



Oracle

Exam 1z0-102

Oracle Weblogic Server 11g: System Administration I

Version: 9.0

[Total Questions: 111]

Question No : 1

Which two statements are true about java EE shared libraries?

- A. A shared library cannot be deployed to a cluster.
- B. An application that is targeted to server1 can use a shared library that is targeted to server 2.
- C. Multiple versions of the same shared library can be deployed and be active at the same time
- D. A shared library is referenced through an application's deployment descriptor.

Answer: C,D

Explanation: C: If more than one copy of myLibrary is registered, WebLogic Server selects the library with the highest specification version. If multiple copies of the library use the selected specification version, WebLogic Server selects the copy having the highest implementation version.

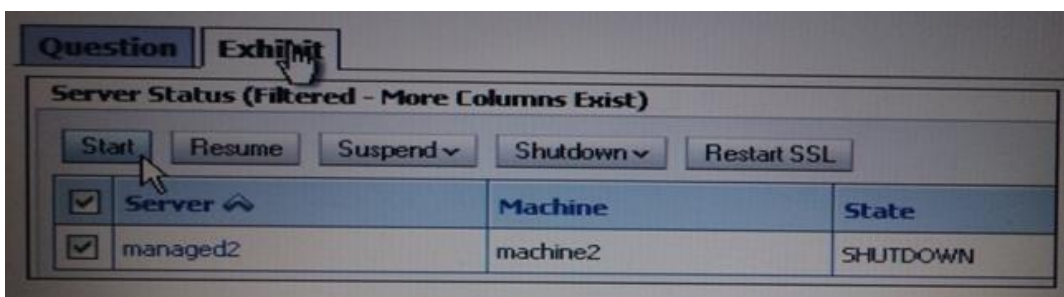
D: A Java EE application can reference a registered shared Java EE library using entries in the application's weblogic-application.xml deployment descriptor.

Reference: Referencing Shared Java EE Libraries in an Enterprise Application

http://docs.oracle.com/cd/E12840_01/wls/docs103/programming/libraries.html

Question No : 2

Refer to the Exhibit.



You use the Start button in the administration console to start managed2 successfully. Which two statements must be true?

- A. Managad2 is part of a cluster.
- B. The Administration Server is running.
- C. Managed2 has applications deployed to it.
- D. The Node Manager for machine2 is configured and running.
- E. The Node Manager for the Administration Server machine is configured and running.

Answer: B,D

Explanation: B: The administration server must be running.

Note: A WebLogic Server administration domain is a logically related group of WebLogic Server resources. Domains include a special WebLogic Server instance called the Administration Server, which is the central point from which you configure and manage all resources in the domain.

D:

Node Manager is a WebLogic Server utility that enables you to start, shut down, and restart Administration Server and Managed Server instances from a remote location. Although Node Manager is optional, it is recommended if your WebLogic Server environment hosts applications with high availability requirements.

A Node Manager process is not associated with a specific WebLogic domain but with a machine. You can use the same Node Manager process to control server instances in any WebLogic Server domain, as long as the server instances reside on the same machine as the Node Manager process. Node Manager must run on each computer that hosts WebLogic Server instances -- whether Administration Server or Managed Server -- that you want to control with Node Manager.

Question No : 3

Within your data center, the Administration and Managed Servers utilize a shared, central storage device, which server's access via NFS.

In this scenario, your application files no longer need to be copied to servers across the network. Which server attribute allows you to override this behavior?

- A. Startup Mode
- B. Staging Mode
- C. Root Directory
- D. Listen Address

Answer: B

Explanation: The server staging mode specifies the default deployment mode for a server if none is specified at deployment time. For example, the server staging mode is used if

you deploy an application or module using weblogic.Deployer and you do not specify a staging mode.

Reference:

http://docs.oracle.com/cd/E12840_01/wls/docs103/ConsoleHelp/taskhelp/deployment/SetAServerStagingMode.html

Question No : 4

An EJB application is targeted to a cluster. Remote EJB clients can therefore take advantage of WebLogic Server's load balancing and failover capabilities.

