

Oracle WebLogic Server 12c: Advanced Administrator II

Oracle 1z0-134

Version Demo

Total Demo Questions: 17

Total Premium Questions: 172

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Configuring WebLogic Server Clusters	8
Topic 2, Configuring Whole Server Migration	16
Topic 3, Configuring WebLogic Server for High Availability	30
Topic 4, Configuring Java Message Service	37
Topic 5, Configuring WebLogic Server Security	16
Topic 6, Monitoring WebLogic Server	8
Topic 7, Tuning WebLogic Server	10
Topic 8, Troubleshooting WebLogic Server	6
Topic 9, Configuring the WebLogic Server Diagnostic Framework	13
Topic 10, Mix Questions	28
Total	172

QUESTION NO: 1

A client application asynchronously receives the message sent to a JMS Queue. The client is going to send a response to the sender.

Which two JMS Header fields from the original message are used to send the response? (Choose two.)

- A. JMSResponseMode
- B. JMSTimeStamp
- C. JMSDeliveryMode
- D. JMSCorrelationID
- E. JMSReplyTo

ANSWER: D E

Explanation:


JMSCorrelationID and **JMSReplyTo** support the standard JMS request/reply messaging pattern. The requester places a destination in the original message's `JMSReplyTo` header. A consumer that processes the request uses that destination as the target when it creates and sends the response, allowing the reply to return to the requester through the intended queue or topic.

`JMSCorrelationID` supplies the identifier needed to associate a response with its specific request, particularly when several requests are in flight concurrently. The responding application normally sets the response message's correlation ID to the correlation ID from the received request. Where an application uses the request message ID as its correlation value, it instead sets the response correlation ID to that request's `JMSMessageID`. This correlation mechanism enables the original sender to identify which outstanding request the asynchronously received response fulfills.

The JMS API defines `getJMSReplyTo()` as the destination to which a reply should be sent and defines `JMSCorrelationID` as a provider/client-supplied value used to link messages. See the [Oracle JMS Message API documentation](#) and the [JMS request/reply documentation](#).

QUESTION NO: 2

Examine the domain diagram:



Assume all servers are up and running and you have just started an offline WLST interactive session. Assume further that the username/password “weblogic/Welcome1” maps to a valid administrator.

You type two lines of WLST code:

```
connect('weblogic','Welcome1','host02.acme.com:7101')
```

```
print cmo.getType() + ':' + cmo.getName()
```

What will be printed? (Choose the best answer.)

- A. Domain:TestDomain
- B. DomainRuntime:server03
- C. Server:server03
- D. Machine:machine02
- E. Session:weblogic

ANSWER: A

Explanation:

`Domain:TestDomain` is printed because a successful WLST `connect` operation establishes an online connection and, unless a different navigation tree is explicitly selected, WLST enters the `serverConfig` tree. The current management object (`cmo`) at the root of that tree is the domain configuration MBean, whose type is `Domain` and whose name is the configured domain name, `TestDomain`. Therefore, concatenating `cmo.getType()`, a colon, and `cmo.getName()` produces `Domain:TestDomain`. The host and port in the connection URL identify the Administration Server endpoint used to establish the management connection; they do not automatically change `cmo` to a `Server`, `Machine`, or runtime MBean. WLST navigation commands such as `domainRuntime()`, `serverRuntime()`, or an explicit `cd()` would be required before `cmo` represented another MBean. Oracle documents that the configuration MBean hierarchy is accessed through the server configuration tree and that the `DomainMBean` is its root. See the [WLST command reference](#) and Oracle's [WLST navigation documentation](#).

QUESTION NO: 3

Review this scenario of an application running on a storage-disabled, managed Coherence server that uses a Coherence two-tier near cache:

Cache data is already written to the distributed cache.

There is a local near cache configured on the application side (storage-disabled node).

Application requests the data from Coherence twice.

Assuming this is the server's first access of this data from the cache, how many network round trips between Coherence cluster instances take place in total? (Choose the best answer.)

