Security+ SY0-701 Free Sample (30 Questions)

Unofficial practice material. Not affiliated with or endorsed by CompTIA.

This free sample contains 30 of the 865 questions from the full Security+ SY0-701 question bank.

Q1. A security team wants to detect data theft by planting a fake spreadsheet containing fabricated customer records on a file share. Any access to this file should trigger an alert. Which deception technique is being used?

A) A honeynet, because it emulates an entire network of decoys

B) A honeyfile, because it is a decoy file designed to attract unauthorized access

C) A honeypot, because it is a full decoy operating system

D) A firewall, because it filters access to the file share

Correct answer: B

Explanation: A honeyfile is a decoy document or file placed where attackers will find it. Legitimate users have no business opening it, so access attempts reliably indicate malicious activity. Why the others are wrong - A: A honeynet is a network of decoy systems, not a single decoy file.; C: A honeypot is a decoy system or service; a single planted spreadsheet is specifically a honeyfile.; D: A firewall controls traffic flow; it is not a deception technique.

Q2. A server runs file integrity monitoring software that alerts administrators whenever a critical operating system file is modified outside of a maintenance window. Which control type is this monitoring?

A) Preventive, because it blocks changes to system files

B) Detective, because it identifies unexpected changes to monitored files

C) Corrective, because it restores the changed files automatically

D) Directive, because it documents the file baseline policy

Correct answer: B

Explanation: File integrity monitoring watches critical system files and alerts when they change unexpectedly, identifying potential compromise in progress - a detective function. Why the others are wrong - A: FIM observes and alerts; it does not block the modification itself.; C: Standard FIM does not restore files; restoration would be a separate corrective action.; D: The baseline is a reference for detection, not a policy directive.

Q3. What is the primary purpose of a Trusted Platform Module (TPM) in a laptop?

A) Hardware-based key storage, sealed storage, and measured boot for platform integrity

B) Acting as a network firewall for the laptop

C) Issuing TLS certificates to public websites

D) Storing the user's browser passwords

Correct answer: A

Explanation: A TPM provides hardware-protected key storage, sealed storage tied to platform measurements, and measured boot, establishing trust in the device's integrity. Why the others are wrong - B: Packet filtering is done by host or network firewalls, not by the TPM.; C: Certificate issuance is the CA's role; a TPM never issues public certificates.; D: Password storage is handled by password managers or OS credential stores, not the TPM's core purpose.

Q4. In a properly designed CA hierarchy, which statement is true?

A) The root CA issues end-entity certificates directly every day
B) The root CA is kept offline and delegates issuance to intermediate CAs
C) The registration authority issues certificates directly to users
D) A single-tier CA with the root kept online is best practice

Correct answer: B

Explanation: Keeping the root CA offline (powered down, in a safe) minimizes its attack surface; day-to-day issuance is delegated to intermediate CAs whose compromise can be contained by revoking the intermediate. Why the others are wrong - A: Daily use of the root maximizes exposure of the most critical key in the PKI.; C: The RA only verifies identities; only a CA signs and issues certificates.; D: An online root is a single point of catastrophic compromise; tiering is the accepted best practice.

Q5. A compromised web server is cleaned and rebuilt, but the attacker returns two weeks later through an undocumented administrator account created during the original breach. The attacker MOST likely left a:

A) Worm
B) Logic bomb
C) Fileless malware implant that survives reboots
D) Backdoor

Correct answer: D

Explanation: A hidden account created to allow re-entry after remediation is a classic backdoor for persistent access. Why the others are wrong - A: A worm spreads between hosts; this was targeted re-entry to one server.; B: A logic bomb triggers on a condition rather than at the attacker's discretion.; C: Fileless implants live in memory and would not survive a full rebuild.

Q6. An accounts-payable clerk receives an email that appears to come from a known vendor. It references an actual invoice number from last month, uses the vendor's real logo and signature block, and asks the clerk to remit payment today to updated bank details. Which attack is this?

A) Whaling
B) Spear phishing
C) Typosquatting
D) Smishing

Correct answer: B

Explanation: This is spear phishing: a highly targeted email built with specific, researched details (a real invoice number, the vendor's identity) to make the lure convincing to one chosen victim or small group. Why the others are wrong - A: Whaling specifically targets senior executives or other high-value individuals as the victim.; C: Typosquatting abuses lookalike or misspelled domain names; the deception here is the personalized message content.; D: Smishing is phishing conducted via SMS text messages, not email.

