFW detects the current SSH session and automatically excludes it from the default policy.
- ☐ **D.** Only the default port 22 is blocked; the custom port remains accessible because UFW does not know the current SSH configuration.

9. Fail2ban analyses system logs rather than live network traffic. What is the concrete consequence of this approach on its behaviour?

- ☐ **A.** Fail2ban can only monitor services that write their logs to /var/log/syslog.
- ☐ **B.** Fail2ban acts after attempts have already reached the target service: it tolerates the first N attempts (up to maxretry) before banning. It therefore complements UFW, which filters upstream at the network layer, and is not a substitute. ← Correct answer
- ☐ **C.** Fail2ban cannot monitor PostgreSQL because its logs are in a proprietary binary format.
- ☐ **D.** Log analysis makes Fail2ban inherently slower than signature-based network IDS systems.

10. What is the precise difference between the findtime and bantime parameters in Fail2ban configuration?

- ☐ **A.** findtime defines the ban duration; bantime defines the monitoring window before counters reset.
- ☐ **B.** findtime defines the time window within which maxretry failures must occur to trigger a ban; bantime defines how long the IP address remains banned after the threshold is exceeded. ← Correct answer
- ☐ **C.** findtime and bantime are two alternative names for the same parameter and can be used interchangeably.
- ☐ **D.** findtime applies exclusively to the SSH jail, while bantime only concerns the PostgreSQL jail.

11. Why should jail.conf be copied to jail.local before any Fail2ban customisation, rather than editing jail.conf directly?

- ☐ **A.** jail.conf is read-only on Debian 12 and cannot be edited directly.
- ☐ **B.** jail.conf is managed by the Fail2ban package: every update overwrites it and erases all customisations. jail.local is preserved during updates and its directives take priority over those in jail.conf. ← Correct answer
- ☐ **C.** Editing jail.conf directly makes the service unstable and requires a full system restart.
- ☐ **D.** jail.local unlocks advanced filtering features that are inaccessible from jail.conf.

12. Why is maxretry=3 configured for the SSH jail rather than maxretry=1, which would ban on the very first error?

- ☐ **A.** With maxretry=1, a legitimate administrator making a simple typo would be immediately banned, causing a self-inflicted denial of service. maxretry=3 tolerates common human errors while effectively blocking automated attack scripts. ← Correct answer
- ☐ **B.** Fail2ban does not support the value maxretry=1; the minimum allowed value is 3.
- ☐ **C.** A value of maxretry=1 silently disables the SSH jail in Fail2ban.
- ☐ **D.** The value maxretry=1 is only available in the commercial version of Fail2ban.

13. Why is the scram-sha-256 method recommended over md5 in pg_hba.conf for network connections to PostgreSQL?

- ☐ **A.** scram-sha-256 is significantly faster than md5 for local Unix socket connections.
- ☐ **B.** md5 sends a static hash over the network, vulnerable to rainbow-table and replay attacks. scram-sha-256 implements a challenge-response protocol (RFC 5802): the password is never transmitted and each session uses a unique random challenge, making replay impossible. ← Correct answer
- ☐ **C.** The md5 method was removed from PostgreSQL since version 14 and causes an error at startup.
- ☐ **D.** scram-sha-256 is a mandatory prerequisite to enable SSL in postgresql.conf.

14. A DBA leaves the following line in pg_hba.conf: 'host all all 0.0.0.0/0 trust'. What is the exact security consequence?

- ☐ **A.** PostgreSQL refuses to start if the trust method is paired with the 0.0.0.0/0 range.
- ☐ **B.** The trust + 0.0.0.0/0 combination is intercepted by UFW before reaching PostgreSQL.
- ☐ **C.** The trust directive is automatically ignored by PostgreSQL when SSL is enabled in postgresql.conf.
- ☐ **D.** Any host reachable on the Internet can connect to any database with any role name without providing a password, constituting a total compromise of data confidentiality, integrity and availability. ← Correct answer

15. An administrator wants to identify all active systemd services on a freshly installed Debian server in order to disable unnecessary ones. Which command does he use?

- ☐ **A.** ps aux | grep systemd
- ☐ **B.** systemctl list-units --type=service --state=running ← Correct answer
- ☐ **C.** service --status-all
- ☐ **D.** systemctl show --all

16. Which directives in /etc/ssh/sshd_config must be configured to: change the port to 49999, forbid root login, restrict access to the sole user dbadmin, and allow only public-key authentication?

- ☐ **A.** Port 49999 — RootLogin no — Users dbadmin — KeyAuth yes
- ☐ **B.** Port 49999 — PermitRootLogin no — AllowUsers dbadmin — PasswordAuthentication no ← Correct answer
- ☐ **C.** SSHPort 49999 — DenyRoot yes — AllowUsers dbadmin — KeyOnly yes
- ☐ **D.** Port 49999 — PermitRoot no — AuthUsers dbadmin — PubkeyOnly yes

17. Which UFW command allows PostgreSQL connections (port 5432) only from the loopback interface (localhost), strictly applying the principle of least privilege?

- ☐ **A.** ufw allow 5432/tcp
- ☐ **B.** ufw allow from 192.168.1.0/24 to any port 5432 proto tcp
- ☐ **C.** ufw allow from 127.0.0.1 to any port 5432 proto tcp ← Correct answer
- ☐ **D.** ufw limit 5432/tcp comment 'PostgreSQL local'

18. After configuring the SSH jail in jail.local, which sequence of commands applies the changes and confirms that the sshd and postgresql jails are active?

- ☐ **A.** systemctl restart fail2ban && fail2ban-client status ← Correct answer
- ☐ **B.** systemctl reload fail2ban
- ☐ **C.** fail2ban-client reload --all && fail2ban-client list
- ☐ **D.** systemctl daemon-reload && systemctl start fail2ban

19. After simulating 3 failed SSH login attempts from a test machine (IP: 192.168.1.50), which command verifies that this IP has been banned, and which command then unbans it manually?

- ☐ **A.** ufw status | grep 192.168.1.50 — then — ufw delete deny from 192.168.1.50
- ☐ **B.** fail2ban-client status sshd — then — fail2ban-client set sshd unbanip 192.168.1.50 ← Correct answer
- ☐ **C.** iptables -L | grep 192.168.1.50 — then — iptables -D INPUT -s 192.168.1.50 -j DROP
- ☐ **D.** tail /var/log/auth.log | grep banned — then — fail2ban-client unban 192.168.1.50

20. In jail.local, which configuration correctly defines the SSH jail on port 49999, with a maximum of 3 attempts and a 2-hour ban?

- ☐ **A.** [sshd] enabled = true port = 22 maxretry = 3 bantime = 120
- ☐ **B.** [ssh] active = true port = 49999 maxfail = 3 bantime = 2h
- ☐ **C.** [sshd] enabled = true port = 49999 maxretry = 3 bantime = 7200 ← Correct answer
- ☐ **D.** [sshd] enabled = true port = 49999 retry = 3 duration = 2h

---

Good luck — Rigour and method are your greatest assets.