DUMPSARENA

Certified Internet of Things Security Practitioner (CIoTSP)

CertNexus ITS-110

Version Demo

Total Demo Questions: 13

Total Premium Questions: 131

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Planning an IoT Implementation	13
Topic 2, Implementing IoT Security	108
Topic 3, Managing IoT Security Operations	5
Topic 4, Responding to IoT Security Incidents	5
Total	131

QUESTION NO: 1

An IoT device which allows unprotected shell access via console ports is most vulnerable to which of the following risks?

A. Directory harvesting

B. Rainbow table attacks

C. Malware installation

D. Buffer overflow

ANSWER: C

Explanation:

Malware installation is correct because an unprotected console port can give a person with physical access an interactive shell on the IoT device without first defeating an authentication control. Shell access enables execution of operating-system commands, inspection and modification of the file system, changing startup configuration, and loading or downloading unauthorized executable code. An attacker can therefore install malware that persists after reboot, such as a remote-access backdoor, botnet client, credential-stealing tool, or code that alters device operation. This is especially serious for IoT devices because they are often deployed unattended in physically accessible locations and may have limited monitoring or endpoint protections. Secure device design should restrict physical debug and console interfaces, require strong authentication and authorization when maintenance access is needed, and ensure that only trusted, cryptographically verified software can execute. NIST identifies protection of device interfaces and software-update mechanisms as important IoT device cybersecurity capabilities; these controls help prevent unauthorized modification of device software. See NISTIR 8259A, IoT Device Cybersecurity Capability Core Baseline and OWASP Internet of Things Project.

QUESTION NO: 2

You work for a business-to-consumer (B2C) IoT device company. Your organization wishes to publish an annual report showing statistics related to the volume and variety of sensor data it collects. Which of the following should your organization do prior to using this information?

A. Confirm the devices they've sold are turned on

B. Ensure all sensors are running the latest software

C. Require customers to sign a subscription license

D. Remove any customer-specific data

ANSWER: D

Explanation:

Remove any customer-specific data is correct because an annual report about aggregate sensor-data volume and variety should be prepared with privacy protection as a core requirement. Before data is used for publication, the organization should de-identify it so that the report cannot disclose information linked to an identifiable customer, household, device owner, or account. This includes direct identifiers, such as names and email addresses, as well as data elements that could reasonably be combined to re-identify an individual. The report can then communicate useful high-level metrics—such as total records processed, categories of sensor readings, or trends in collection—without unnecessarily exposing personal information. This practice supports privacy-by-design principles for IoT systems, where data minimization and appropriate handling of personal data are built into the lifecycle of collection, analysis, and disclosure. It also helps the organization meet its responsibility to protect consumer data and reduce privacy risk when information is shared outside the company. The U.S. National Institute of Standards and Technology describes de-identification as a privacy-enhancing technique intended to reduce the risk of associating data with individuals; see NISTIR 8053, De-Identification of Personal Information. The UK Information Commissioner's Office also provides practical guidance on anonymisation and managing re-identification risk at ICO anonymisation guidance.

QUESTION NO: 3

An IoT systems integrator has a very old IoT gateway that doesn't offer many security features besides viewing a system configuration page via browser over HTTPS. The systems integrator can't get their modern browser to bring up the page due to a cipher suite mismatch. Which of the following must the integrator perform before the configuration page can be viewed?

A. Upgrade the browser, as modern browsers have stopped allowing connections to hosts that use only outdated cipher suites.

B. Downgrade the browser, as modern browsers have stopped allowing connections to hosts that use only outdated cipher suites.

C. Upgrade the browser, as older browsers have stopped allowing connections to hosts that use only outdated cipher suites.

D. Downgrade the browser, as modern browsers have continued allowing connections to hosts that use only outdated cipher suites.

ANSWER: B

Explanation:

Downgrade the browser, as modern browsers have stopped allowing connections to hosts that use only outdated cipher suites. is correct because a cipher suite mismatch means the browser and gateway cannot negotiate a mutually supported set of cryptographic algorithms for the HTTPS session. Older IoT gateways commonly support only deprecated TLS versions, weak key-exchange methods, or legacy cipher suites that current browsers intentionally disable to protect users from known cryptographic weaknesses. If the gateway cannot be upgraded or reconfigured to support a current TLS configuration, the integrator needs a suitably older browser that still implements the gateway's legacy protocol and cipher-suite requirements in order to establish the management connection.