However, a proxy server exists between the clients and the cluster, which performs IP address translation. Which cluster attribute should you modify to ensure that load balancing and failover work correctly?








- A. Multicast Address
- B. Persistent Store
- C. Cluster Address
- D. Migration Basis
- E. Replication Channel

Answer: C

Explanation: Note:

Updating Proxy Service Configurations for an Expanded Cluster

If your AquaLogic Service Bus configuration includes one or more proxy services that use JMS endpoints with cluster addresses, then you must also perform the following procedure using the AquaLogic Service Bus Console after adding the new managed server to the cluster:

-  In the Change Center, click Create to create a session.
-  Using the Project Explorer, locate and select a proxy service that uses JMS endpoints with cluster addresses.
-  At the bottom of the View Details page, click Edit.
-  If there is a cluster address in the endpoint URI, add the new server to the cluster address.
-  On the Edit a Proxy Service - Summary page, click Save.
-  Repeat step 2. through step 5. for each remaining proxy service that uses JMS endpoints with cluster addresses.
-  In the Change Center, click Activate.

The proxy services are now configured for operation in the extended domain.

Reference: eDocs Home > BEA AquaLogic Service Bus 2.0 Documentation > Deployment Guide > Configuring a Clustered Deployment

Question No : 5

You run the default startmanageWeblogic (.cmd in Windows) script as shown here:
startmanageWeblogic.sh server1 http://192.168.1.102:8001.

What does this do?

- A. It starts the administration server named server1, which is running at 192.168.1.102.8001.
- B. It starts the managed server named server1, which is running at 192.168.1.102.8001.
- C. It starts the managed server named server1 whose Node Manager is running at 192.168.1.102.8001.
- D. It starts the Managed Server named server1 whose Administration Server is running at 192.168.1.102:8001.

Answer: D

Explanation: A Managed Server is a WebLogic Server instance that runs deployed applications. It refers to the Administration Server for all of its configuration and deployment information. Usually, you use Managed Servers to run applications in a production environment.

See step 4 and 5 below in particular:

To use the WebLogic Server scripts to start Managed Servers:

✍ Refer to Starting Servers: Before You Begin for prerequisite tasks.

✍ If you have not already done so, create one or more Managed Servers.

See Creating WebLogic Domains Using the Configuration Wizard or "Create Managed Servers" in the Administration Console Online Help.

✍ Start the domain's Administration Server.

✍ In a shell (command prompt) on the computer that hosts the Managed Server, change to the directory that contains

the startManagedWebLogicscript:DOMAIN_NAME\bin\startManagedWebLogic.cmd (Windows)DOMAIN_NAME/bin/startManagedWebLogic.sh (UNIX)

where *DOMAIN_NAME* is the directory in which you located the domain. By default, this directory is *BEA_HOME*\user_projects\domains*DOMAIN_NAME*.

✍ Enter one of the following commands:

where *managed_server_name* specifies the name of the Managed Server and *admin_url* specifies the listen address (host name or IP address) and port number of the domain's Administration Server.

For example, the following command uses `startManagedWebLogic.cmd` to start a Managed Server named `myManagedServer`. The listen address for the domain's Administration Server is `AdminHost:7001`:

```
c:\bea\user_projects\domains\mydomain\bin\startManagedWebLogic.cmd
```

```
myManagedServer http://AdminHost:7001
```

- ✍ For each Managed Server that you want to start, open a separate command shell and follow steps 4 and 5. If you are starting Managed Servers on another machine, log in to that machine (remotely or locally) and then follow steps 4 and 5.

Reference: Starting Managed Servers with a Startup Script

Question No : 6

An application requires a JDBC data source, which you create in the domain. You name it but the administrator allot another domain name the same data source `DatasourceB`.

Without modifying either domain, which is the best method to ensure that the application deploys successfully in both?

- A. Create a separate application archive to use with each domain.
- B. Create a separate deployment plan file to use with each domain.
- C. Create and register a custom deployment listener.
- D. Modify the application code to determine the current domain.

