PassQueen
a.k.a dumpsqueen.com

Developing Solutions Using Cisco IoT and Edge Platforms (DEVIOT)

Cisco 300-915

Version Demo

Total Demo Questions: 8

Total Premium Questions: 89

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cisco IoT and Edge Architecture	18
Topic 2, Edge Compute	8
Topic 3, Edge Data Networking	22
Topic 4, Cisco IOx Applications	24
Topic 5, Cisco Kinetic	14
Topic 6, Cisco IoT Operations Dashboard	3
Total	89

QUESTION NO: 1

Refer to the exhibit. When a laptop is connected to a device via a serial connection, which action makes the output readable?

A. Change the baudrate in the connection settings.

B. Reboot the device in case a firmware error is the issue.

C. Change the power level of the device.

D. Change the unicode settings to UTF-8.

ANSWER: A

Explanation:

Change the baudrate in the connection settings. Garbled or unreadable characters in a serial-console session are typically caused by a mismatch between the terminal emulator's line settings and the serial settings configured on the connected device. Baud rate defines the signaling speed used by both ends of the serial connection; if the laptop uses a different value, it samples the transmitted bits at the wrong timing and displays meaningless characters. Set the terminal emulator to the device's configured console baud rate, commonly 9600 bps by default on many Cisco platforms, then reconnect if necessary. The remaining serial parameters must also match the device—typically 8 data bits, no parity, 1 stop bit, and no flow control—but adjusting the baud rate is the action that resolves the displayed symptom when the speed is mismatched. Cisco's console-access guidance documents the required terminal-emulation settings and notes the common default console speed. See Cisco: Connecting a Terminal to the Console Port and Cisco terminal and console-line configuration documentation.

QUESTION NO: 2

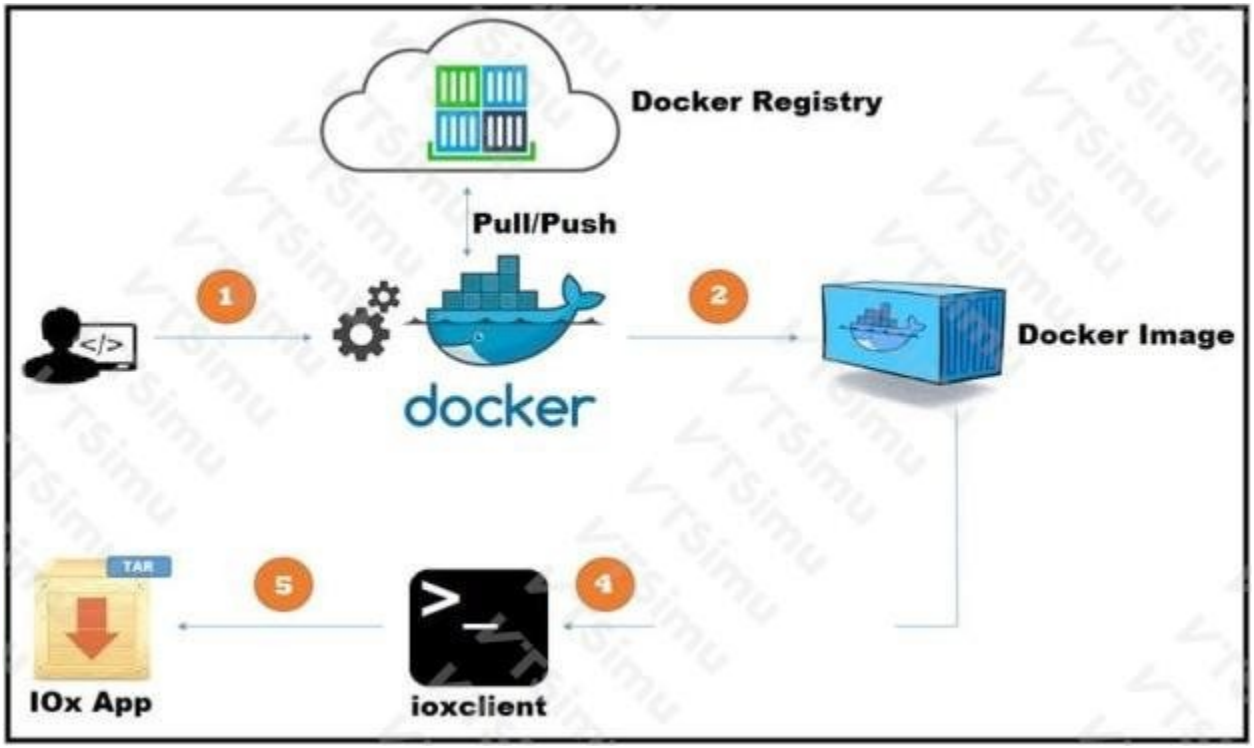

Refer to the exhibit. A new application has just been developed as a Docker container.