Q7. Which of the following is the MOST effective defense against SQL injection?

A) Parameterized queries (prepared statements)
B) Client-side input length limits
C) Hiding database error messages
D) Blocking the word SELECT at the firewall

Correct answer: A

Explanation: Parameterized queries (prepared statements with bound parameters) separate SQL code from

user data at the database driver level, so injected input can never alter query logic. It is the definitive fix. Why the others are wrong - B: Client-side length limits are easily bypassed by sending requests directly to the server.; C: Suppressing error messages reduces information leakage but does not prevent injection.; D: Keyword blocking at the firewall is a fragile signature that attackers evade with encoding and alternative syntax.

Q8. A security engineer discovers the company's DNS servers answer recursive queries from any internet host. Which action BEST reduces the risk that these servers are abused in amplification attacks?

A) Deploying DNSSEC on the authoritative zones

B) Increasing the TTL on all DNS records

C) Blocking all inbound port 53 traffic at the perimeter

D) Restricting recursion to internal clients and enabling response rate limiting

Correct answer: D

Explanation: Open recursion is what makes a resolver abusable for amplification. Restricting recursion to internal clients removes the abuse vector, and response rate limiting (RRL) further caps reflected volume. Why the others are wrong - A: DNSSEC provides authenticity and integrity of DNS data; it does nothing to limit response volume or open recursion.; B: Longer TTLs do not prevent external hosts from using the resolver for reflection.; C: Blocking all inbound port 53 would break legitimate DNS resolution for the organization's own services.

Q9. A web application concatenates user input directly into database queries without checking it. Which vulnerability class does this represent at its root?

A) Broken authentication

B) Security misconfiguration

C) Improper input validation

D) Use of vulnerable components

Correct answer: C

Explanation: Trusting and processing raw user input without validation or sanitization is the root cause class - improper input validation - which enables injection attacks. Why the others are wrong - A: Broken authentication concerns flaws in identity verification and session handling, not query construction.; B: Security misconfiguration concerns insecure settings and defaults, not how input is handled in code.; D: Vulnerable components concern insecure third-party libraries, not the application's own input handling.

Q10. After deploying an emergency patch for a critical vulnerability, what is the BEST way to validate that the remediation worked?

A) Close the ticket as soon as the patch package finishes deploying

B) Wait for the next quarterly vulnerability scan

C) Ask the system owner whether the system seems to run faster

D) Run a targeted follow-up vulnerability scan against the host

Correct answer: D

Explanation: A follow-up scan verifies that the vulnerability is actually no longer detectable on the host, confirming the patch applied correctly and the exposure is closed. Why the others are wrong - A: Deployment does not guarantee success; a patch can fail to install or require a reboot.; B: Waiting months leaves the team uncertain whether the critical exposure is actually closed.; C: System performance is not a measure of

whether a vulnerability was remediated.

Q11. An organization wants a single enforcement point that can apply data loss prevention and access policies to all of its employees' sanctioned and unsanctioned SaaS applications. Which technology BEST meets this need?

A) A cloud access security broker (CASB)

B) A hardware security module (HSM)

C) A network intrusion prevention system

D) A traditional perimeter firewall

Correct answer: A

Explanation: A cloud access security broker (CASB) sits between users and cloud services as a policy enforcement point, providing visibility into SaaS usage (including shadow IT) and enforcing DLP, access control, and compliance policies across cloud applications. Why the others are wrong - B: An HSM manages cryptographic keys; it does not broker or inspect SaaS traffic.; C: A NIPS inspects network traffic for attacks but does not enforce SaaS-level DLP policy per application.; D: A perimeter firewall cannot see or control traffic to cloud/SaaS apps once users are off-network or using encrypted sessions.

Q12. A security assessment flags a fleet of smart cameras as a high risk. The cameras ship with a documented default admin password that cannot be changed. Which risk does this present?

