

TIBCO Enterprise Message Service 8 Certification Exam

Tibco TCP-EMS8

Version Demo

Total Demo Questions: 19

Total Premium Questions: 193

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

Which two statements are true about configuring security for a fault-tolerant pair that does not share state? (Choose two.)

Options are :

- A. When SSL is deployed, both servers must use the same certificate(s).
- B. acl.conf settings can differ between servers.
- C. When SSL is deployed, servers can use different certificate(s), but must use the same root certificate authority.
- D. All access control settings should be the same on both servers.

ANSWER: C D

Explanation:

When an EMS fault-tolerant pair does not use shared state, security-related configuration must be kept consistent manually so that a client has the same authorization outcome after failover. In particular, all access-control settings should be the same on both servers. This includes the effective user, group, and permission rules used to authorize connections, administration, and destinations. Matching these settings prevents a client that was permitted on the active server from being denied, or granted different rights, after reconnecting to the standby server.

For SSL, each member of the pair can have its own server certificate. This is appropriate because the servers are separate TLS endpoints and can have distinct host identities. Their certificates must, however, chain to the same trusted root certificate authority so that clients can consistently validate either server when a failover occurs. Client trust configuration can therefore recognize both members of the pair through a common trust anchor while each server retains its own certificate and private key.

These practices are part of maintaining equivalent security behavior across independently configured FT servers. See the [TIBCO Enterprise Message Service documentation](#) and the [EMS 8.5 documentation library](#).

QUESTION NO: 2

Which type of wiring can be done at deployment time?

Options are :

- A. from a service in one SCA composite to a reference of another SCA composite
- B. from a reference in one SCA composite to a service of another SCA composite
- C. from a component reference to a component service in the same SCA composite
- D. from a component service to a component reference in the same SCA composite

ANSWER: B

Explanation:

from a reference in one SCA composite to a service of another SCA composite is the deployment-time wiring model. An SCA composite exposes references as its required dependencies and services as its provided capabilities. When separately developed composites are assembled into an application or deployed to an SCA runtime, a composite reference can be bound to a compatible service supplied by another composite. This lets the deployer select the target implementation, endpoint, or environment-specific provider without changing the source composite design.

This separation is useful when the same application archive must be promoted through development, test, and production environments, where the appropriate provider may differ. The reference carries the consumer-facing contract, while the target composite service supplies the matching contract. Deployment configuration then establishes the connection between them, allowing independently packaged composites to be integrated at runtime.

The SCA Assembly specification describes composite-level services and references and their assembly/wiring semantics: [OASIS SCA Assembly Model Specification](#). TIBCO ActiveMatrix BusinessWorks documentation is available from [TIBCO ActiveMatrix BusinessWorks documentation](#).

QUESTION NO: 3

Which two objects are referred to as EMS administered objects?

Options are :

- A. A & B
- B. connection factories
- C. groups
- D. topics

ANSWER: B D

Explanation:

In TIBCO Enterprise Message Service, the JMS administered objects relevant here are connection factories and destinations. Therefore, **connection factories** and **topics** are the correct selections. A connection factory is a configured administrative object that clients use to create JMS connections; it centralizes connection settings rather than requiring those settings to be hard-coded in each client application. A topic is a JMS destination used with the publish-and-subscribe messaging model. EMS administrators configure and manage topics on the server, including properties such as permissions and routing-related behavior where applicable.

JMS defines connection factories and destinations as administered objects because they are created and maintained by an administrator, then made available to client applications through naming or direct configuration. In EMS, destinations include both queues and topics, so a topic fits the administered-object category. This division supports portable applications: client code obtains a connection factory and destination, then uses standard JMS APIs to connect, publish, or consume messages without embedding server-specific administrative details. See the [TIBCO Enterprise Message Service documentation](#) and the [Jakarta Messaging specification](#) for the JMS administered-object model.

QUESTION NO: 4

Which two statements about channels are true? (Choose two.)

Options are :

- A. All destinations for a channel must use different server connections.
- B. A Channel can contain many destinations.
- C. A destination can contain many channels.
- D. All destinations for a channel use the same server connection.

ANSWER: B D

Explanation:

In TIBCO Enterprise Message Service, a channel is a logical grouping used to associate one or more destinations with a server connection. Therefore, *A Channel can contain many destinations*. A channel is not limited to a single queue or topic; multiple destinations can be grouped where they are to be handled through that channel's configured connection.