Answer: B

Explanation: A WebLogic Server deployment plan is an optional XML document that you use to configure an application for deployment to a specific WebLogic Server environment, such as a domain.

Reference: Deploying Applications To WebLogic Server, Configuring Applications for Production Deployment


Question No : 7

Identify three options for Keystores in a server's configuration.

- A. Demo Identity and Demo Trust
- B. Demo Identity and Java Standard Trust
- C. Java Standard Identity and Custom Trust
- D. Custom Identity and Demo Trust
- E. Custom Identity and Custom Trust
- F. Custom Identity and Command Line Trust

Answer: A,E,F

Explanation: Please refer to the screenshot below:



Reference:

http://docs.oracle.com/cd/E23549_01/apirefs.1111/e13952/pagehelp/Corecoreserverserverconfigkeystore.html#attributes

Question No : 8

When installing Oracle Weblogic Server 11g with the graphical installer, which three statements are true?

- A. You must choose either a typical or a custom installation.
- B. You install under Microsoft Windows because the graphical installer is available only for Window
- C. You may create a new middleware home directory or choose an existing one.
- D. You must register for critical security updates.
- E. You may install a JDK or choose one that was previously installed.

Answer: A,C,E

Explanation: Tasks WebLogic Server Product Installation Procedure

Step 1 - Obtain the appropriate installation file for your platform

Step 2 - Complete the installation planning requirements

Ensure that your system environment meets the requirements for the installation. Also determine your Oracle Middleware home directory, and product installation home directories.

(C) When you are installing WebLogic Server, you are prompted to choose an existing Middleware home directory or specify a path to create a new Middleware home directory. If you choose to create a new directory, the installation program automatically creates it for you.

Step 3 - Determine the appropriate installation mode for your situation

Step 4 - Determine the type of installation

(A) There are two types of installation Typical or Custom. The type of installation depends on the products and sub-components you want to install.

Step 5 - Install the software

Step 6 - Create a WebLogic domain

E: The JDK selection screen.

This screen is displayed only if you selected a Custom installation. It is not displayed if you are running an Upgrade installer.

This screen displays a list of JDKs. This list contains JDKs that are available for (and common across) previous installations of all the selected WebLogic Server components. Select the JDK that you want to install with the product.

Reference: Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server 11g Release 1, Running the Installation Program in Graphical Mode

Reference: Oracle Fusion Middleware Installation Guide for Oracle WebLogic Server 11g Release 1, Installation Overview

Question No : 9

Each domain has a directory called servers that contains subdirectories for each server in the domain. Your domain has a managed server named managed1. When will the managed1 directory be created?

A. When the managed1 server is defined in the configuration Wizard

B. When the domain is created

- C. When the administration server comes up for the first time
- D. When the managed1 server1 comes for the first time

Answer: D

Explanation: This directory is created when the server is started for the first time.

Question No : 10

Identify three properties required by the domain Configuration Wizard when creating a new domain.

- A. machine name
- B. Managed Server name
- C. domain startup mode
- D. domain name
- E. administrator username and password

Answer: C,D,E

Explanation: C: Choose the Startup Mode

Specify the startup mode for your domain.

Either Development or Production

D: The Create WebLogic Domain window prompts you to specify the name and pathname for the domain, and initiate its creation.

E: The Configure Administrator Username and Password window prompts you to specify a username and password to be used for starting the Administration Server.

Reference: Creating WebLogic Domains Using The Configuration Wizard

http://docs.oracle.com/cd/E13179_01/common/docs92/configwiz/newdom.html

Question No : 11

Which two statements are true about adding servers to a cluster?

- A. When you create a cluster, you must add at least one server to it.
- B. Only managed servers can be in a cluster.
- C. The administration server is automatically added to a new cluster.
- D. You must explicitly identify which servers belong to the cluster.

Answer: A,B

Explanation: A cluster must include at least one server.

A clustered is administered through an administration server, but the cluster consists of managed servers only.

Question No : 12

Consider a single JMS module that includes a connection factory as well as queue.

Based on your system's requirements, the connection factory needs to be made available to all servers in the domain. The queue, on the other hand, can be targeted to only a single JMS server domain.