A) The cameras transmit video over unencrypted HTTP, exposing footage to eavesdropping on the network

B) The cameras' firmware is no longer supported by the vendor, so known vulnerabilities cannot be patched

C) The cameras are assigned publicly routable IP addresses, making them directly reachable from the internet

D) The devices use well-documented default credentials that cannot be changed

Correct answer: D

Explanation: IoT devices frequently ship with hard-coded or default credentials that cannot be changed, so attackers can take over entire fleets with publicly known passwords - the first check is always whether defaults can be changed or the device belongs on an isolated segment. Why the others are wrong - A: Unencrypted video is a real camera risk, but the assessment specifically flagged the default credential weakness.; B: Unsupported firmware is a different IoT lifecycle risk; the immediate finding is the unchangeable default password.; C: Public IP exposure worsens the risk, but the core finding - an unchangeable, publicly known admin password - is the weakness itself.

Q13. Two manufacturing plants need their industrial control networks to exchange monitoring data continuously. The security policy requires the traffic to be encrypted, but plant devices cannot run VPN client software. Which solution meets these requirements?

A) Require each device to establish its own TLS VPN session

B) Open the plant firewalls to allow unencrypted control traffic

C) Deploy a site-to-site IPsec VPN between the plants' gateway firewalls

D) Email the monitoring data as encrypted attachments

Correct answer: C

Explanation: A gateway-to-gateway IPsec tunnel encrypts all inter-site traffic transparently, so the constrained plant devices need no client software. Why the others are wrong - A: The scenario states the devices cannot run VPN clients.; B: Unencrypted control-system traffic across the internet violates the encryption requirement and exposes industrial systems.; D: Email is not a continuous, automated data-exchange

mechanism.

Q14. A web farm's application servers are CPU-saturated, and profiling shows TLS encryption is consuming most of the cycles. Which load balancer feature BEST relieves the backend servers?

A) Add more backend servers to the farm so the TLS processing load is spread across additional systems

B) Move TLS termination to the database servers, which have spare CPU capacity during business hours

C) Disable TLS encryption entirely so no server in the farm spends cycles on cryptography

D) SSL/TLS offloading at the load balancer, terminating client encryption there and using plain HTTP to backends over the trusted internal network

Correct answer: D

Explanation: Offloading moves the expensive cryptographic work to the load balancer - often with hardware acceleration - freeing application servers. Backend traffic stays on a trusted internal segment. Why the others are wrong - A: Scaling out without offloading just spreads the same cryptographic cost across more servers instead of removing it.; B: Databases should not terminate web TLS sessions; that misplaces the workload rather than centralizing it.; C: Removing encryption exposes customer data in transit - the goal is to move the work, not drop the protection.

Q15. A payment processor must store cardholder data so that a database breach reveals nothing usable, yet authorized systems can still retrieve original values. The architect chooses tokenization over encryption primarily because...

A) Tokens have no mathematical relationship to the original values, reducing the scope of systems that handle sensitive data

B) Tokenization is always faster than any encryption algorithm

C) Tokens can be cracked more easily to aid forensic analysis

D) Tokenization removes the need for access controls

Correct answer: A

Explanation: Tokens have no mathematical relationship to the original values, so systems handling tokens fall outside sensitive-data scope and a breach reveals nothing usable. Why the others are wrong - B: Performance is not the driver, and tokenization is not guaranteed to be faster than encryption.; C: Tokens are random substitutes with no derivable link to the originals, so there is nothing to crack for forensics.; D: Access controls are still required around the token vault and the systems authorized to detokenize.

Q16. An administrator regularly restores a random backup set to an isolated system to confirm it works. This practice is best described as...

A) Backup encryption

B) Backup compression

C) Backup deduplication

D) Verification through test restores

Correct answer: D

Explanation: Regularly restoring backups in an isolated environment is the only way to prove the backups are complete and actually recoverable. Why the others are wrong - A: Encryption protects backup confidentiality; it does not prove the backup will restore.; B: Compression reduces backup size but says nothing about recoverability.; C: Deduplication saves storage space; it does not validate that restores succeed.

Q17. A team shares one 'admin' login for a network appliance. What is the primary security problem?

A) Password complexity requirements cannot be enforced reliably on a shared credential

B) The account cannot be protected with multi-factor authentication
C) Actions cannot be attributed to an individual, destroying accountability
D) Password rotation policies become difficult to enforce without disrupting the whole team
Correct answer: C
Explanation: Non-repudiation requires unique identities. With a shared login, no action can be traced to a person, making incident investigation and insider-threat detection impossible. Why the others are wrong - A: Complexity rules can technically apply to a shared password; the deeper failure is that no action can ever be attributed to an individual.; B: MFA can technically be attached to a shared account, and even then actions still could not be traced to one person.; D: Rotation is an operational headache with shared credentials, but the primary security failure is the loss of individual accountability, not the schedule.