This should be performed only as a tightly controlled, temporary administrative measure. The older browser and gateway should be isolated on a trusted management network, and the browser should not be used for ordinary web access. After configuration access is restored, the preferred remediation is to replace or upgrade the gateway, or update its firmware and cryptographic settings, so that it supports modern HTTPS protections. Mozilla's guidance on secure server configurations is available at Mozilla Server Side TLS, and Google documents the removal of obsolete TLS support at Chromium Blog: Deprecating TLS 1.0 and TLS 1.1.

QUESTION NO: 4

An embedded engineer wants to implement security features to be sure that the IoT gateway under development will only load verified images. Which of the following countermeasures could be used to achieve this goal?

A. Implement Over-The-Air (OTA) updates

B. Enforce a secure boot function

C. Enforce a measured boot function

D. Harden the update server

ANSWER: B

Explanation:

Enforce a secure boot function is correct because secure boot establishes a hardware-anchored chain of trust from the gateway's immutable boot code through each subsequent boot stage and the operating-system or firmware image. Before code is loaded or executed, the secure-boot process verifies its cryptographic signature against a trusted public key or certificate stored in protected device hardware or firmware. A valid signature demonstrates that the image was produced by an authorized signer and has not been altered since signing. If signature verification fails, the device can halt the boot process or fall back to a known-good recovery image, preventing an unverified image from running. This is particularly important for IoT gateways because they often bridge less-trusted device networks with enterprise or cloud services; compromise at the boot layer could undermine all higher-level protections. Secure boot therefore directly enforces the requirement that only verified images are loaded, including images installed through a legitimate update workflow. NIST

describes secure boot as a mechanism for verifying the authenticity and integrity of firmware and software before execution, while the U.S. Cybersecurity and Infrastructure Security Agency identifies secure boot as a foundational firmware security control. See NIST's Secure Boot definition and CISA guidance on Secure Boot.

QUESTION NO: 5

An IoT cloud service exposes an application programming interface (API) used by mobile applications and device-management portals. Which of the following controls should the service implement to reduce the risk of API abuse? (Choose two.)

A. Implement rate limiting for API requests

B. Require strong authentication and authorization for API requests

C. Trust authorization decisions supplied by the mobile application

D. Allow unlimited unauthenticated requests to management functions

E. Return detailed internal error messages to every caller

ANSWER: A B

Explanation:

Rate limiting restricts the number of requests a client can make during a defined period. This helps reduce the impact of automated credential attacks, excessive resource consumption, API scraping, and some denial-of-service conditions. Rate limits should be tuned so that legitimate device-management and operational activity can continue while abnormal request patterns are constrained.

Requiring strong authentication and authorization for API requests ensures that callers are identified and can access only the functions and data permitted for their identity or role. Authorization should be enforced server-side for every request, rather than relying on a mobile application or portal to hide unauthorized functions. OWASP identifies broken authentication, broken authorization, and unrestricted resource consumption as significant API security risks; see the OWASP API Security Project.

QUESTION NO: 6

An IoT security administrator confirms that several internet-connected cameras are communicating with known command-and-control infrastructure. Which of the following actions should the administrator take first to contain the incident and support investigation? (Choose two.)

A. Isolate the affected cameras from the network

B. Preserve relevant logs and device evidence

C. Immediately reconnect the cameras to confirm that the alert was accurate

D. Delete all camera logs to free storage capacity

E. Publish the administrator credentials in the incident ticket

ANSWER: A B

Explanation:

Isolating the affected cameras from the network is correct because it limits further command-and-control communications and reduces the opportunity for the compromised devices to spread malware, participate in attacks, or transmit additional data. Isolation should be performed in a way that considers operational and safety requirements.

Preserving relevant logs and device evidence is also correct because incident records, network captures, device configurations, and available volatile data can help determine the scope, initial access path, malware behavior, and affected

systems. Evidence should be collected and handled according to established incident-response procedures. NIST describes containment and evidence collection activities in NIST SP 800-61 Rev. 2.

QUESTION NO: 7

An IoT healthcare service needs to allow support personnel to troubleshoot individual devices while limiting unnecessary exposure of patient-associated telemetry. Which of the following controls support this objective? (Choose two.)

A. Implement Role-Based Access Control (RBAC)

B. Pseudonymize patient identifiers

C. Use a single shared support account for all technicians

D. Retain all telemetry indefinitely for possible future analysis

E. Disable audit logging for support sessions

ANSWER: A B

Explanation:

Role-Based Access Control (RBAC) is correct because it can limit support personnel to the device records, functions, and data types required for their assigned duties. Rather than granting broad access to all telemetry, the service can assign permissions based on a defined support role and the principle of least privilege.