What is the best way to achieve the desired JMS setup?

- A. Create two subdeployments within the JMS module.
- B. Target the JMS server to the JMS module.
- C. Target the JMS module to the entire domain.
- D. Replace the queue with a topic.

Answer: A

Explanation: A subdeployment is a mechanism by which JMS module resources (such as queues, topics, and connection factories) are grouped and targeted to a server resource (such as JMS servers, server instances, SAF agents, or a cluster).

For example, you can group a connection factory with stand-alone queues or topics in a subdeployment targeted to a specific JMS server, which guarantees that all these resources are co-located to avoid extra network traffic. Another advantage of such a configuration would be if the targeted JMS server needs to be migrated to another WebLogic server instance, then the connection factory and all its connections will also migrate along with the JMS server's destinations. However, when stand-alone queues or topics are members of a subdeployment, a connection factory can only be targeted to the same JMS server.

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, Configure subdeployments in JMS system modules

Question No : 13

Which three attributes are part of the configuration of an existing Managed Server and can be modified through the " administration console?

- A. name of the server
- B. the cluster it is in
- C. the machine it runs on
- D. its replication group
- E. whether it is in Development Mode or Production Mode
- F. that it should become the Administration Server for the domain

Answer: B,C,D

Explanation: B: Attribute label: Cluster

The cluster to which this server belongs. If set, the server will listen for cluster multicast events.

This is a configurable attribute.

C: Attribute label: Machine

The WebLogic Server host computer (machine) on which this server is meant to run.

This is a configurable attribute.

D: Server --> Configuration --> Cluster

A WebLogic Server cluster is a group of servers that work together to provide a more scalable, more reliable application platform than a single server.

Use this page to specify a server's primary and secondary replication group, cluster weight, and the interface address used to handle multicast traffic.

Attribute label: Replication Group

Description: Defines preferred clustered instances considered for hosting replicas of the primary HTTP session states created on the server.

Reference: BEA WebLogic Server 8.1 Documentation, Server --> Configuration --> General

Question No : 14

Which can be associated with multiple domains?

- A. Cluster
- B. Server Log
- C. Node Manager
- D. Administration Server

Answer: C

Explanation: A Node Manager process is not associated with a specific WebLogic domain but with a machine. You can use the same Node Manager process to control server instances in any WebLogic Server domain, as long as the server instances reside on the same machine as the Node Manager process.

Reference: Using Node Manager to Control Servers

Reference: Understanding WebLogic Server Clustering

http://docs.oracle.com/cd/E11035_01/wls100/cluster/overview.html

Question No : 15

Identify two supported methods of deploying a JMS module to a domain.

- A. Create a module by using the administration console.
- B. Load a module into the WebLogic database.
- C. Include a module file within a web application archive.
- D. Include a module file within an enterprise application archive.
- E. Define a module within an existing JDBC module.

Answer: A,D

Explanation: A: Main Steps for Creating Packaged JMS Application Modules

Follow these steps to configure a packaged JMS module:

If necessary, create a JMS server to target the JMS module to, as explained in "Configure JMS Servers" in the Administration Console Online Help.

Create a JMS system module and configure the necessary resources, such as queues or topics, as described in "Configure JMS system modules and add JMS resources" in the Administration Console Online Help.

The system module is saved in config\jms subdirectory of the domain directory, with a "-

jms.xml" suffix.

Copy the system module to a new location, and then:

Give the module a unique name within the domain namespace.

Delete the JNDI-Name attribute to make the module application-scoped to only the application.

Add references to the JMS resources in the module to all applicable J2EE application component's descriptor files, as described in Referencing a Packaged JMS Application Module In Deployment Descriptor Files.

Package all application modules in an EAR, as described in Packaging an Enterprise Application With a JMS Application Module.

Deploy the EAR, as described in Deploying a Packaged JMS Application Module.

D: JMS application modules can be packaged as part of an Enterprise Application Archive (EAR), as a packaged module. Packaged modules are bundled with an EAR or exploded EAR directory, and are referenced in the weblogic-application.xml descriptor.