Q18. At 3:00 AM the SIEM generates a critical alert for data exfiltration from a database server. What should the analyst do FIRST?
A) Shut down the database server immediately
B) Call law enforcement to report a breach
C) Validate the alert against raw logs to rule out a false positive
D) Delete the alert and retune the rule
Correct answer: C
Explanation: Always verify before acting: check the raw logs to confirm the alert is genuine. Premature shutdown destroys volatile evidence, and false positives are common in SIEM alerting. Why the others are wrong - A: Shutting down the server destroys volatile evidence and disrupts the business; containment comes after confirmation.; B: Law enforcement notification is a later management and legal decision, not a first response step.; D: Deleting an unvalidated alert could erase evidence of a genuine incident.

Q19. A vulnerability scanner runs nightly and generates 500 SIEM alerts each run, drowning out real alerts. What should the analyst do?
A) Delete the scanner
B) Turn off the SIEM
C) Ignore all future alerts
D) Tune the SIEM correlation rule to whitelist or suppress the scanner's known-benign traffic
Correct answer: D
Explanation: Known-benign sources should be baselined or whitelisted in correlation logic so genuine anomalies stand out instead of drowning in scanner noise. Why the others are wrong - A: The scanner provides essential visibility; removing it weakens the security program.; B: Disabling the SIEM eliminates all detection capability.; C: Ignoring everything guarantees missed incidents - the alert-fatigue failure mode.

Q20. The chain of custody is primarily used to:
A) Prioritize security alerts based on their severity and impact
B) Identify the attack vector used during the initial compromise
C) Speed up system recovery after the incident is contained
D) Document who handled evidence and when, preserving its admissibility
Correct answer: D
Explanation: Chain of custody is a documented, unbroken record of every person who handled evidence and what was done with it, proving the evidence was not altered - required for it to be admissible. Why the others are wrong - A: Alert prioritization is a detection/triage activity, unrelated to evidence handling.; B: Attack

vectors are determined during analysis, not by custody records.; C: Recovery speed is an operations concern; custody protects legal integrity, not speed.

Q21. Which metric is MOST useful to an IR team assessing how long an attacker operated undetected?

A) Dwell time (time from initial compromise to detection)

B) The number of office printers

C) Total disk capacity

D) Average office temperature

Correct answer: A

Explanation: Dwell time measures the gap between compromise and detection - longer dwell means deeper potential compromise and guides scoping. Why the others are wrong - B: Printer counts have no bearing on incident timelines.; C: Disk capacity does not measure attacker dwell.; D: Temperature is irrelevant to incident timelines.

Q22. A call center uses shared workstations where agents rotate every few hours. Agents sometimes walk away without logging out. Which session control best protects the unattended sessions?

A) Disable screen locks so agents can work faster

B) Set an automatic screen lock after 5 minutes of inactivity

C) Tape each agent's password to the monitor for convenience

D) Share one user account among all agents

Correct answer: B

Explanation: An automatic inactivity lock protects unattended shared machines: after five idle minutes the session requires re-authentication, so the next agent cannot access the previous agent's data. Why the others are wrong - A: Disabling locks removes the only protection for unattended sessions.; C: Writing passwords down defeats authentication entirely.; D: Shared accounts destroy accountability and let anyone access the previous session.

Q23. How should WSUS approvals be structured to balance timely patching with the risk of a faulty update?

A) Auto-approve all updates for every system immediately

B) Disable automatic updates entirely

C) Use update rings: pilot group first, then production

D) Manual updates only, performed one machine at a time

Correct answer: C

Explanation: Update rings stage approvals: a pilot group receives updates first and is monitored for issues, then the approval widens to production. Immediate fleet-wide auto-approval risks a bad update everywhere at once. Why the others are wrong - A: Immediate fleet-wide approval turns one bad update into an enterprise outage.; B: Disabling updates guarantees growing vulnerability exposure.; D: One-at-a-time manual updates cannot scale to hundreds of systems.

Q24. An organization uses geofencing on its MDM platform. What capability does this provide?