Also, *All destinations for a channel use the same server connection*. The server connection is an attribute of the channel itself, so destinations included in that channel share that configured connection. This design provides a consistent communication path for the grouped destinations and allows the administrator to manage their connection behavior as one channel definition.

These channel concepts are part of EMS administration and server-to-server connectivity configuration. Consult the TIBCO Enterprise Message Service documentation library for the applicable EMS 8 Administration guide and configuration terminology: [TIBCO Enterprise Message Service documentation](#). Additional official product and documentation access is available through the [TIBCO Enterprise Message Service product page](#).

QUESTION NO: 5

What are three valid acknowledgment modes in TIBCO Enterprise Message Service?

(Choose three.) Options are :

- A. DUPS_OK_ACKNOWLEDGE
- B. IMPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE
- C. EXPLICIT_CLIENT_ACKNOWLEDGE
- D. NO_ACKNOWLEDGE
- E. IMPLICIT_CLIENT_ACKNOWLEDGE

ANSWER: A C D

Explanation:

DUPS_OK_ACKNOWLEDGE, EXPLICIT_CLIENT_ACKNOWLEDGE, and NO_ACKNOWLEDGE are valid TIBCO Enterprise Message Service acknowledgment modes. DUPS_OK_ACKNOWLEDGE is the JMS mode that permits the provider to acknowledge messages lazily, which can improve throughput but allows duplicate delivery under some failure conditions. This is appropriate when an application can safely tolerate and handle duplicate messages.

EXPLICIT_CLIENT_ACKNOWLEDGE is a TIBCO EMS extension that gives an application explicit control over acknowledgment behavior. It is useful when message processing must complete successfully before the application confirms receipt, enabling more precise recovery behavior for consumers.

NO_ACKNOWLEDGE is also an EMS extension. It avoids acknowledgment processing, reducing acknowledgment overhead in situations where reliable redelivery is not required and the application accepts the resulting delivery semantics. These modes are documented as part of EMS session and message-consumer acknowledgment support. See the [TIBCO Enterprise Message Service documentation](#) and the JMS definition of [DUPS_OK_ACKNOWLEDGE](#).

QUESTION NO: 6

Which character is valid separator in the URL for a client to connect to a fault-tolerant EMS server pair?

Options are :

- A. .
- B. &
- C. ;

D. |

E. ,

ANSWER: E

Explanation:

The comma (,) is the valid separator for listing the two server URLs that make up a TIBCO Enterprise Message Service fault-tolerant pair. A client connection URL contains each server endpoint as a complete URL and separates the endpoints with a comma, for example: `tcp://ems-primary.example.com:7222,tcp://ems-secondary.example.com:7222`. Providing both endpoints lets the EMS client initially connect to an available member of the pair and reconnect appropriately after a failover. The individual URLs can use the applicable EMS transport scheme, such as `tcp` or `ssl`, while the comma remains the delimiter between the endpoint URLs. This client-side URL list is distinct from the fault-tolerant configuration performed on the EMS servers themselves; it supplies the client with the alternate location needed to maintain connectivity when the active server changes. TIBCO's Enterprise Message Service documentation describes connection URLs and fault-tolerant server support in its product documentation. See the [TIBCO Enterprise Message Service documentation portal](#) and the [EMS 8.6 documentation collection](#).

QUESTION NO: 7

In Cache Object Management settings, what does the property cache Agent Quo rum specify?

Options are :

- A. The size of the limited cache
- B. The maximum size of the object table cache
- C. The minimum number of storage-enabled nodes that must be active in the cluster at startup
- D. The number of members in the cluster that hold the backup data

ANSWER: C

Explanation:

The minimum number of storage-enabled nodes that must be active in the cluster at startup is correct. The cache-agent quorum is a startup-safety threshold for a distributed cache cluster. Storage-enabled nodes are the members that maintain persisted or replicated cache data and participate in making the cluster's data available. By requiring a defined minimum number of these members before startup proceeds, the quorum setting helps ensure that the cache is brought online only when sufficient storage membership is present to establish a reliable cluster state. This prevents a cache from starting with an insufficient subset of the expected storage participants, which could otherwise risk serving an incomplete view of distributed data after a restart or failure event. The value is therefore expressed as a required count of active storage-enabled cluster nodes, not as a memory or object-table capacity and not as a description of backup-data placement. In an EMS deployment using cache object management, administrators should select the quorum value in accordance with the cluster's planned number of storage-enabled members and its availability requirements. TIBCO Enterprise Message Service documentation is available through the [TIBCO Enterprise Message Service documentation portal](#); related distributed-cache concepts are documented in the [TIBCO ActiveSpaces documentation portal](#).