Pseudonymizing patient identifiers also reduces privacy risk by replacing direct identifiers with values that do not directly identify the patient. The mapping between the pseudonym and the person must remain separately protected, because pseudonymized data can still be personal data when re-identification is possible. NIST describes access control and least privilege in NIST SP 800-53 Rev. 5, and the European Data Protection Board discusses pseudonymisation in EDPB guidance on data-processing principles.

QUESTION NO: 8

An IoT security administrator needs to improve the ability to investigate unauthorized configuration changes on remote gateways. Which of the following logging practices should the administrator implement? (Choose two.)

A. Forward logs to a centralized logging system

B. Record administrative authentication events and configuration changes

C. Allow administrators to erase logs after each maintenance session

D. Store only successful administrative login events

E. Disable logging to conserve storage capacity

ANSWER: A B

Explanation:

Forwarding logs to a centralized logging system helps preserve evidence when a gateway is compromised, damaged, or reset. Central collection also enables correlation among gateway, network, identity, and cloud-service events. The logging system should restrict access and preserve records according to the organization’s retention requirements.

Recording administrative authentication events and configuration changes provides accountability for security-relevant actions. These records can help investigators determine which account accessed a gateway, when access occurred, what was changed, and whether the activity was authorized. NIST recommends collecting and protecting logs that support operational monitoring, incident response, and forensic analysis; see NIST SP 800-92.

QUESTION NO: 9

An IoT security administrator is concerned about an external attacker using the internal device management local area network (LAN) to compromise his IoT devices. Which of the following countermeasures should the security administrator implement? (Choose three.)

A. Require the use of Password Authentication Protocol (PAP)

B. Create a separate management virtual LAN (VLAN)

C. Ensure that all IoT management servers are running antivirus software

D. Implement 802.1X for authentication

E. Ensure that the Time To Live (TTL) flag for outgoing packets is set to 1

F. Only allow outbound traffic from the management LAN

G. Ensure that all administrators access the management server at specific times

ANSWER: B D F

Explanation:

Creating a separate management virtual LAN (VLAN), implementing 802.1X for authentication, and only allowing outbound traffic from the management LAN together establish layered protection for an IoT management environment. A dedicated management VLAN separates administrative systems and device-management traffic from ordinary user and production networks. This reduces the number of systems that can reach high-value management interfaces and provides a clear boundary at which network security controls can be enforced.

802.1X adds port-based network access control: before a device or user is granted access to the management network, it must authenticate through an appropriate authentication infrastructure. This helps prevent an unauthorized system from simply connecting to an internal switch port or wireless network and gaining management-LAN access. Restricting the management LAN to outbound traffic further limits unsolicited inbound connections from other internal segments or external networks, reducing opportunities for attackers to directly reach management servers and IoT administration interfaces. Such segmentation, explicit access control, and least-privilege communications align with zero-trust principles and industrial/operational network guidance. See NIST SP 800-207, Zero Trust Architecture and NIST SP 800-82 Rev. 3, Guide to Operational Technology Security.

QUESTION NO: 10

An IoT security practitioner should be aware of which common misconception regarding data in motion?

A. That transmitted data is point-to-point and therefore a third party does not exist.

B. The assumption that all data is encrypted properly and cannot be exploited.

C. That data can change instantly so old data is of no value.

D. The assumption that network protocols automatically encrypt data on the fly.

ANSWER: B

Explanation:

The assumption that all data is encrypted properly and cannot be exploited is the key misconception. Data in motion includes information transmitted among IoT devices, gateways, cloud services, management platforms, and user applications. Practitioners cannot simply presume that every connection has encryption enabled, that the selected cryptographic protocol is appropriate, or that its implementation and configuration securely protect the communication. Encryption must be intentionally designed into each relevant communication path and supported by sound certificate, key, identity, and endpoint-management practices.

Even when transport encryption is present, protection depends on such details as validation of peer identities, use of current protocol versions and cipher suites, secure key storage and rotation, and resistance to compromised endpoints. Therefore, encryption is an essential control for data in motion, but it is not an automatic or absolute assurance that transmitted IoT data cannot be exposed, modified, or otherwise abused. NIST's guidance on secure transport configurations emphasizes selecting and correctly configuring approved transport-layer protections: NIST SP 800-52 Rev. 2. NIST also identifies data protection and secure device communications as core IoT cybersecurity capabilities: NISTIR 8259A.

QUESTION NO: 11

The network administrator for an organization has read several recent articles stating that replay attacks are on the rise. Which of the following secure protocols could the administrator implement to prevent replay attacks via remote workers' VPNs? (Choose three.)