The packaged JMS module is deployed along with the Enterprise Application, and the resources defined in this module can optionally be made available only to the enclosing application (i.e., as an application-scoped resource). Such modules are particularly useful when packaged with EJBs (especially MDBs) or Web Applications that use JMS resources. Using packaged modules ensures that an application always has required resources and simplifies the process of moving the application into new environments.

Reference: Packaging JMS Application Modules In an Enterprise Application

Question No : 16

Which statement is true about the java Node Manager?

- A. It runs on Windows not on Linux
- B. In its secure form, it uses one-way SSL.
- C. It does not use the nodemanager.domains file, whereas the script-based Node Manager.
- D. Unlike the script based Node Manager, it should not be run as an operating system service

Answer: B

Explanation: Node Manager security relies on a one-way SSL connection between the

client and server.

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, Configuring Java-based Node Manager Security

http://docs.oracle.com/cd/E21764_01/web.1111/e13740/java_nodemgr.htm#i1068397

Question No : 17

Which two statements are true about a WebLogic domain?

- A. It is the basic administrative unit of WebLogic Server.
- B. It is an optional administrative construct to organize groups of WebLogic Administration Servers.
- C. Administration Servers are defined within a domain, but Managed Servers are not.
- D. You must first have a domain to define a WebLogic Server cluster.
- E. Every domain has one Administration Server and at least one managed server.

Answer: A,D

Explanation: A: A domain is the basic administration unit for WebLogic Server instances.

D: To define a WebLogic Server cluster there must be a domain:

Note: Before you start the Managed Servers in a domain, start the Administration Server. When you start a standalone or clustered Managed Server, it contacts the Administration Server for its configuration information. In this way, the Administration Server operates as the central control entity for the configuration of the entire domain.

Reference: Configuring and Managing WebLogic Server

http://docs.oracle.com/cd/E13222_01/wls/docs81/adminguide/overview_domain.html

Question No : 18

Which two statements are true about multicast communication within a cluster?

- A. Multicast can be used for cluster heartbeat messages.

- B. Multicast is used to replicate session state between servers.
- C. Multiple clusters can share the same multicast address and port.
- D. You may disable multicast on individual cluster members.
- E. Multicast traffic is often limited to a single subnet.

Answer: A,E

Explanation: A: WebLogic Server instances in a cluster communicate with one another using two basic network technologies:

* IP sockets, which are the conduits for peer-to-peer communication between clustered server instances.

* IP unicast or multicast, which server instances use to broadcast availability of services and heartbeats that indicate continued availability.

E: In many deployments, clustered server instances reside within a single subnet, ensuring multicast messages are reliably transmitted. However, you may want to distribute a WebLogic Server cluster across multiple subnets in a Wide Area Network (WAN) to increase redundancy, or to distribute clustered server instances over a larger geographical area.

Reference: Using WebLogic Server Clusters, Communications In a Cluster

http://docs.oracle.com/cd/E11035_01/wls100/cluster/features.html

Question No : 19

Identify two true statements about an Oracle WebLogic Server JMS Connection Factory.

- A. A default JMS Connection Factory exists and is named weblogic.jms.connections
- B. A Connection Factory can be targeted to only one WebLogic Server.
- C. A JMS Connection Factory's Default Priority and Default Time-to-Live attribute cannot be modified after the Connection Factory is created.
- D. A JMS Connection factory can be created as a system module resource within the administration console.

Answer: C,D

Reference: Programming WebLogic JMS, Understanding WebLogic JMS

Question No : 20

In WLST, you enter the shutdown () command with no arguments. Which statement is true in this case?

- A. The server you are connected to will shut down.
- B. The command does a "Force" shutdown.
- C. The command blocks user interaction.
- D. The command has timeout of 30,000 milliseconds.

Answer: A

Explanation: The WLST life cycle shutdown command gracefully shuts down a running server instance or cluster.

Reference: WebLogic Scripting Tool ,Life Cycle Commands

Question No : 21

Which three statements are true about WebLogic users and groups?

