

Oracle WebLogic Server 12c: Administration I

Oracle 1z0-133

Version Demo

Total Demo Questions: 15

Total Premium Questions: 154

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, WebLogic Server Architecture	3
Topic 2, Installing WebLogic Server	5
Topic 3, Configuring WebLogic Server Domains	25
Topic 4, Configuring WebLogic Server	28
Topic 5, Deploying Applications	11
Topic 6, Configuring JDBC Resources	20
Topic 7, Configuring Java Messaging Service (JMS)	3
Topic 8, Configuring Security	10
Topic 9, Monitoring WebLogic Server	20
Topic 10, Configuring Clusters	29
Total	154

QUESTION NO: 1

Your developers created Plain Old Java Objects (POJO) that provide useful functionality. The code is packaged in a JAR file. You want to make this code accessible from all the applications deployed to the managed servers in your domain. What are the four ways to accomplish this? (Choose four.)

- A. Include the JAR file as an optional package
- B. Deploy the JAR file as an EJB
- C. Deploy the JAR file as a shared library
- D. Place the JAR file in the lib folder of the domain
- E. Deploy the JAR file to the administration server. All managed servers have access to their admin server
- F. Add the JAR file to the CLASSPATH of the domain servers by editing the domain's `setDomainEnv.sh` script
- G. Deploy the JAR file as another web application

ANSWER: A C D F

Explanation:

Including the JAR file as an optional package is a supported class-loading approach for making reusable classes available to applications that declare a dependency on that package. Optional packages use JAR manifest extension metadata, allowing the dependent application to locate and load the shared code without bundling duplicate copies into every application archive.

Deploying the JAR file as a shared library is the standard WebLogic Server deployment mechanism for reusable application code. Applications reference the deployed library through their deployment descriptors, and WebLogic makes the specified library available to each referencing application on the managed servers to which it is targeted.

Placing the JAR file in the domain `lib` directory makes it available on the class path for servers in that domain. This is appropriate for classes intended to be common infrastructure used broadly by deployed applications.

Adding the JAR file to the domain servers' class path through `setDomainEnv.sh` is also a valid domain-level technique. The domain environment script establishes server startup environment settings, so a JAR added to the relevant class path is visible to applications running on those managed servers. Oracle documents both the domain library directory and class-path configuration as class-loading mechanisms; shared-library deployment is documented separately. See [Oracle WebLogic Server Classloading](#) and [Oracle WebLogic Server Shared Libraries](#).

QUESTION NO: 2

An administrator in your organization says you do not need to back up the domain files on a machine which only managed servers run.

Select the best explanation for this statement.

- A. This is true because the domain files can be re-created with the pack and unpack utilities.
- B. This is true (with an exception) because the domain files can be re-created with the pack and unpack utilities. The exception is if your security realm uses the embedded LDAP system. Which is not included in the JAR file created by the pack utility.
- C. This is true (with an exception) because the domain files can be re-created with the pack and unpack utilities. The exception is for applications that are deployed to the managed server. Those files are not included in the JAR file created by the pack utility. If you back those up manually, then pack and unpack can be used.
- D. This is not true. You need to back up the domain files on a machine on which only managed servers run by using operating system commands (such as tar or zip).

ANSWER: C

Explanation:

The statement is true with the stated application-file exception because a Managed Server host's domain directory can be reconstructed rather than restored from a conventional filesystem backup. The WebLogic Server `pack` utility can create a Managed Server template JAR from the domain, and `unpack` can use that template to re-create the domain files required on another machine that runs the Managed Server. This makes the Administration Server's domain configuration and the generated template the important recovery sources for a Managed Server-only host.

Application artifacts require separate protection when they are deployed to the Managed Server and are not included in the template JAR. Before relying on `pack` and `unpack` for recovery, those application files must therefore be backed up independently and restored to the appropriate location. With the application artifacts preserved, `unpack` can rebuild the Managed Server domain structure and the restored applications can again be available to the server. Oracle documents `pack` and `unpack` as the mechanism for creating and extending domains, including Managed Server templates, in its [Creating Templates and Domains documentation](#). See also Oracle's [domain configuration overview](#) for the domain components that must be managed during recovery.

QUESTION NO: 3

Your Java-based Node Manager is unable to start a managed server named `myserver`. You investigate and realize that the `boot.properties` file Node Manager uses to store the credentials of the managed server has been corrupted in the file system. From past experience you know that if you delete the file, Node Manager will re-create it.