QUESTION NO: 8

How do you achieve high availability for a TIBCO Active matrix Java component providing a service with a SOAP/HTTP binding type?

Options are :

- A. deploy the service in the system node so the Active matrix infrastructure can manage availability
- B. deploy the service across multiple nodes and place a load balancer in front of the service availability

- C. specify virtualization upon service deployment allowing Active matrix to transparently route requests
- D. create a single node that spans multiple hosts to provide resiliency in the event of host failure

ANSWER: B

Explanation:

For a SOAP/HTTP service, high availability is achieved by deploying equivalent service instances on multiple ActiveMatrix nodes and placing an HTTP-capable load balancer in front of those instances. The load balancer presents a stable endpoint to consumers and distributes incoming SOAP/HTTP requests among the available service providers. If a host or node running one provider instance becomes unavailable, the load balancer can stop directing traffic to that instance and continue sending requests to healthy instances. This protects service access from a single node or host failure while also allowing requests to be spread across the deployment.

This approach fits the request/response nature of HTTP bindings: clients connect to a virtual service address managed by the load-balancing tier rather than relying on one particular runtime node. The ActiveMatrix application must therefore be deployed redundantly, with each deployed instance able to provide the same service contract and behavior. Appropriate health checks and load-balancer configuration are also needed so that only ready service instances receive traffic. TIBCO ActiveMatrix Service Grid documentation describes deployment and runtime administration concepts for distributed nodes, while the product documentation portal provides the relevant release documentation: [TIBCO ActiveMatrix Service Grid documentation](#). For HTTP load-balancing concepts and health-based routing, see the [Load Balancing 101 overview](#).

QUESTION NO: 9

Which three commands will update EMS configuration files when an exclusive commit command is issued from EMS? (Choose three.) Options are :

- A. set server password=password
- B. set server track_message_ids=enabled
- C. set server routing=enabled
- D. set server client_trace=enabled
- E. set server authorization=enabled

ANSWER: B C D

Explanation:

`set server track_message_ids=enabled`, `set server routing=enabled`, and `set server client_trace=enabled` are server configuration changes that can be made in an exclusive administration session and then persisted with `commit`. The exclusive mode provides controlled access while the administrator modifies configuration, and committing writes the eligible in-memory server-property changes to the applicable EMS configuration files. Enabling message-ID tracking controls the server's tracking of message identifiers, which is used for administration and diagnostic capabilities. Enabling routing activates EMS routing behavior for message exchange across configured routes. Enabling client tracing activates tracing information associated with client activity, supporting operational troubleshooting. These are configuration-level server properties rather than merely transient administrative observations, so an exclusive commit persists their enabled state and allows it to survive a server restart. TIBCO's EMS documentation describes the administration commands, exclusive mode, and commit behavior in the EMS administration reference; see the [TIBCO Enterprise Message Service documentation portal](#) and the [EMS 8.6 documentation collection](#).

QUESTION NO: 10

Which two statements are true about the Delegation pattern? (Choose two.)

Options are :

- A. The service consumer does not receive the work result.
- B. The service consumer receives confirmation that the work was completed successfully.
- C. The service consumer receives confirmation that the request was accepted.
- D. The service provider sends the work result to the service consumer.

ANSWER: A C

Explanation:

In the Delegation pattern, a service consumer hands work to a service provider for asynchronous processing rather than waiting for a business-level response containing the completed work result. The immediate interaction is therefore an acknowledgment of the handoff: the service consumer receives confirmation that the request was accepted. Acceptance confirms that the provider has taken responsibility for the request, not that the requested business operation has already finished successfully.

Because processing occurs independently after the request is delegated, the service consumer does not receive the work result as part of this messaging exchange. This pattern is useful when the caller must be decoupled from the duration of the work and only needs assurance that the request entered the provider's processing flow. TIBCO Enterprise Message Service supports this type of asynchronous decoupling through its queue-based messaging capabilities, in which senders and receivers do not need to process a message at the same time. See the [TIBCO Enterprise Message Service documentation](#) and the [TIBCO Enterprise Message Service 8.6 documentation](#) for EMS messaging and destination concepts.