- A. A user is associated with a single security provider.
- B. A user can be a member of several groups.
- C. A group can contain other groups.
- D. A group consists of a name and a password.
- E. A group is associated with multiple security providers.
- F. Both users and groups are assigned a keystore.

Answer: A,B,C

Explanation: B: For efficient security management, BEA recommends adding users to groups. A group is a collection of users who usually have something in common, such as working in the same department in a company.

C: Example of group nesting:

Every user is a member of the everyone group.

The users group is nested within the everyone group.

Note: Security Providers - are modules that provide security service to application to protect Weblogic resource. Types of security providers in WebLogic Server are Authentication Provider, Authorization Provider, Auditing Providers, Credential Mapping Provider, Identity Assertion Provider, Principal Validation Provider, Adjudication Providers, Role Mapping Providers, Certificate Lookup and Validation Providers, Keystore Providers and Realm Adapter providers

Question No : 22

Your production JMS server and/or its consumers are not able to handle the incoming message workload. The number of messages on the server never stabilizes and the server eventually becomes overload.

Which JMS server attribute will best help prevent the JMS server from being overloaded by producers?

- A. Producer Pause High
- B. Messages Threshold High
- C. Reconnect Polity
- D. Paging Directory
- E. Pool Maximum Capacity

Answer: B

Explanation: JMS Configuration option ' Messages Threshold High':

The upper threshold (number of messages stored in this JMS server) that triggers flow control and logging events. A value of -1 disables the events for this JMS server.

If the number of messages exceeds this threshold, the triggered events are:

Log Messages

- A message is logged on the server indicating a high threshold condition.

Flow Control

- If flow control is enabled, the JMS server becomes armed and instructs producers to begin decreasing their message flow.

Reference: Administration Console Online Help, JMS Server: Configuration: Thresholds and Quotas

Question No : 23

Which three statements are true about WebLogic clusters?

- A. Clusters provide clients with transparent failover.
- B. EJB applications can be targeted to an entire cluster.
- C. All cluster members must bind to the same port number.
- D. Cluster members replicate application data by using heartbeats.
- E. Cluster members can be associated with one or more domains.
- F. Web applications require a proxy in order to be used in a cluster.

Answer: A,B,C

Explanation: A: If an object is clustered, failover and load balancing for that object is available.

In a WebLogic Server cluster, application processing can continue when a server instance fails. You “cluster” application components by deploying them on multiple server instances in the cluster—so, if a server instance on which a component is running fails, another server instance on which that component is deployed can continue application processing.

The choice to cluster WebLogic Server instances is transparent to application developers and clients.

B: A clustered application or application component is one that is available on multiple WebLogic Server instances in a cluster. If an object is clustered, failover and load balancing for that object is available. Deploy objects homogeneously—to every server instance in your cluster—to simplify cluster administration, maintenance, and troubleshooting.

Web applications can consist of different types of objects, including Enterprise Java Beans (EJBs), servlets, and Java Server Pages (JSPs).

Note: Load balancing and failover for EJBs and RMI objects is handled using *replica-aware stubs*, which can locate instances of the object throughout the cluster.

C: Cluster Multicast Address and Port

Identify the address and port you will dedicate to multicast communications for your cluster.

Reference: Using WebLogic Server Clusters, Understanding WebLogic Server Clustering

Question No : 24

Which three statements are true when an Administration Server is compared to a managed server?

- A. A domain must have at least one of each type of server.
- B. Administration Server keeps a domain log, but a Managed Server does not.
- C. A Managed Server keeps a server log, but an Administration Server does not.
- D. When a Managed Server comes up, it asks its Administration Server for the latest configuration.
- E. An Administration Server manages the domain configuration.
- F. JEE applications are deployed only to a Managed Server.

Answer: D,E,F

Explanation: D: When a Managed Server starts, it contacts the Administration Server to retrieve its configuration information. If a Managed Server is unable to connect to the specified Administration Server during startup, it can retrieve its configuration directly by reading a copy of the config.xml file and other files located on the Managed Server's own file system.