- A. 3
- B. 4
- C. 1
- D. 2

ANSWER: C

Explanation:

1 is correct. A Coherence near cache is a two-tier cache consisting of a local front cache and a distributed back cache. Because this is the storage-disabled application member's first request for the entry, its local front cache does not yet contain the value. Coherence therefore obtains the value from the distributed cache on the storage-enabled member that owns the relevant partition. That remote retrieval requires one request to the owning cluster member and one response back to the requesting member, which together constitute one network round trip. Once the value is returned, the near cache stores it in the local front cache, subject to its configured invalidation and eviction policies. The application's second request can then be served directly from that local front cache, without contacting another Coherence cluster member and without another network round trip. This behavior is the purpose of a near cache: it retains frequently accessed data locally while maintaining consistency with the distributed cache through its configured invalidation strategy. Oracle documents that NearCache uses a local cache in front of a distributed cache and automatically populates the front cache when values are requested. See the [Oracle Coherence Developing Applications guide](#) and the [Oracle Coherence application-development documentation](#).

QUESTION NO: 4

You need to ensure that when a failed admin server instance is restarted on another machine with a different IP address, the managed server is able to establish connection with the newly started instance.

Which technique will ensure that the connection can be made? (Choose the best answer.)

- A. Configure a DNS name for AdminServer which maps to multiple IP addresses and specify this DNS name during managed server startup.
- B. Specify multiple admin server IP addresses during managed server startup.
- C. Use WLST to update the AdminServer IP address for a running managed server.
- D. Update the host file for a running managed server instance.

ANSWER: A

Explanation:

Configure a DNS name for AdminServer which maps to multiple IP addresses and specify this DNS name during managed server startup. A Managed Server identifies its Administration Server through the administration-server URL supplied at startup, commonly by using the `-Dweblogic.management.server` system property. Using a stable logical DNS host name in that URL decouples the Managed Server configuration from the physical IP address of the machine currently hosting the Administration Server. When the Administration Server is restarted on a different host, DNS can resolve the same logical name to the applicable address, allowing the Managed Server to reconnect to the Administration Server without changing its startup configuration. This approach is especially appropriate for an Administration Server recovery or relocation design because the stable name remains the connection contract while the underlying network endpoint can change. The

DNS entry and any applicable DNS caching settings should be designed so that the Managed Server can obtain the current address after the Administration Server moves. Oracle documents the Administration Server URL as a Managed Server startup requirement and describes the relationship between Managed Servers and the Administration Server in its server startup and cluster administration documentation. See [Starting and Stopping Oracle WebLogic Server](#) and [Administering Clusters for Oracle WebLogic Server](#).

QUESTION NO: 5

One server in your WebLogic domain is experiencing a problem. You want to use the WebLogic Diagnostic Framework (WLDF) to diagnose the problem. The default built-in WLDF module is currently configured on the server, but it does not capture the metrics you need.

You have the following requirements:

You do not want to delete the existing module.

You want to minimize metric data output while troubleshooting the issue.

You want to focus your analysis on only the server that is experiencing the problem.

What WLDF configuration would achieve this result? (Choose the best answer.)

- A.** Create a new module and add metrics to it, target the new module to the server with the problem, and then activate the new module alongside the existing module.
- B.** Deactivate the existing module, create a new module and add metrics to it, set an instance parameter to the server with the problem, and then activate the new module.
- C.** Deactivate the existing module, create a new module and add metrics to it, target the new module to the cluster with the problem, and then activate the new module.
- D.** Deactivate the existing module, create a new module and add metrics to it, target the new module to the server with the problem, and then activate the new module.

ANSWER: A

Explanation:

Create a new module and add metrics to it, target the new module to the server with the problem, and then activate the new module alongside the existing module. WLDF supports multiple diagnostic modules in a domain, including custom system modules that can operate at the same time as the built-in diagnostic module. This permits the existing module to remain intact while a dedicated troubleshooting configuration is introduced.