QUESTION NO: 11

Which statement is true about acknowledgment modes for non-transacted sessions?

Options are :

- A. The acknowledge mode is set by the administrator of the EMS server.
- B. Only EMS extensions of the acknowledge modes can be set by the administrator of the EMS server.
- C. The acknowledge mode is set on the message consumer side.
- D. The acknowledge mode is set on the message producer side.

ANSWER: C

Explanation:

The acknowledge mode is set on the message consumer side. For a non-transacted JMS session, the client application selects the session's acknowledgment behavior when it creates the session or JMS context. That setting controls how messages delivered to consumers associated with that session are acknowledged and when the messaging provider may regard a delivered message as successfully consumed. The available standard modes include automatic acknowledgment, client acknowledgment, and duplicate-tolerant automatic acknowledgment. TIBCO Enterprise Message Service implements the JMS programming model, so this is a client-side session characteristic rather than an EMS server administrative setting. In practical terms, the consuming application chooses the reliability and acknowledgment semantics appropriate for its message processing: automatic acknowledgment can acknowledge according to the provider's delivery rules, whereas client acknowledgment allows application code to explicitly acknowledge received messages. A producer sends messages but does not establish the acknowledgment mode used for their later consumption. The Jakarta Messaging Session API describes acknowledgment mode as a parameter used when creating a non-transacted session and defines the associated client behavior. See the [Jakarta Messaging Session API](#) and TIBCO's [TIBCO Enterprise Message Service documentation](#).

QUESTION NO: 12

When wrapping a back end system, what does a Web Service always provide that the native system interface does not?

Options are :

- A. access independent of implementation technology
- B. a common data model
- C. standardized service operation semantics
- D. an HTTP-based interface

ANSWER: A

Explanation:

access independent of implementation technology is correct because a web service exposes a service contract at an interoperability boundary rather than requiring consumers to use the back-end system's native client libraries, programming language bindings, proprietary protocol, or platform-specific APIs. The wrapper translates between the service-facing interface and the underlying application, database, messaging system, or legacy technology. As a result, a consumer can invoke the exposed capability using the published service description and supported message/protocol standards without being coupled to the implementation details of the native system.

This separation is a core service-oriented design principle: the service contract provides an abstraction over the implementation, allowing the provider to evolve or replace its internal technology while preserving the interface expected by consumers. In web-service architectures, a description such as WSDL identifies the operations, messages, bindings, and endpoints needed by a client, making the interaction contract explicit and technology-neutral at the integration boundary. The World Wide Web Consortium describes web services as software systems designed to support interoperable machine-to-machine interaction, and its Web Services Architecture also emphasizes interfaces and message exchanges independently of a specific implementation. See the [W3C Web Services Architecture](#) and the [WSDL 2.0 specification](#).

QUESTION NO: 13

Which section is required in a JMS message? Options are :

- A. Header
- B. Footer
- C. Properties
- D. Body

ANSWER: A

Explanation:

Header is correct. The Java Message Service message model consists of three conceptual parts: a header, an optional set of properties, and a body. Every JMS message has a header, which contains standardized JMS-defined fields used by the provider and applications to control and identify message delivery. These fields include identifiers and delivery-related metadata such as `JMSMessageID`, `JMSTimestamp`, `JMSDestination`, `JMSDeliveryMode`, and `JMSPriority`. Although some individual header-field values may be provider-assigned or have default values, the header itself is an intrinsic part of each JMS message.

In TIBCO Enterprise Message Service, JMS clients use this standard JMS message structure when sending messages to queues or topics. The EMS documentation describes message headers as JMS-defined information associated with every message and distinguishes them from application-defined message properties and the message payload. Understanding the mandatory header is important when configuring selectors, tracing messages, setting delivery characteristics, and correlating request/reply interactions. See the JMS message API documentation at [Oracle JMS Message interface](#) and the TIBCO EMS documentation portal at [TIBCO Enterprise Message Service documentation](#).

QUESTION NO: 14

TIBCO Enterprise Message Service Certification Exam Set 1

Which two datastores are supported by EMS servers? (Choose two.)

Options are :

data grid

dbstore

A,D

dbstore

mstore

Answer : A,D

Which two statements are true about EMS tracing and logging capabilities? (Choose two.)

Options are :

The location of log file is set in tibemsd.conf file via the log_trace parameter.

B,D

Message tracing applies only to headers.