What is the missing step for packaging the container ready for deployment in IOx?

A. Pull/push the image to the Docker registry.

B. Build the package.yaml files.

C. Build the package.cert file to sign the app.

D. Log in to device manager and load the application.

ANSWER: B

Explanation:

Build the package.yaml files. is correct because an IOx application package requires a package descriptor, conventionally named `package.yaml`, in addition to the Docker application image. The descriptor supplies the metadata and deployment requirements that IOx needs to recognize and install the application, including such information as the application name, version, description, architecture, and resource or networking-related configuration references. The Docker image alone is not an IOx deployable application artifact. During the packaging workflow, the image is saved and combined with the package descriptor into the IOx package archive, which can then be installed on a compatible IOx-enabled device. This descriptor therefore bridges the container development stage and the creation of an installable IOx package. Cisco's IOx documentation describes application packaging as the process of preparing the container image together with the required package metadata for deployment. See the Cisco IOx Developer Guide and the Cisco IOx application packaging documentation.

QUESTION NO: 3

Refer to the exhibit.

```
"/Users/venv/bin/python" "/Users/Documents/Coding Projects/fkf3/main.py"
log: Sending CONNECT (u1, p1, wr0, wq0, wf0, c1, k60) client_id=b'python'
log: Sending SUBSCRIBE (d0, m1) [(b'eschbornoffice/things/FCD6BD0001913F4/uplink', 0)]
log: Sending SUBSCRIBE (d0, m2) [(b'eschbornoffice/things/FCD6BD0000191405/uplink', 0)]
log: Sending SUBSCRIBE (d0, m3) [(b'eschbornoffice/things/FCD6BD0001913FC/uplink', 0)]
log: Sending SUBSCRIBE (d0, m4) [(b'eschbornoffice/things/FCD6BD0001913F2/uplink', 0)]
log: Received CONNACK (0, 0)
Result code 0
log: Received SUBACK
log: Received SUBACK
log: Received SUBACK
log: Received SUBACK
```

Which two statements about MQTT messaging are true? (Choose two.)

A. Result code 5 means the connection is accepted; result code 0 means the connection is " refused, not authorized " .

B. Redundant ClientIDs are supported by MQTT brokers.

C. MQTT does not support a " keep alive " functionality because it runs on top of TCP.

D. Result code 0 means the connection is accepted; result code 5 means the connection is " refused, not authorized " .

E. Brokers can send the Last Will and Testament messages on behalf of the client.

ANSWER: D E

Explanation:

“Result code 0 means the connection is accepted; result code 5 means the connection is ‘refused, not authorized’” is correct for MQTT 3.1.1 CONNACK processing. A broker returns a CONNACK return code after receiving a client CONNECT packet. Return code 0 indicates that the connection has been accepted. Return code 5 identifies a connection refusal because the client is not authorized, such as when supplied credentials or permissions do not allow the client to connect. This status lets an MQTT client distinguish an authorization failure from successful session establishment and respond appropriately, for example by obtaining valid credentials or contacting an administrator.

“Brokers can send the Last Will and Testament messages on behalf of the client” is also correct. When establishing an MQTT connection, a client may provide a Will message, its topic, retained-message setting, and quality of service. The broker stores that information for the duration of the connection. If the client disconnects unexpectedly—for example, due to loss of network connectivity or failure to send traffic within the negotiated Keep Alive interval—the broker publishes the stored Will message for the client. This allows subscribers to be informed that the client became unavailable. See the MQTT 3.1.1 specification’s CONNACK return-code and Will Message requirements.

QUESTION NO: 4

Refer to the exhibit.