A custom module can contain only the harvested metric definitions needed for the immediate investigation. Limiting the selected metrics, rather than broadly collecting additional data across the domain, reduces diagnostic output and makes the resulting information easier to analyze. Targeting that module directly to the affected server scopes its instrumentation and collection to that specific server, avoiding unnecessary collection from healthy managed servers.

After the issue is resolved, the custom module can be deactivated or removed without changing the baseline configuration supplied by the existing WLDF module. Oracle documents that diagnostic system modules are targetable resources and that the Harvester component collects configured runtime MBean attributes at defined sample intervals. See the [Oracle WLDF configuration documentation](#) and the [Oracle Harvester documentation](#).

QUESTION NO: 6

You are using the WebLogic Diagnostic Framework (WLDF) to monitor your domain for certain scenarios based on runtime data of your running servers. Different groups are responsible for different aspects of the application.

As a result, your configuration must fulfill the following requirements:

- Notifications must be sent only to the correct group that is responsible for the respective part of the application.

- You must ensure that you can enable and disable monitoring for each level whenever it is needed.

How would you configure WLDF watches and notifications to achieve this result? (Choose the best answer.)

- A. Configure a single module with a watch for each area, and use a notification filter to send an email to the correct group for each area.
- B. Configure a module for each area with a watch that triggers a notification email for the group responsible for that area.
- C. Configure a single module with a watch for each area, and each watch triggers a notification email for the group responsible for that area.
- D. Configure a module for each area with a watch that uses a notification filter to send an email to the correct group for each area.

ANSWER: B

Explanation:

Configure a module for each area with a watch that triggers a notification email for the group responsible for that area. A WLDF diagnostic system module is a self-contained configuration unit that can contain watches and notification definitions. Creating one module per monitored application area establishes a clear operational boundary: each module contains the watch rules applicable to that area and an email notification configured with the recipient group responsible for it. Consequently, a watch firing in one area invokes only that area's notification, rather than distributing an event to unrelated teams.

This design also provides the requested independent control over monitoring. WLDF diagnostic system modules can be enabled or disabled independently, allowing administrators to turn monitoring on or off for one area without altering diagnostic behavior for the other areas. The module remains the appropriate level of isolation when the requirement concerns both the watches and their associated notification configuration, rather than merely suppressing a single event condition. Oracle documents diagnostic modules as the containers for WLDF monitoring configuration and documents watches as rules that trigger one or more configured notifications. See Oracle's [Configuring Diagnostic System Modules](#) and [Configuring Watches and Notifications](#).

QUESTION NO: 7

You are in the process of validating the architecture of a domain that will serve applications to automate various business activities for a large retailer. The applications rely heavily on JMS and you need to pay special attention to JMS storage. You argue in favor of JDBC-based storage rather than file storage.

Your management has asked you to come up with a few key points about why this persistence option would be a better option in case of failure.

Which two justifications would support your choice? (Choose two.)

- A. JDBC-based stores can be accessed by any server and usually do not reside on the machine that hosts the server.
- B. It is easier to recover data from a database rather than from a file.
- C. JDBC-based stores are the default choice when creating a new persistent store so it must be an Oracle-recommended choice.
- D. JDBC-based stores offer better throughput than file stores.

ANSWER: A B

Explanation:

JDBC-based stores can be accessed by any server and usually do not reside on the machine that hosts the server is a key resilience benefit. A JDBC store persists JMS messages and other persistent-store records in a database, which is normally a separately managed, network-accessible service rather than local server disk. This allows a migrated WebLogic service to reconnect to the same persistent data after a host or server failure, provided that the database remains available. It therefore

supports architectures that need service migration and recovery without requiring the message data to remain tied to the failed machine.