Message tracing can be applied at destination and message levels.

Access control violations can be logged to console or log file.

Answer : B,D

Which statement is true about configuration file usage for an EMS fault tolerant configuration?

Options are :

Primary and Secondary EMS servers must share all configuration files excluding the stores.conf file.

Primary and Secondary EMS servers must share all configuration files.

Primary and Secondary EMS servers should share the tibemsd.conf file.

C

Primary and Secondary EMS servers can share all configuration files excluding the tibemsd.conf file.

Answer : C

Which command displays message count delivered to a queue receiver?

Options are :

Show consumers

Show connections

Show users

Show receivers

Answer : Show consumers

What are three valid acknowledgment modes in TIBCO Enterprise Message Service?

(Choose three.)

Options are :

DUPS_OK_ACKNOWLEDGE

IMPLICIT_CLIENT_DUPS_OK_ACKNOWLEDGE

EXPLICIT_CLIENT_ACKNOWLEDGE

NO_ACKNOWLEDGE

A,B,E

IMPLICIT_CLIENT_ACKNOWLEDGE

Answer : A,B,E

Which Administration Tool command creates the queue without reporting any errors?

Options are :

Create queue test.queue store=\$sys.fail-safe, secure,local

Create queue test.queue store=\$sys.fail-safe,trace=header

Create queue test.queue secure,global,store=\$sys.fail-safe

Create queue test.queue secure,store=\$sys.fail-safe,channel=multicast-l

Answer : Create queue test.queue secure,global,store=\$sys.fail-safe

Which parameter must be configured for primary and secondary EMS to operate in fault-tolerant mode?

Options are :

ft_passive parameter in the tibemsd.conf file

store parameter in the store, conf file

server parameter in the tibemsd.conf file

ft_factory in the factories.conf file

Answer : server parameter in the tibemsd.conf file

Which section is required in a JMS message?

Options are :

Header

Footer

Properties

Body

Answer : Header

The customer has identified an interface called "New Customer". When a new customer has been created in the Customer Care Application, the customer data must be distributed to other applications such as billing, payments, etc.

What is the most efficient mechanism to implement the message delivery?

Options are :

Publish/subscribe using queues

Point-to-point using queues

Publish/subscribe using topics

Point-to-point using topics

Answer : Point-to-point using queues

A topic, conf.file contains the following entry: mycompany. *. hr. > secure

Which two topics will inherit the secure property? (Choose two.)

Options are :

mycompany.south.hr.vacation.approval

mycompany.south.eng.hr.vacation

A,F

mycompany.south.hr

mycompany.hr.vacation

mycompany.south.region.hr.vacation

mycompany_office.hr.vacation

mycompany.south.hr.vacation

Answer : A,F

Which two statements are true about topics and queues? (Choose two.)

Options are :

A sender can send to a wildcard queue.

A consumer can listen for wildcard queue names.

C,E

A queue can only have one exclusive connected consumer at a time.

A publisher can publish to a wildcard topic.

A consumer can listen for wildcard topic names.

A queue can only have one consumer per message.

Answer : C,E

What are the three options for the user_auth parameter in EMS? (Choose three.)

Options are :

A,C,D

local

ldap

ssl

jaas

x.509

Answer : A,C,D

A bridge has been configured from topic T1 to topic T2.

Which two statements are true if a message producer sends messages to T1 within a transaction? (Choose two.)

Options are :

The transaction can succeed if the producer only has permission to send to topic T1.

The transaction will fail if the producer does not have permission to send to topic T2.

A,B

The bridge needs to be configured to support transactions.

All messages sent across the bridge are part of the transaction.

Answer : A,B

Which two statements are true about configuring security for a fault-tolerant pair that does not share state? (Choose two.)

Options are :

When SSL is deployed, both servers must use the same certificate(s).

B,D

acl.conf settings can differ between servers.

When SSL is deployed, servers can use different certificate(s), but must use the same root certificate authority.

All access control settings should be the same on both servers.

Answer : B,D

Which two statements are true about the listen parameter? (Choose two.)

Options are :

Multiple listen parameters can be set in the TIBCO EMS configuration file.

A,B

The listen parameter is set in the tibemsd.conf file.

TCP is the only protocol supported in the listen parameter.

The listen parameter is set in the transports, conf file.

Answer : A,B

Which statement is true about acknowledgment modes?

Options are :

They determine the way the EMS server sends confirmations to message producers.