E: One instance of WebLogic Server in each domain acts as an Administration Server. The Administration Server provides a central point for managing a WebLogic Server domain.

F: For a typical production system, BEA recommends that you deploy your applications only on Managed Servers. This practice allows you to dedicate the Administration Server to configuration and monitoring of the domain, while one or more Managed Servers service your applications.

Reference: Overview of WebLogic Server System Administration

Question No : 25

Which three tasks can be performed by the Node manager?

- A. Start a server.
- B. Define a node server.
- C. Host the Node Manager console.

- D. Automatically restart a failed server.
- E. Kill a failed application on a server.
- F. Kill the process of a server that did not shut down properly.

Answer: A,B,D

Explanation: Node Manager enables you to perform these tasks:

- * Start and stop remote Managed Servers. (A)
- * Monitor the self-reported health of Managed Servers and automatically kill server instances whose health state is "failed".
- * Automatically restart Managed Servers that have the "failed" health state, or have shut down unexpectedly due to a system crash or reboot. (D)

Question No : 26

You have successfully created a global data source and connection pool via the administration console.

In which directory relative to the domain will this new data source and connection pool definition located?

- A. in a directory named config/dataSource
- B. in a directory named config/connPool
- C. in a directory named config/resources
- D. in a directory named config/jdbc
- E. in a directory named config

Answer: D

Explanation: When you create a JDBC resource (data source or multi data source) using the Administration Console or using the WebLogic Scripting Tool (WLST), WebLogic Server creates a JDBC module in the config/jdbc subdirectory of the domain directory, and adds a reference to the module in the domain's config.xml file.

Reference: Configuring WebLogic JDBC Resources

Question No : 27

What is the maximum number of targets (standard or migrate) that you can configure for a JMS server?

- A. 0
- B. 1
- C. 2
- D. No limit

Answer: D

Explanation: To configure the migratable target servers for JMS-related service migration:

1. If you have not already done so, in the Change Center of the Administration Console, click Lock & Edit (see Use the Change Center).
2. In the Domain Structure tree, expand Environment, then select Migratable Targets.
3. On the Summary of Migratable Targets page, click New.

On the Create a new Migratable Target page:

In Name, enter a name for the migratable target.

In Cluster, select a configured cluster for the migratable target.

Reference: Configure migratable targets for JMS-related services

Question No : 28

You are deploying a web application called inventory.war. It has been installed and has the Stat* "Prepared." It is targeted to the Managed Server named managed1. The web application's context route 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. What host and port would you use?

- A. managed, host and port
- B. Administration Server host and port
- C. Managed1 host and Administration port
- D. Administration Server host and Administration port

Answer: D

Explanation: Servicing only administrative requests: Specifies that WebLogic Server make the application available in Administration Mode only.

While in Administration mode, the application can be accessed only by internal clients through a configured Administration port.

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, A weblogic.Deployer Command-Line Reference

Question No : 29

You deployed a simple web application WAR by using the administration console. Its state is currently "Active."

In the administration console, under Configuration of the application, you change some deployment descriptor values.

What happens when you save such changes?

- A. The changes are in memory and temporary
- B. This is not possible because changes cannot be made to an "Active" application.
- C. You are prompted to select a location for a new deployment plan where the changes will be stored.
- D. The archive is opened and new versions of the deployment descriptors are placed within it.

Answer: C

Explanation: Note: You can use the Administration Console to configure selected deployment descriptor element values for an EJB that is deployed as an exploded application.

To configure editable deployment descriptor values for an EJB that is deployed as an exploded application:

1. If you have not already done so, in the Change Center of the Administration Console, click Lock & Edit (see Use the Change Center).
2. In the left pane of the Administration Console, select Deployments.
3. In the right pane, click the desired EJB module.

4. Select the Configuration > General tab.

Selected deployment descriptor elements for the EJB are listed.

5. Enter new values for the elements you want change, and click Save.

6. To activate these changes, in the Change Center of the Administration Console, click Activate Changes.

Not all changes take effect immediately—some require a restart

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, Configure deployment descriptor values

Question No : 30

Your server includes an application that uses a generic JDBC data source to connect to a database. However, the network connection between the server and the database is not 100% reliable.