It is easier to recover data from a database rather than from a file is also a valid operational justification. Database platforms provide established backup, restore, replication, monitoring, transaction recovery, and administrator tooling. These capabilities can make recovering and validating persistent JMS data more manageable following storage or host failures. Oracle describes JDBC stores as using a JDBC-accessible database for persistent-store data and identifies their value where database availability and administrative capabilities are required. The database and its recovery procedures must, of course, be designed appropriately for the application's recovery objectives.

See Oracle's [Persistent Stores overview](#) and [JMS advanced configuration documentation](#).

QUESTION NO: 8

You have an Oracle Fusion Middleware domain that uses an Oracle database to store application data and metadata repositories. As part of your disaster recovery configuration, you have set up a passive secondary site. You need to ensure that the data tier components of the primary and secondary sites are in sync so that you can perform a switchover to secondary in case of a disaster scenario. You check with your technical team regarding how to accomplish this requirement.

Which Oracle product is the technical team most likely to recommend for achieving data tier synchronization across the two sites? (Choose the best answer.)

- A. Oracle Service Bus
- B. Oracle ExaLogic
- C. Oracle Tuxedo
- D. Oracle Data Guard

ANSWER: D

Explanation:

Oracle Data Guard is the appropriate product because it provides disaster-recovery protection and synchronization for Oracle Database instances across a primary site and a standby site. It maintains one or more standby databases by transmitting and applying redo data generated by the primary database. As a result, the standby database can remain synchronized with the application data and the Fusion Middleware metadata repositories held in the primary database.

For an Oracle Fusion Middleware disaster-recovery topology with a passive secondary site, Data Guard is used to protect the database tier and supports role transitions such as switchover and failover. A switchover is a planned role reversal in which the primary and standby databases exchange roles without data loss, enabling the secondary site to take over in a controlled disaster-recovery or maintenance scenario. Oracle's Fusion Middleware disaster-recovery guidance identifies Oracle Data Guard as the database replication mechanism for keeping the primary and standby database sites synchronized. Oracle Database documentation further describes Data Guard's redo transport, redo apply, and role-transition capabilities.

See the [Oracle Fusion Middleware Disaster Recovery Guide](#) and the [Oracle Data Guard Concepts and Administration guide](#).

QUESTION NO: 9

You are responsible for creating scripts to use Node Manager with a WebLogic Server domain.

Which four statements correctly describe the WLST command: nmEnroll? (Choose four.)

- A. WLST must NOT be connected to a Node Manager when invoking nmEnroll.
- B. The nmEnroll command enables the nmConnect command to be used in a production environment.
- C. The nmEnroll command downloads the nm_password.properties and SerializedSystemIni.dat files.
- D. WLST must be connected to a Node Manager when invoking nmEnroll.

E. WLST must be connected to an Administration Server when invoking nmEnroll.

F. The nmEnroll command is run once per WebLogic domain per machine unless that domain shares the root directory of the Administration Server.

G. WLST must be offline when invoking nmEnroll.

ANSWER: A B C E F

Explanation:

`nmEnroll` is an online WLST command that enrolls the local machine's Node Manager configuration with a WebLogic domain. WLST must be connected to the domain's Administration Server so it can obtain the domain-specific Node Manager credentials and security artifacts; it must not be connected to Node Manager when the command is invoked. As part of enrollment, the command downloads the `nm_password.properties` and `SerializedSystemIni.dat` files needed for secure Node Manager use on that machine. This enrollment supplies the configuration that permits `nmConnect` to be used in a production environment, where authenticated Node Manager access is required.

Enrollment is normally performed once for each WebLogic domain on each machine that will run Node Manager. The documented exception applies when the domain shares the Administration Server's root directory, because the required domain configuration is already shared and a separate per-machine enrollment is not necessary. Therefore, the supplied choices contain five technically correct statements rather than the four stated in the question. Oracle documents the command behavior, connection requirements, downloaded files, and per-machine enrollment guidance in the [WLST Command Reference](#) and the [Node Manager documentation](#).