They determine the way message consumers send confirmations to message producers

They determine the way message consumers send confirmations to the EMS server.

They determine the way the EMS server sends confirmations to message consumers.

Answer : They determine the way message consumers send confirmations to the EMS server.

What are three valid delivery modes in EMS? (Choose three.)

Options are :

NON PERSISTENT

PERSISTENT

TRANSACTIONAL

GUARANTEED

RELIABLE_DELIVERY

A,C,E

Answer : A,C,E

What are two characteristics of multicasting? (Choose two.)

Options are :

The EMS server uses TCP to transmit messages to receivers,

The EMS server uses PGM to transmit messages to receivers.

B,C

It reduces network traffic generated between EMS server and EMS client machines.

The producer can choose to use PGM or TCP to send messages to the EMS server,

Multicasting is used in messaging scenarios where high bandwidth delivery of real-time data is a requirement,

Answer : B,C

Which Context.PROVIDER_URL is used to look up EMS administered objects stored in the EMS server running on host EMSHOST on port 7222?

Options are :

tibjmsnaming://EMSHOST:7222

emsnaming://EMSHOST:7222

jndi://EMSHOST:7222

ems://EMSHOST:7222

tcp: //EMSHOST: 7222

Answer : jndi://EMSHOST:7222

User Bob is in the ENG group with the following permissions:

Larger image

Which statement is true for user Bob and the topic quote.tibx.us.new?

Options are :

The EMS server starts and Bob can publish and subscribe to quote.tibx.us.new.

The wildcard topics collide, so the EMS server fails to start.

The EMS server starts and Bob overrides all the ENG group permission, so Bob cannot subscribe to quote.tibx.us.new.

The EMS server starts and Bob inherits all permissions from the ENG group, so Bob cannot publish to quote.tibx.us.new.

Answer : The EMS server starts and Bob can publish and subscribe to quote.tibx.us.new.

What must you do to set an EMS Java client application to attempt reconnection 10 times with 1000 ms delay between attempts to the EMS server during fault tolerance switch or when a connection is being terminated by the network? (Choose two.)

Larger image

Options are :

Option A

Option B

Option D

Option C

A,D

Answer : A,D

A dynamic queue named q1 has been created by an EMS client.

What will be the result of executing the following command: create queue q1?

Options are :

The command will fail saying the queue already exists.

The command will succeed and dynamic queue will be renamed.

The command will succeed and queue will be promoted to static.

The command will fail if queue has connected receivers.

Answer : The command will succeed and queue will be promoted to static.

Which statement is true about message compression with EMS?

Options are :

Body is the only message part that is compressed

Properties are compressed, but headers are not.

Properties and headers are also compressed.

Headers are compressed, but properties are not.

Answer : Body is the only message part that is compressed

What are two valid delivery modes in EMS? (Choose two.)

Options are :

B,E

GUARANTEED

DURABLE

PERSISTENT

RELIABLE_DELIVERY

EXPLICIT

Answer : B,E

Producer A has published a message with expiration time as 30 seconds to a queue with the property JMS_TIBCO_PRESERVE_UNDELIVERED set to true.

What happens when the message expires?

Options are :

The message is moved to Ssys.expired queue.

The message is deleted from the undelivered queue.

The message is moved to Ssys.undelivered queue.

The message is completely deleted from the server.

Answer : The message is moved to Ssys.undelivered queue.

If authorization is enabled in the tibemsd.conf file and the secure property is set for a destination, when are permissions checked? (Choose two.)

Options are :

A,B

When creating non-durable subscription for wildcard topics

When performing a JNDI lookup of a destination

When a producer sends a message to the destination

When attempting to become a consumer for a queue

Answer : A,B

Which parameter is used to preallocate disk space for the store file?

Options are :

file_minimum

store_min_size

store_minimum

file_min_size

file_truncate

Answer : file_minimum

What are three limitations in the use of wildcards by EMS clients? (Choose three.)

Options are :

Clients cannot subscribe to wildcard topics,

Clients cannot send to wildcard queues.

Clients cannot use wildcards for dynamic topics.

B,C,E

Clients cannot receive from wildcard queues.

Clients cannot publish to wildcard topics.

Answer : B,C,E

Which event causes EMS to return a delivered but unacknowledged queue message to be available to another receiver?

Options are :

The client application terminates the connection.

Ssys.redelivery.delay queue is purged.

redeliverydelay property is set to zero.