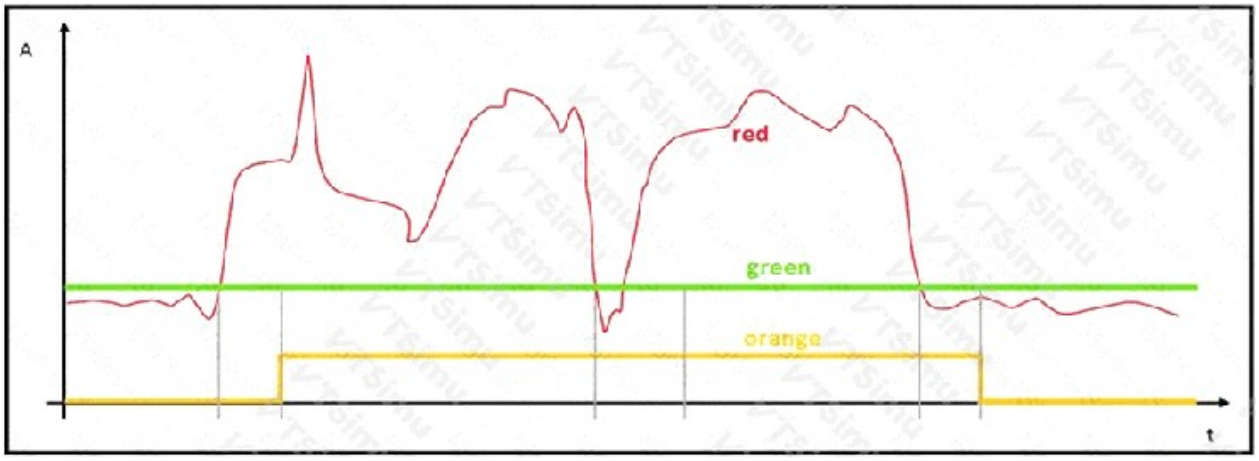

The red line represents the energy consumption of an industrial machine. Which two statements about the graphs are true? (Choose two.)

A. The orange graph is expected to drop to zero for a brief period around the midpoint on the x axis.

B. The red graph contains the raw data points.

C. The green graph is the mean value of the data points.

D. The green graph does not give us any valuable information.

E. The orange graph has three values on the y-axis.

ANSWER: B C

Explanation:

The red graph contains the raw data points because it follows the individual, rapidly changing energy-consumption measurements collected from the industrial machine. Raw time-series telemetry commonly contains short-duration fluctuations caused by changing load, operating state, sampling variation, and other real-world process effects. Retaining that series is important because it represents the actual observed machine behavior over time.

The green graph is the mean value of the data points because it is a smoothed representation of the red telemetry. A mean, often calculated over a moving window in time-series analysis, reduces the visual impact of individual measurement fluctuations and exposes the underlying consumption trend. This makes it useful for operational analysis: an engineer can compare instantaneous consumption against the normal average pattern, identify sustained deviations, and use the trend as a baseline for monitoring or alerting. Cisco IoT architectures emphasize collecting device telemetry and deriving operational insight from data analysis at the edge or in connected platforms; see Cisco's IoT overview. For the statistical basis of a moving mean used to smooth time-series observations, see NIST's moving-average guidance.

QUESTION NO: 5

In a secure software development lifecycle process, what must be created as soon as the architecture is in place?

A. integration plan

B. test plan

C. functional specification document

D. threat modeling

ANSWER: D

Explanation:

threat modeling must be created as soon as the architecture is available because it uses the system's design, trust boundaries, components, data flows, interfaces, and external dependencies to identify security threats before implementation is too far advanced. At this point, the development team can document likely attack paths, assess the impact and likelihood of each threat, define security requirements, and select mitigations that can be incorporated directly into the design. Performing this work early is central to a secure software development lifecycle: architectural decisions such as authentication placement, authorization boundaries, encryption, network segmentation, and service exposure are substantially easier and less costly to address before code is built and deployed. The threat model is also a living security artifact; it should be revisited whenever the architecture, data flows, deployment model, or exposed interfaces change. Cisco's secure-development approach emphasizes building security into products throughout their lifecycle, rather than treating it as a final testing activity. See Cisco's Security Programs overview and Microsoft's practical threat-modeling guidance for the role of architecture-based threat analysis.

QUESTION NO: 6 - (DRAG DROP)

The network team of a large international airport is creating a sensor network on their site using the Cisco LoRaWAN solution. Drag and drop the actions from the left to right needed to implement this project. Not all options are used.