- A. under the domain directory in `servers/myserver/nodemanager`
- B. under the Node Manager home directory in `myserver`
- C. under the domain directory in `servers/myserver/security`
- D. under the domain directory in `servers/myserver/data/nodemanager`
- E. under the Node Manager home directory

ANSWER: C

Explanation:

The correct location is *under the domain directory in `servers/myserver/security`*. A server's boot identity file is named `boot.properties` and, for a Managed Server named `myserver`, its standard domain-relative location is `DOMAIN_HOME/servers/myserver/security/boot.properties`. This file supplies the server startup identity, including the username and password that allow the server to connect to its Administration Server. When Node Manager starts the Managed Server, it starts that server using the server-specific configuration and boot identity associated with the domain.

Initially, a boot identity file can contain clear-text credentials. During startup, WebLogic Server encrypts those values and rewrites the file, provided the server can successfully read it and establish its domain security context. Therefore, when this server-specific file is corrupted, removing the corrupted `boot.properties` from the Managed Server's `security` directory allows the expected boot-identity-file handling to occur again when valid startup credentials are supplied. Oracle documents the default boot identity file location as the server's `security` directory and describes Node Manager's role in starting Managed Servers. See [Oracle WebLogic Server server startup documentation](#) and [Oracle Node Manager documentation](#).

QUESTION NO: 4

Over time your server log files have increased in size to over 500 MB. Identify three techniques to help decrease the size of these log files. (Choose three.)

- A. Increase the severity level
- B. Redirect standard output to the log
- C. Increase the rotation time

D. Create and use a custom log filter

E. Decrease the rotation file size

ANSWER: A D E

Explanation:

Increasing the severity level reduces the volume of messages written to the server log because WebLogic records only messages at the configured threshold and above. For example, using a more restrictive threshold than an informational level prevents lower-severity operational messages from being persisted, which directly slows log growth.

Creating and using a custom log filter also reduces logged output by selecting the message records that should be written based on defined criteria. This permits targeted suppression of unnecessary messages while retaining the categories that are operationally important.

Decreasing the rotation file size limits the maximum size of each individual log file before WebLogic rotates it and begins a new file. This is particularly appropriate when the concern is that a single server log has become too large to open, transfer, search, or retain conveniently. Log rotation settings control the size and lifecycle of individual log files, while message severity and filters reduce the volume generated in the first place. Oracle documents server-log severity thresholds, log filters, and rotation policies in its WebLogic Server logging configuration documentation. See [Oracle WebLogic Server logging configuration help](#) and [Oracle WebLogic Server log configuration documentation](#).

QUESTION NO: 5

View the following file excerpt, which configures HTTP session replication:

```
< session-descriptor >
```

```
< persistent-store-type > replicated
```

From which file is this excerpt taken?

A. Weblogic –application.xml

B. Weblogic.xml

C. Session.xml

D. Config.xml

E. Web.xml

ANSWER: B

Explanation:

`weblogic.xml` is correct because it is the WebLogic Server deployment descriptor for a Web application and contains the `session-descriptor` configuration element. Within that element, `persistent-store-type` controls how HTTP session state is persisted or replicated by WebLogic Server. Setting it to `replicated` instructs the server to replicate eligible in-memory HTTP session data among members of a cluster, supporting session failover when a server instance becomes unavailable.

This is WebLogic-specific behavior, so it is configured in the WebLogic-specific Web application descriptor rather than in a standard Java EE deployment descriptor. The descriptor is packaged in the Web application's `WEB-INF` directory as `weblogic.xml`. Oracle documents `session-descriptor` as an element of the `weblogic-web-app` descriptor root and documents `persistent-store-type` as the setting that selects the session persistence mechanism, including replicated storage. See Oracle's [weblogic.xml Deployment Descriptor Reference](#) and its [HTTP Sessions in WebLogic Server](#) documentation.

QUESTION NO: 6

Which session persistence type is invalid?

- A. `async_replicated`
- B. `async_jdbc`
- C. `replicated`
- D. `async_memory`
- E. `cookie`
- F. `file`
- G. `replicated_if_clustered`

ANSWER: D

Explanation:

`async_memory` is invalid because WebLogic Server does not define an asynchronous in-memory HTTP-session persistence mechanism. WebLogic supports memory-based session storage, but that mode is simply configured as `memory`; it keeps session state locally and does not provide an asynchronous persistence or replication variant under the name `async_memory`.