Ssys.undelivered queue is disabled.

Answer : The client application terminates the connection.

Click the Exhibit button.

An EMS server is configured with a bridge from the topic ORDER.NEW to the queue SALES.ORDERS,Q1. You enable flow control at the server and destination levels (order, new and SALES, ORDERS, Q1) .The source destination order, new blocks publishers.

Which two statements are true about flow control? (Choose two.)

Options are :

- A. It does not discard the messages.
- B. It blocks the sender but does not send back an error.
- C. It discards the earliest messages (FIFO).
- D. It discards the latest messages (FILO).

ANSWER: A B

Explanation:

When EMS flow control is configured on a source destination to block publishers, it provides back-pressure rather than a discard policy. The server temporarily suspends the producer's send activity once the applicable flow-control limit is reached. Consequently, already accepted messages are retained; the action is not to remove either the oldest or newest messages. This is why *It does not discard the messages.* is correct.

The corresponding producer-side behavior is described by *It blocks the sender but does not send back an error.* A send call waits while flow control remains active and proceeds when the server can again accept traffic. This permits the application to slow naturally in response to constrained server or destination capacity, preserving message delivery rather than treating the condition as a send failure. In the stated scenario, the source topic is explicitly configured to block publishers, so this blocking behavior applies at the point where producers publish to the bridged topic.

TIBCO's Enterprise Message Service documentation describes destination and server flow control as a mechanism for regulating producer traffic when resource thresholds are exceeded. See the [TIBCO Enterprise Message Service documentation portal](#) and the [EMS 8.6 documentation](#).

QUESTION NO: 15

What must you do to set an EMS Java client application to attempt reconnection 10 times

with 1000 ms delay between attempts to the EMS server during fault tolerance switch or when a connection is being terminated by the network? (Choose two.)

Options are :

- A. Call `setReconnectAttempt(10)` on the `TibjmsConnectionFactory` before creating the connection.

- B. Set the EMS server's connection timeout to 10 seconds.
- C. Configure the client to use a 10,000 ms reconnect delay.
- D. Call `setReconnectDelay(1000)` on the `TibjmsConnectionFactory` before creating the connection.

ANSWER: A D

Explanation:

An EMS Java client controls its reconnection behavior through properties of `TibjmsConnectionFactory`. Calling `setReconnectAttempt(10)` configures the client connection factory to make up to ten reconnection attempts after a recoverable connection loss, including a loss encountered while an EMS fault-tolerant pair switches active servers. The reconnect-attempt setting is a client-side setting; it determines how long the client continues trying to re-establish its connection rather than immediately reporting the connection failure to the application.

Calling `setReconnectDelay(1000)` sets the interval between those attempts to 1,000 milliseconds. Used together on the same connection factory before creating the connection, these settings produce the requested behavior: ten retries with a one-second pause between retries. EMS client reconnection supports continuity across temporary network interruptions and server failover when the connection factory has been configured with an appropriate server URL and fault-tolerant connection information. TIBCO documents the Java client connection-factory reconnection controls in the EMS API and client programming documentation. See the [TIBCO Enterprise Message Service documentation](#) and the [EMS Java API documentation](#).

QUESTION NO: 16

What is the Recovery Point Objective (RPO) of a system?

Options are :

- A. the point in time to which the application state must be restored
- B. the complete collection of data that must be restored
- C. a formal contract detailing the critical system to be restored
- D. the time period that can elapse before normal operation must be restored

ANSWER: A

Explanation:

The point in time to which the application state must be restored is correct. A Recovery Point Objective (RPO) defines the acceptable amount of data loss measured in time: after a disruption, the recovered system must contain data at least as current as the specified recovery point. For example, an RPO of 15 minutes means the organization can accept losing no more than 15 minutes of transactions or message-state changes, so backups, replication, or other recovery mechanisms must provide a recoverable state within that interval.

In a TIBCO Enterprise Message Service deployment, the RPO is especially relevant to durable server state, including persisted messages and configuration or store data needed to resume service consistently. The selected recovery point guides the required frequency and design of backups, store replication, or fault-tolerant arrangements. It is a data-freshness target rather than a target for the duration of service restoration. The distinction is important in disaster-recovery planning because a recovery solution must satisfy both the acceptable point to which data is restored and the operational timing requirements established for the service. See the [NIST definition of Recovery Point Objective](#) and IBM's [RTO versus RPO overview](#).