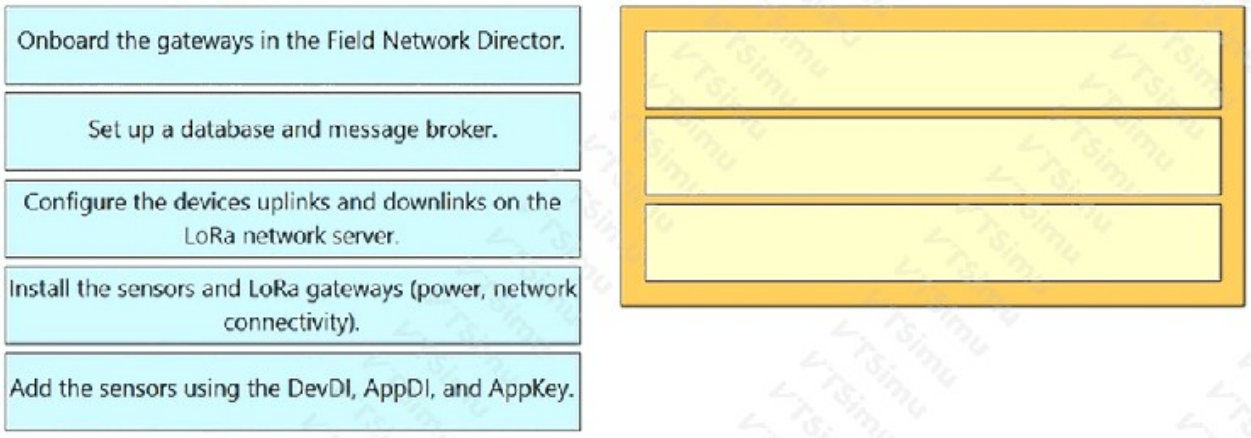

ANSWER:

Explanation:

A Cisco LoRaWAN deployment starts with the field infrastructure because the gateways need power and IP network connectivity before they can participate in the LoRaWAN solution. Installing the sensors and LoRa gateways establishes the radio coverage layer for the airport and provides the backhaul path through which gateway traffic can reach the central management and network-server components. This first step also ensures the gateway is physically available and connected when it is added to the operational environment.

Once the gateways are installed and reachable, they are onboarded in Cisco Field Network Director. Field Network Director is used to manage the gateway fleet and provides the operational control point for gateway onboarding and lifecycle management. Onboarding after installation lets the management system identify the reachable gateways and apply the relevant management configuration. Cisco describes Field Network Director as part of its LoRaWAN solution architecture for managing the field network; see the Cisco LoRaWAN solution overview.

After the radio and gateway-management layers are ready, configuring device uplinks and downlinks on the LoRa network server enables the actual LoRaWAN communications service. Uplinks carry sensor telemetry from the airport devices through the gateways into the solution, while downlinks provide the controlled return path for device configuration or commands. This sequencing creates a usable end-to-end service: deployed connectivity first, managed gateways second, and device traffic configuration last. The result is an operational sensor network in which installed devices can reliably exchange LoRaWAN traffic through the managed gateway infrastructure.

QUESTION NO: 7

Which statement explains the correct position of a PLC within the CPwE architecture?

A. PLCs are always part of Level 0 – process control.

B. PLCs are part of Level 4 – site business planning.

C. PLCs are part of Level 1 – basic control.

D. PLCs are placed at Level 2 – area supervisory control.

ANSWER: C

Explanation:

PLCs are part of Level 1 – basic control. CPwE, which applies the Purdue Enterprise Reference Architecture to converged industrial networks, places the control devices that directly execute automated control logic in the basic-control layer. A PLC receives input signals from field devices, evaluates its programmed logic, and issues outputs that operate equipment. This makes it the core controller for the process rather than the process equipment itself or a supervisory/business system.

Within the Purdue model used by CPwE, Level 0 contains the physical process and the field instruments, sensors, and actuators interacting with that process. Level 1 contains the intelligent industrial controllers that provide basic, real-time control, including PLCs and similar control platforms. This placement is important when designing industrial zones, segmentation, and communication paths because Level 1 control assets have distinct operational and security requirements. Cisco's industrial automation design guidance describes CPwE as using the Purdue model to organize manufacturing systems and their networked functions. See Cisco Industrial Automation Design Zone and Cisco Manufacturing Solutions.

QUESTION NO: 8

An edge gateway publishes telemetry over MQTT through a cellular connection. The connection is unreliable, and the application must continue delivery after it reconnects.

Which two MQTT client settings support this requirement? (Choose two.)

A. Use a persistent session.

B. Use Quality of Service 1 for telemetry that must be delivered.

C. Use a clean session for every reconnect.

D. Disable MQTT keep-alive processing.

E. Use Quality of Service 0 for messages that require acknowledgement.

ANSWER: A B

Explanation:

Use a persistent session is correct because the broker can retain session state, including subscriptions and eligible queued messages, when the client disconnects. On reconnect, the client can resume the existing session rather than creating a completely new one.

Use Quality of Service 1 for telemetry that must be delivered is also correct. QoS 1 requires acknowledgement of a published message and supports redelivery when an acknowledgement is not received. The application must tolerate possible duplicate delivery because QoS 1 provides at-least-once delivery semantics. See the OASIS MQTT Version 5.0 specification.