For session failover and durability requirements, WebLogic provides explicitly named persistence mechanisms such as asynchronous replication and asynchronous JDBC persistence. These modes use background or asynchronous handling to replicate session state to another managed server or save it to a database, respectively. The valid persistence-type vocabulary is defined by WebLogic's HTTP session-management configuration and does not include `async_memory`. Administrators should select the supported persistence type that corresponds to the desired availability, clustering, and storage behavior, then configure the related replication or JDBC resources as necessary.

Oracle documents the available HTTP session persistence approaches and their configuration in the [WebLogic Server HTTP Sessions documentation](#). The relevant descriptor attribute is also described in the [weblogic.xml deployment descriptor reference](#).

QUESTION NO: 7

All administrators that deal with the production domain in your organization are extremely experienced. As the lead administrator, you have been asked to change the administration console preferences for all administrators so that inline help no longer displays, since no one needs it.

Select the true statement about this scenario.

- A. There is nothing you can do, since inline help cannot be disabled.
- B. There is nothing you can do for all users, since disabling inline help must be done per user.
- C. Log into the administration console and disable inline help under "help."
- D. Access WLST, connect to the domain's admin server, and use the command `EnableInlineHelp` (false).

ANSWER: B

Explanation:

There is nothing you can do for all users, since disabling inline help must be done per user. Inline help is controlled through an individual administrator's Administration Console preferences rather than through a domain-wide configuration setting. An administrator can open the Console's preferences and clear the setting that displays inline help; however, that selection

applies to that administrator's own Console experience. It does not propagate to other administrators, even when they connect to the same Administration Server and manage the same production domain.

This behavior is intentional because Console preferences govern user-interface presentation choices, such as help display behavior, instead of operational domain configuration. Domain configuration changes are centrally managed and activated through the change-management process, whereas user preferences let each authenticated administrator tailor the Console interface to their working needs. Therefore, the lead administrator cannot centrally suppress inline help for every administrator using a single Console or WLST domain-level action. Each administrator who wants inline help hidden must change the relevant preference in their own Console session. Oracle's Administration Console documentation is available in the [Administration Console Online Help](#); see also the [Administration Console overview](#).

QUESTION NO: 8

In the certification matrix for your version of WebLogic Server , you find the row for your database and its version. In that row, in the "type of database use" column, it says "application data access only".

Your database cannot be used _____.

- A. By the administration server
- B. For internal data storage by WebLogic Server
- C. By a multidata source
- D. In a production environment
- E. By an Enterprise JavaBean

ANSWER: B

Explanation:

"For internal data storage by WebLogic Server" is correct. The certification matrix classification "application data access only" means that the database and version are supported for application-level JDBC access: deployed applications may connect to it through WebLogic Server data sources and use it to read and write their own business data. It does not certify that database combination for WebLogic Server's own internal database-dependent services.

WebLogic Server internal storage includes server-managed persistent data, such as data held in JDBC persistent stores. These stores can support WebLogic messaging and transaction-related persistence requirements, depending on the configured services. Because this kind of usage is part of WebLogic Server infrastructure rather than ordinary application access, it requires a database/version combination whose certification entry explicitly permits internal data storage. Administrators should therefore consult the applicable WebLogic Server certification matrix before selecting a database for JDBC stores or other server-internal persistence. Oracle publishes supported configuration and certification information through its [Fusion Middleware certification resources](#), and its WebLogic documentation describes the role of JDBC data sources in providing database connectivity to applications at [Understanding JDBC Data Sources](#).

QUESTION NO: 9

Which three statements are true about the administration console? (Choose three.)

- A. The administration console can be disabled.
- B. The context path of the administration console may be changed.
- C. If a domain is in Development Mode, the administration console no longer requires an administration-level user when logging in.
- D. If a domain is in Production Mode, access to the administration console must be made through a secure port.
- E. If the domain's Administration Port is enabled, access to the administration console must be made only through that port.

ANSWER: A B E

Explanation:

The administration console can be disabled because it is an optional management application deployed on the Administration Server. WebLogic provides a domain-level Console configuration MBean with an enabled setting, allowing administrators to turn off console access when management is performed through other supported mechanisms, such as scripting or management APIs. Disabling it does not eliminate the Administration Server's underlying management capabilities.