QUESTION NO: 17

Which parameter must be configured for primary and secondary EMS to operate in fault-

tolerant mode?

Options are :

- A. ft_passive parameter in the tibemsd.conf file
- B. store parameter in the store, conf file
- C. server parameter in the tibemsd.conf file
- D. ft_factory in the factories.conf file

ANSWER: C

Explanation:

The `server` parameter in the `tibemsd.conf` file is required because it defines the EMS server's listening URL or URLs. In a fault-tolerant deployment, the primary and secondary servers must be configured with the appropriate server endpoints so that each EMS instance has its network identity and can participate correctly in the fault-tolerant pair. The URL configuration is also the basis on which clients and peer EMS components can address the server service. Consequently, correctly setting `server` is a fundamental part of configuring both members of an EMS fault-tolerant setup, alongside the other fault-tolerance and persistent-state settings that a complete deployment requires.

TIBCO EMS uses `tibemsd.conf` as the principal server configuration file, and its server-configuration documentation describes the `server` setting as the mechanism for specifying the protocols, hosts, and ports on which the EMS server accepts connections. Review the EMS documentation for the applicable product release when preparing the paired configurations, as exact examples and supported settings can differ by version. See the [TIBCO Enterprise Message Service 8.6 documentation](#) and the [TIBCO Enterprise Message Service 8.5 documentation](#).

QUESTION NO: 18

You have a routed queue named ORDER. NEW defined between the US-EMS (home queue) and MEXICO-EMS servers

What must be done for messages to route properly? Options are :

- A. Send a message to the queue named ORDER, NEW
- B. Send a message to the queue named order. NEW@US-EMS
- C. Define ORDER.NEW@home in the US-EMS server's queues, conf file
- D. Send a message to the routed queue named order. NEW@MEXICO-EMS

ANSWER: A

Explanation:

Sending a message to the queue named ORDER, NEW is the required action. In a TIBCO Enterprise Message Service routed-queue configuration, applications use the routed queue's ordinary destination name when they produce messages. Here, US-EMS is identified as the home server, so a producer connected to US-EMS sends to the home queue, ORDER.NEW. EMS then applies the routed-queue configuration and forwards messages to the corresponding routed queue on MEXICO-EMS as necessary. The application does not need to address a server-qualified destination or create an additional local queue definition to cause routing; routing is an EMS server responsibility defined by the routed-queue setup. Using the configured queue name also preserves the normal JMS programming model: a producer creates or looks up the queue destination and sends its message, while the EMS servers manage inter-server delivery, persistence behavior, and route availability. Queue and destination naming should match the name defined in the EMS configuration, including its punctuation and case. TIBCO's Enterprise Message Service documentation provides the product documentation set and administration material covering queues, routes, and server configuration: [TIBCO Enterprise Message Service documentation](#). The JMS destination model used by EMS is also described in the [JMS Queue API documentation](#).

QUESTION NO: 19

Which two statements are true about EMS management of persistent messages? (Choose two.) Options are :

- A. Persistent messages published to a topic are always written to disk,
- B. Persistent messages sent to a queue are written to disk only if that queue has at least one durable consumer or one consumer with a fault-tolerant connection to the EMS server.
- C. Persistent messages published to a topic are written to disk only if that topic has at least one durable subscriber or one subscriber with a fault-tolerant connection to the EMS server.
- D. Persistent messages sent to a queue are always written to disk.

ANSWER: B C

Explanation:

Persistent messages sent to a queue are written to disk only if that queue has at least one durable consumer or one consumer with a fault-tolerant connection to the EMS server. EMS can optimize persistent-message handling when all queue consumers are currently connected through non-fault-tolerant connections, because those consumers cannot resume consumption after a server failure. When durable or fault-tolerant consumption is present, EMS must retain the message on disk so it remains available across a server failure or connection failover.

Persistent messages published to a topic are written to disk only if that topic has at least one durable subscriber or one subscriber with a fault-tolerant connection to the EMS server. A durable subscriber requires EMS to preserve publications while the subscriber is disconnected, and a fault-tolerant subscriber must be able to continue correctly after failover. Disk persistence supplies the recovery guarantee needed for both cases. This behavior reflects EMS's distinction between the persistent delivery mode requested by a producer and the server's need to retain a message for recoverable consumers. See the [TIBCO Enterprise Message Service documentation](#) and the [EMS 8.6 documentation set](#).