QUESTION NO: 10

Which statement is correct regarding JMS Template configuration? (Choose the best answer.)

A. Message Priority defined for the Template configuration overrides that specified by the producer.

B. Message Expiration Policy cannot be set on the template and needs to be configured for actual destinations that reference the template.

C. JMS Template configuration must include the quota associated with the template.

D. JMS Template configuration must include the message sort order associated with the template.

ANSWER: A

Explanation:

Message Priority defined for the Template configuration overrides that specified by the producer. A JMS template supplies reusable destination settings for queues and topics that reference it, including delivery-related defaults and overrides. The template's Priority Override setting controls whether WebLogic Server uses a configured destination priority instead of the priority submitted by a JMS producer. When a priority override value is configured, messages delivered to destinations using that template are assigned the configured priority, providing consistent priority handling regardless of individual producer settings. This is useful where administrators must enforce a destination-level service policy, such as ensuring that traffic sent to a particular queue is processed at a prescribed JMS priority. The override is inherited by destinations that use the template unless the relevant setting is explicitly configured at the destination level. Oracle documents this behavior through the JMS template Priority Override attribute, which specifies the priority assigned to messages independently of the producer-specified priority. See the [Oracle JMS Template MBean reference](#) and the [WebLogic JMS system resource configuration guide](#).

QUESTION NO: 11

Your developers have decided to use a deployment plan with their applications that will be deployed by WebLogic administrators into one or more WebLogic Server production environments.

Which two primary goals can be accomplished by leveraging a deployment plan? (Choose two.)

- A. Expose the external resource requirements of the application, such as JNDI names of datasources.
- B. Expose additional configurable properties, such as tuning parameters.
- C. Expose the targeted servers for deployment.
- D. Expose the specific Java EE modules within the application, such as EJBs.
- E. Expose access points into the application, such as a WSDL for web services.

ANSWER: A B

Explanation:

Deployment plans separate environment-specific deployment configuration from an application's deployable archive. This lets developers package the application while allowing WebLogic administrators to supply values appropriate to each target environment without editing the archive itself. A principal use is exposing external resource requirements as deployment-plan variables, including resource references and JNDI names. For example, an application can retain its resource-reference configuration while an administrator maps it to the datasource JNDI name available in a particular test, staging, or production domain.

Deployment plans also expose other configurable descriptors as variables, including operational and tuning-related settings selected during the export process. Administrators can then adjust those values for a production deployment and preserve them in a plan that can be reused or maintained independently from the application binary. This supports portability across environments and keeps configuration changes under deployment administration control. Oracle describes deployment plans as XML documents that define deployment configuration and allow descriptor properties to be overridden externally; its deployment configuration export process specifically supports exposing resource requirements and additional configurable properties as variables. See [Oracle WebLogic Server Deployment Plans documentation](#) and [Oracle documentation on exporting deployment configuration](#).

QUESTION NO: 12

You are creating a Domain Template to simplify the process of deploying a domain across multiple hosts.

Which five types of information can be included in a Domain Template? (Choose five.)

- A. domain configuration, those included in the config.xml file
- B. Java Enterprise Edition (Java EE) applications and shared libraries
- C. Windows and UNIX server start scripts
- D. Windows Start menu entries
- E. custom folders and files
- F. LDAP data
- G. JTA log
- H. e-mail notification settings

ANSWER: A B C D E

Explanation:

A WebLogic Domain Template can package the domain definition and reusable deployment content needed to create consistently configured domains on other hosts. The domain configuration, those included in the config.xml file, is included automatically by Domain Template Builder; this encompasses the domain's configured resources and settings represented by its central configuration. Java Enterprise Edition (Java EE) applications and shared libraries can be selected for inclusion so that deployments travel with the template and can be recreated by the Configuration Wizard.