The context path of the administration console may be changed through the Console configuration. This supports environments that require a nondefault management URL or need to align the console endpoint with organizational URL conventions. The console's context path is part of the Administration Server URL used by browser clients.

If the domain's Administration Port is enabled, access to the administration console must be made only through that port. The administration port is a domain-wide secure listener dedicated to administrative traffic. When enabled, administrative protocols, including Administration Console access, use that secure port rather than regular server listen ports. This separation helps protect management communications by requiring SSL-based access for administrative operations. Oracle documents both the console configuration and the administration-port behavior in its WebLogic Server administration and security guides: [Administration Console overview](#) and [Using the Administration Port](#).

QUESTION NO: 10

You require high availability for HTTP sessions in a web application. In addition, you have the following constraints:

- There is little free memory on the machines hosting the cluster.
- HA must be guaranteed even when multiple machines fail simultaneously.
- You are not allowed to update the production databases schema.

Which is the best session persistence option?

- A. Database persistence
- B. File persistence
- C. In-memory replication
- D. Whole-server migration

ANSWER: B

Explanation:

File persistence is the best fit because it writes HTTP session state to durable files rather than retaining replicated session copies in each managed server's heap. This minimizes the memory consumed on cluster machines, which is important when little free memory is available. For clustered high availability, the persistence directory should be located on shared, highly available storage that is accessible to every server capable of handling the application. Session data therefore remains available even if several hosting machines fail at the same time, provided that the shared storage itself is designed for the required availability level.

File persistence also avoids creating or modifying tables in the production database schema. The WebLogic Server session-persistence configuration identifies the file-based persistence mechanism and its storage location; the application's production database need not be used to hold HTTP session state. Operationally, the shared directory must provide reliable locking, adequate capacity, and resilience appropriate to the availability commitment. Oracle documents file-based persistence as a supported HTTP-session persistence mechanism and describes configuring session persistence at the Web application level. See the [Oracle WebLogic Server documentation on HTTP session state](#) and the [weblogic.xml session-descriptor configuration reference](#).

QUESTION NO: 11

You need to scale out an existing domain to include a new machine. You install WebLogic Server on the new machine. Which three steps are required to create the domain directory on the new machine? (Choose three.)

- A. Use the Configuration Wizard to clone the domain
- B. Run the pack tool on the administration server's machine
- C. Transfer the installation certificate to the administration server's machine
- D. Transfer the managed server template to the new machine
- E. Use the Administration Console to clone the administration server
- F. Run the unpack tool on the new machine

ANSWER: B D F

Explanation:

To create the domain directory for a new machine that will host Managed Servers in an existing WebLogic domain, the domain must be packaged on the Administration Server host, moved to the target host, and unpacked there. Running the `pack` tool on the administration server's machine with the managed-server template option creates an archive containing the domain configuration and files needed by a Managed Server installation. This archive is commonly called a managed server template.

That managed server template must then be transferred to the new machine, where WebLogic Server has already been installed. Running `unpack` on the new machine expands the template and creates the local domain directory structure. The resulting directory supplies the configuration required for the Managed Server to be started on that machine and subsequently communicate with the Administration Server. Oracle documents this pack-and-unpack procedure as the supported method for distributing an existing domain's configuration to remote Managed Server hosts. See the [Oracle pack and unpack command documentation](#) and the [Oracle guidance for configuring Managed Servers on remote machines](#).

QUESTION NO: 12

You are deploying a web application called `inventory.war`. It has been installed and has the State of "Prepared". It is targeted to the Managed Server named `managed1`. The web application's context root is `/inventory` and the starting page is `index.jsp`.

However, you want to test the application before allowing end users access to it. In the administration console, you select `inventory.war`, and then select Start and "Servicing only administration requests."

In your web browser, you need to enter a host and port followed by `/inventory/index.jsp`.

Which host and port would you use?

- A. `managed1` host and port
- B. Administration Server host and port
- C. `managed1` host and Administration port
- D. Administration Server host and Administration port

ANSWER: C

Explanation:

managed1 host and Administration port is correct because selecting "Servicing only administration requests" starts the application in WebLogic Server administration mode. In this mode, the deployment is available only to administrative clients, rather than to normal application users through the server's regular listen port. The application remains targeted to `managed1`, so the request must be directed to `managed1`'s network host name. To reach an administration-mode application, the request must use that Managed Server's Administration Port, which is reserved for administrative traffic when the domain has enabled administration ports.