A. Internet Protocol Security (IPSec)

B. Enhanced Interior Gateway Routing Protocol (EIGRP)

C. Password Authentication Protocol (PAP)

D. Challenge Handshake Authentication Protocol (CHAP)

E. Simple Network Management Protocol (SNMP)

F. Layer 2 Tunneling Protocol (L2TP)

G. Interior Gateway Routing Protocol (IGRP)

ANSWER: A D F

Explanation:

Internet Protocol Security (IPSec), Challenge Handshake Authentication Protocol (CHAP), and Layer 2 Tunneling Protocol (L2TP) are the applicable technologies for a VPN solution designed to resist replay attacks. IPSec provides explicit anti-replay protection through sequence numbers and a sliding receive window in its Authentication Header and Encapsulating Security Payload processing. A receiver can identify and reject packets that are duplicated or fall outside the acceptable sequence window, preventing captured protected traffic from being successfully resent.

Challenge Handshake Authentication Protocol uses a challenge-response exchange rather than sending a reusable password across the connection. The authenticator supplies a new, unpredictable challenge, and the response is derived from that challenge and the shared secret. Consequently, a response captured from an earlier authentication exchange cannot simply be replayed for a new challenge.

Layer 2 Tunneling Protocol is a VPN tunneling protocol commonly deployed with IPSec as L2TP/IPSec. In that architecture, L2TP supplies the remote-access tunneling capability while IPSec supplies confidentiality, integrity, peer authentication, and anti-replay safeguards for the tunneled traffic. This combination is a recognized remote-worker VPN deployment model. See RFC 4303: IP Encapsulating Security Payload and RFC 1994: Challenge Handshake Authentication Protocol.

QUESTION NO: 12

An IoT service collects massive amounts of data and the developer is encrypting the data, forcing administrative users to authenticate and be authorized. The data is being disposed of properly and on a timely basis. However, which of the following countermeasures is the developer most likely overlooking?

A. That private data can never be fully destroyed.

B. The best practice to only collect critical data and nothing more.

C. That data isn't valuable unless it's used as evidence for crime committed.

D. That data is only valuable as perceived by the beholder.

ANSWER: B

Explanation:

The best practice to only collect critical data and nothing more is the overlooked countermeasure. The stated safeguards protect data after collection: encryption helps preserve confidentiality, authentication and authorization restrict administrative access, and timely disposal limits retention. However, a secure IoT data-handling program must also address whether each data element needs to be collected in the first place. Data minimization means limiting collection to data that is relevant, necessary, and proportionate to the defined service purpose. This reduces the amount of sensitive information exposed if a device, cloud service, administrative account, or data pipeline is compromised. It also reduces storage, governance, retention, and compliance burdens throughout the data lifecycle. In an IoT environment, where sensors can continuously generate detailed telemetry, location information, audio, video, or behavioral data, limiting collection is especially important because apparently routine data can become sensitive when aggregated or correlated. The NIST Privacy Framework identifies data minimization as a privacy-protective data processing practice, and the FTC's IoT guidance similarly emphasizes collecting only data needed for a legitimate purpose. Applying minimization at design time complements the encryption, access controls, and disposal practices already described. See the NIST Privacy Framework and the FTC IoT security guidance.

QUESTION NO: 13

A compromised IoT device is initiating random connections to an attacker's server in order to exfiltrate sensitive data. Which type of attack is being used?

A. Man-in-the-middle (MITM)

B. SSL session hijack

C. Reverse shell

D. Honeypot

ANSWER: C

Explanation:

Reverse shell is correct because it causes a compromised device to initiate an outbound connection from the victim environment to infrastructure controlled by the attacker. Once this connection is established, the attacker can interact with a command shell on the device remotely, issue commands, collect data, and transfer that data out of the network. Outbound connections are particularly useful to attackers because perimeter firewalls and network-address translation commonly make unsolicited inbound access to an IoT device difficult, while outbound traffic may be allowed or less closely scrutinized.

For IoT environments, a reverse shell can be delivered after exploiting weak credentials, an exposed service, vulnerable firmware, or an application flaw. The resulting channel may be intermittent or use varying connection patterns to evade detection, while still serving as a command-and-control path and a route for exfiltration. Security teams should monitor unusual outbound connections from IoT assets, especially connections to unapproved external hosts, unexpected ports, or destinations with no legitimate business purpose. MITRE ATT&CK documents command and scripting interpreter activity at https://attack.mitre.org/techniques/T1059/ and describes application-layer command-and-control channels at https://attack.mitre.org/techniques/T1071/.