Windows and UNIX server start scripts are eligible template content. Domain Template Builder processes standard scripts and replaces host-specific hard-coded domain environment values with replacement variables, allowing the generated domain to use its local environment. Windows Start menu entries may also be defined as template content for applicable Windows installations. Finally, custom folders and files can be added explicitly, enabling administrators to package supporting artifacts, such as organization-specific scripts, configuration files, or other resources required by the domain.

These capabilities make templates portable while retaining both the WebLogic configuration and selected operational or application assets. Oracle documents the template-content selection workflow, automatic inclusion of configuration files, application/library handling, script processing, Start menu entries, and custom file inclusion in the [WebLogic Server Administration Guide](#) and the [Domain Template Reference](#).

QUESTION NO: 13

You are using the WebLogic Diagnostic Framework (WLDF) to monitor your domain for certain scenarios based on runtime data of your running servers. Different groups are responsible for different aspects of the application.

As a result, your configuration must fulfill the following requirements:

Notifications must be sent only to the correct group that is responsible for the respective part of the application.

You must ensure that you can enable and disable monitoring for each level whenever it is needed.

How would you configure WLDF watches and notifications to achieve this result? (Choose the best answer.)

- A.** Configure a single module with a watch for each area, and use a notification filter to send an email to the correct group for each area.
- B.** Configure a module for each area with a watch that triggers a notification email for the group responsible for that area.
- C.** Configure a single module with a watch for each area, and each watch triggers a notification email for the group responsible for that area.
- D.** Configure a module for each area with a watch that uses a notification filter to send an email to the correct group for each area.

ANSWER: B

Explanation:

Configure a module for each area with a watch that triggers a notification email for the group responsible for that area. This design provides an independent WLDF configuration boundary for every monitored application area. Each module contains the watch rules and notification configuration relevant to one responsible group, so a watch firing in that module invokes the email notification configured specifically for that group. The resulting separation is operationally useful: monitoring for one area can be enabled, disabled, deployed, or otherwise managed without changing the watches and notifications for unrelated areas. It also keeps ownership clear, because the diagnostics artifacts for an area are grouped together rather than shared in one broad module.

WLDF watches evaluate runtime data against configured rules and can trigger configured notification mechanisms, including email. A WLDF system module is a deployable, manageable diagnostics resource; creating one per independently managed monitoring area satisfies the requirement for granular lifecycle control while maintaining direct routing of alerts to the appropriate responsible team. Oracle's WLDF documentation describes defining watches and associating notifications within a diagnostic module, while the administration documentation describes WLDF modules as domain configuration resources. See [Configuring Watches and Notifications](#) and [Configuring Diagnostic System Modules](#).

QUESTION NO: 14

You have defined Uniform Distributed Destination (UDD) topics targeted to multiple JMS servers.

Which three statements are correct about UDDs? (Choose three.)

- A.** UDDs guarantee the order of messages.

- B. UDDs increase the overall availability of JMS.
- C. Weight-based load balancing policy distributes messages evenly to servers with preassigned weights.
- D. JMS destinations do not need to run on each server in a cluster.
- E. You no longer need to create or designate destination members.

ANSWER: B D E

Explanation:

UDDs increase the overall availability of JMS because a uniform distributed topic is represented by member destinations on the targeted JMS servers. This distribution avoids depending on a single physical destination member and lets WebLogic route production and consumption through the distributed destination as members are available. UDDs also simplify deployment topology: JMS destinations do not need to run on each server in a cluster, because the UDD can be targeted to the relevant JMS servers or a cluster rather than requiring every cluster server to host a JMS destination.

You no longer need to create or designate destination members because WebLogic automatically creates and manages the uniform member destinations for the JMS servers included in the UDD target. When the target is a cluster, WebLogic creates a corresponding member on each eligible JMS server and adjusts membership as JMS servers are added or removed. This automatic membership management is a defining administrative advantage of UDDs over manually configured distributed destinations. Oracle documents UDD behavior and automatic member creation in its [Configuring JMS System Resources](#) guide; the broader distributed-destination architecture is covered in the [Advanced JMS System Resource Configuration](#) guide.