A) Tracking device battery levels across regions

B) Encrypting email attachments automatically

C) Blocking all third-party applications

D) Enforcing policies based on the device's physical location

Correct answer: D

Explanation: Geofencing applies different policies depending on where the device is - for example, disabling cameras or requiring VPN inside a restricted facility, and relaxing rules elsewhere. Why the others are wrong - A: Battery monitoring is a device-health feature, not geofencing.; B: Email encryption is handled by encryption policies, not location.; C: App blocking is an application-control policy, independent of location.

Q25. The primary purpose of the PCI-DSS is to:

A) Mandate a specific encryption algorithm for all industries

B) Protect cardholder data and reduce payment card fraud

C) Require all companies to undergo annual financial audits

D) Standardize password lengths across all operating systems

Correct answer: B

Explanation: PCI-DSS is a set of security requirements created by the payment card industry to protect cardholder data and reduce fraud wherever payment cards are accepted. Why the others are wrong - A: PCI-DSS requires strong cryptography for cardholder data but does not mandate one specific algorithm for all industries.; C: Financial audits are a separate accounting matter; PCI-DSS assessments focus specifically on cardholder data security.; D: PCI-DSS addresses authentication broadly but does not set password standards for all operating systems.

Q26. A compliance officer discovers that required access reviews have not been performed for six months. What is the BEST next step?

A) Hide the finding to avoid trouble

B) Delete the access logs

C) Wait until the external audit to mention it

D) Document the gap, report it to management, and schedule the overdue reviews

Correct answer: D

Explanation: Governance requires transparency: document the finding, escalate it to management, and remediate by scheduling the overdue reviews and tracking completion. Why the others are wrong - A: Hiding a finding is unethical and worsens liability when the gap is later discovered.; B: Deleting logs destroys audit evidence and is itself a compliance violation.; C: Waiting for the external audit delays remediation and turns a fixable gap into an audit finding.

Q27. A company decides that the cost of fixing a low-rated vulnerability exceeds its benefit. What is the correct risk response?

A) Removing the vulnerability from the risk register since it will not be fixed

B) Buying insurance

C) Formal documented acceptance by the risk owner, acknowledging residual risk

D) Accepting the risk informally without documentation, since the rating is low

Correct answer: C

Explanation: Acceptance is a conscious, documented decision by the risk owner to live with the residual risk - not ignoring it. Why the others are wrong - A: Removing the entry hides the risk instead of managing it; the register must track accepted risks.; B: Insurance shifts the financial impact to an insurer, which is risk transfer - not acceptance.; D: Informal acceptance is not acceptance at all; the response must be documented by the risk owner to be valid.

Q28. How often should an organization perform risk assessments?

A) On a regular schedule and after significant changes to systems, threats, or business processes
B) Once at company founding
C) Only after a breach
D) Every ten years
Correct answer: A
Explanation: Risk assessments are repeated on a schedule and re-run when systems, threats, vendors, or business processes change materially. Why the others are wrong - B: A one-time assessment goes stale immediately.; C: Waiting for a breach is reactive, not risk management.; D: A decade between assessments is far too infrequent.

Q29. A bank wants to test how its SOC analysts detect and respond to a stealthy, multi-week intrusion, giving the testers only the limited information a real external attacker would have. Which engagement best fits?
A) A quarterly unauthenticated vulnerability scan of the public-facing IP range, performed to satisfy audit requirements
B) A compliance audit reviewing the SOC's policies and procedures against a security framework
C) Unknown-environment (black-box) penetration test, possibly as a red team exercise
D) A known-environment review of the SIEM detection rules to find logic errors
Correct answer: C
Explanation: Testing detection and response against a stealthy, attacker-realistic intrusion calls for a black-box (or red-team) engagement with minimal starting information. Why the others are wrong - A: Vulnerability scans enumerate weaknesses; they do not exercise the SOC's detection and response against a live intrusion.; B: Compliance audits check adherence to requirements; they do not test live detection.; D: A white-box rule review finds detection logic flaws but never exercises the analysts or the response process.

Q30. A CISO has budget for only two of five proposed controls. How should she choose?
A) Fund the controls that reduce the most risk per dollar spent
B) Fund the two most expensive controls
C) Fund the controls the vendors recommend most
D) Split the budget equally among all five
Correct answer: A
Explanation: With constrained budgets, funding the controls that reduce the most risk per dollar is the core of risk-based program management. Why the others are wrong - B: Cost does not equal risk reduction value.; C: Vendor recommendations carry sales bias and must be evaluated independently.; D: Equal splits dilute funding and ignore risk differences.