Therefore, the appropriate test URL has the form `http(s)://managed1-host:administration-port/inventory/index.jsp`. This provides an administrator a controlled way to validate the deployed application before it is started to service all requests. Oracle documents that administration-mode deployments accept requests only through the Administration Port, and that the Administration Port is configured per server for administrative communication. See Oracle's [Deploying Applications to Oracle WebLogic Server](#) and [Configuring Server Startup and Shutdown](#) documentation.

QUESTION NO: 13

You are configuring a non-XA JDBC data source to participate with XA resources in global transactions. You decide to use the Logging Last Resource transaction option.

Which two statements are true? (Choose two.)

- A. An LLR data source can participate with XA resources in a global transaction.
- B. Only one LLR data source can participate in a global transaction.
- C. An LLR data source requires an XA JDBC driver.
- D. Any number of LLR data sources can participate in the same global transaction.
- E. LLR eliminates the need for transaction recovery records.

ANSWER: A B

Explanation:

Logging Last Resource (LLR) allows a non-XA JDBC data source to participate in a global transaction with XA resources. WebLogic Server records the transaction outcome in a table in the non-XA database, allowing the transaction manager to recover the transaction if a failure occurs.

Only one LLR data source can participate in a global transaction. This restriction exists because the LLR resource acts as the transaction's last resource and is coordinated differently from XA participants. The data source must use a non-XA JDBC driver and be configured with the Logging Last Resource transaction protocol. Oracle documents the restrictions and recovery behavior in [Using JDBC Data Sources](#).

QUESTION NO: 14

In memory replication in your cluster seems to be slow. You decide to create a replication channel for your cluster. Which two statements are true? (Choose two.)

- A. The channel is used for heartbeats so they do not interfere with replication communication
- B. Each server in the cluster must have this named replication channel configured on it
- C. If no replication channel is configured, then sessions are persisted to the file system or the database
- D. The channel can be used for stateful session bean replication
- E. The channel cannot use a secure protocol

ANSWER: B D

Explanation:

A replication channel is a dedicated WebLogic Server network channel that isolates in-memory replication traffic from normal application traffic. To use it at the cluster level, the same named replication channel must be configured on every server that participates in the cluster. This consistency lets each clustered server use the designated channel when it communicates replicated state with its peers, rather than relying on a channel that might not exist or have equivalent settings on another member.

The channel can be used for stateful session bean replication because WebLogic's in-memory replication mechanism replicates state for supported highly available services, including HTTP session state and stateful session EJBs. Directing this traffic through a dedicated channel can improve predictability and reduce contention when replication is slow due to competing network traffic. The channel is selected in the cluster's Replication settings, while the corresponding network channel is created and targeted on each cluster server. Oracle's administration guidance describes configuring a replication channel for a cluster and identifies the supported replicated state traffic: [Oracle WebLogic Server Administration Console Help: Replication Configuration](#). See also [Oracle documentation on configuring replication for a cluster](#).

QUESTION NO: 15

You are carrying out the RDBMS Security Store Configuration using the configuration Wizard Which three actions can you perform with the Configuration Wizard during the process? (Choose three.)

- A. Create tables in the data store.
- B. Execute SQL scripts in WebLogic server's server/lib directory.
- C. Populate the database tables with external LDAP users.
- D. Test the connection to a back-end database.
- E. Allocate tablespace for the data store.

ANSWER: A B D

Explanation:

The Configuration Wizard supports the essential database setup tasks required when configuring an RDBMS-based security store. It can validate the supplied database URL, credentials, and connectivity by testing the connection to the back-end database. This allows the administrator to confirm that the wizard can reach the target database before proceeding with repository configuration.

It can also execute the supplied SQL scripts used to establish the security-store repository. Those scripts create the database objects needed by the security services, including the required tables. Consequently, creating tables in the data store and executing SQL scripts in the WebLogic Server `server/lib` directory are supported setup actions in this configuration workflow. The database account used for the operation must have the appropriate database privileges to create and manage the required repository objects.

These tasks are part of establishing the RDBMS security store and its usable database connection; they enable Oracle Platform Security Services data to be stored in the configured database. Oracle documentation describes configuring a database security store and the associated database schema requirements in the [WebLogic Server Security Guide](#) and the domain configuration workflow in the [WebLogic Server Configuration Guide](#).