QUESTION NO: 15

You are tasked with doing a rolling upgrade of a production domain that will minimize user interruptions.

What are two characteristics of a rolling upgrade? (Choose two.)

- A. Configuration changes can be made during the Rolling Update process.
- B. A major upgrade, for example 11g to 12c, can be accomplished with a Rolling Update.
- C. The machines can be upgraded in any order.
- D. The machine hosting the Administration Server must be upgraded first.
- E. Avoid configuration changes until all the servers in the domain have been upgraded.

ANSWER: D E

Explanation:

In a WebLogic Server rolling upgrade, the Administration Server is upgraded first. The Administration Server provides the domain's central management and configuration services, so it must run the updated software before the Managed Servers are upgraded and restarted one at a time. This sequencing lets the remaining Managed Servers continue servicing applications while an individual server is taken out of service, updated, and returned to the cluster, thereby reducing interruption for users.

It is also important to avoid configuration changes until every server in the domain has been upgraded. During the transition, servers can be running different software patch or release levels. Changing the domain configuration—particularly by using newly introduced configuration attributes—can cause incompatibilities for servers that have not yet been upgraded. Deferring changes maintains a stable, consistent configuration throughout the rolling procedure. After all servers run the upgraded version, administrators can safely make and activate configuration changes appropriate to that version. Oracle's upgrade guidance describes the rolling process and its required sequencing in the [WebLogic Server Rolling Upgrade documentation](#); related upgrade planning guidance is available in the [WebLogic Server Upgrade Guide](#).

QUESTION NO: 16

You are configuring Java Node Manager for Managed Servers in a production domain.

Which two statements are true about the following Node Manager properties? (Choose two.)

- A. StartScriptEnabled determines whether Node Manager uses a server start script to start servers.
- B. CrashRecoveryEnabled enables Node Manager recovery of servers that were running before Node Manager stopped.
- C. StartScriptEnabled configures the maximum number of Managed Server restart attempts.
- D. CrashRecoveryEnabled enables automatic migration of all migratable targets.
- E. StartScriptEnabled applies only to script-based Node Manager.

ANSWER: A B

Explanation:

`StartScriptEnabled` controls whether Java Node Manager invokes a server start script when it starts a server. When this property is enabled, Node Manager uses the configured start script rather than directly constructing the server-start command. This is useful when the startup environment requires custom operating-system commands or environment preparation.

`CrashRecoveryEnabled` controls Node Manager crash recovery. When it is enabled, Node Manager can attempt to restart servers that were running when Node Manager stopped unexpectedly, such as after a host or Node Manager outage. The applicable servers must be configured for Node Manager control.

These properties do not replace the server-specific restart controls, such as `RestartMax` and `RestartInterval`, which govern the number and timing of restart attempts. Oracle documents Java Node Manager properties in [Advanced Node Manager Configuration](#).

QUESTION NO: 17

You are configuring a WebLogic JMS application to preserve message ordering for messages that share the same business-order identifier when they are sent to a distributed destination.

Which message property is used by WebLogic JMS for a Unit-of-Order value? (Choose the best answer.)

- A. JMS_BEA_UnitOfOrder
- B. JMSXGroupID
- C. JMSMessageSequence
- D. JMS_BEA_DeliveryOrder

ANSWER: A

Explanation:

The correct property is `JMS_BEA_UnitOfOrder`. WebLogic JMS uses this property to associate messages with a named unit of order. Messages carrying the same unit-of-order value are routed and processed in a manner that preserves ordering for that unit, including in supported distributed-destination configurations.

The unit-of-order value is a string identifier. It can be configured through a connection factory or set by the sending application, depending on the desired design. Different unit-of-order values allow independent streams of work to be processed concurrently while preserving order within each individual stream.

Oracle documents units of order and the `JMS_BEA_UnitOfOrder` property in [Using Message Unit-of-Order](#).