You wish to avoid scenarios in which the application fails due to a failed JDBC connection. Identify two data source attributes to help accomplish this.

- A. Statement Cache
- B. Pool Reserve Mode
- C. Test Table Name
- D. Test Connection Mode
- E. Test Connections on Reserve

Answer: A,E

Explanation: A: Each connection in a data source has its own individual cache of prepared and callable statements used on the connection. However, you configure statement cache options per data source. That is, the statement cache for each connection in a data source uses the statement cache options specified for the data source, but each connection caches its own statements.

E: When Test Connections On Reserve is enabled, when your application requests a connection from the data source, WebLogic Server tests the connection using the query specified in Test Table Name before giving the connection to the application.

Testing reserved connections can cause a delay in satisfying connection requests, but it makes sure that the connection is viable when the application gets the connection.

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, Configuring JDBC Data Sources

http://docs.oracle.com/cd/E12839_01/web.1111/e13737/jdbc_datasources.htm#i1192665

Question No : 31

You have selected Configuration Archive Enabled for the domain. When a configuration change and activated, where is the previous version of the configuration archived?

- A. in the Oracle database
- B. in the config directory
- C. in the pending directory
- D. in the console-ext directory
- E. in the configArchive directory

Answer: E

Explanation: You can configure WebLogic Server to make backup copies of the configuration files. This facilitates recovery in cases where configuration changes need to be reversed or the unlikely case that configuration files become corrupted. When the Administration Server starts up, it saves a JAR file named config-booted.jar that contains the configuration files. When you make changes to the configuration files, the old files are saved in the configArchive directory under the domain directory, in a JAR file with a sequentially-numbered name like config-1.jar.

Reference: Understanding Domain Configuration, Configuration File Archiving

http://docs.oracle.com/cd/E12840_01/wls/docs103/domain_config/config_files.html

Question No : 32

When a typical application utilizes a JDBC data source, the process involves several steps. Identify three of these steps.

- A. The application returns the connection to the data source.

- B. The application looks up the data source using the JNDI name.
- C. The application tests the connection to verify the database's availability.
- D. The application adds a new connection to the data source if none are available.
- E. The application requests a connection from the data source.

Answer: B,D,E

Explanation: Applications look up the data source on the JNDI tree or in the local application context (B) and then reserve a database connection with the getConnection (E) method. Data sources and their connection pools provide connection management processes that help keep your system running and performant.

B: JNDI Name - Enter the JNDI path to where this JDBC data source will be bound. Applications look up the data source on the JNDI tree by this name when reserving a connection.

Note: If you've done any work with JDBC DataSources in a Web application, you already understand the basic process: create an initial JNDI (Java Naming and Directory Interface) context, then use the context to perform a "lookup" to retrieve your DataSource.

Reference: Administration Console Online Help, Create JDBC data sources

http://docs.oracle.com/cd/E15051_01/wls/docs103/ConsoleHelp/taskhelp/jdbc/jdbc_datasources/CreateDataSources.html

Question No : 33

As an Oracle WebLogic Server administrator, you are given an EAR file application to deploy. Identity three reasons for creating an EAR file first.

- A. After the .ear file has been deployed, the developer will not have to ask for your assistance to deploy it again.
- B. The EAR file is created to avoid namespace clashes in the Java code.
- C. The EAR file can contain multiple WAR, EAR, and other resources to easily bundle together.
- D. The EAR file is required if you are deploying an application.
- E. The EAR file can declare application-wide security roles definitions.

Answer: A,B,E

Explanation:

E: An EARfile contains:

- * Web module

A WAR file containing the view logics like JSPs, Servlets, Html contents. Every web module must have one Web Deployment descriptor placed inside “WEB-INF/web.xml”

web.xml is a J2EE standard deployment descriptor. Among other settings, it has a set of elements for configuring security for the web application.

Note:

- * An EAR file provides a convenient way to bundle up all pieces of a J2EE application. It may contain Java Module, WAR Module or an EJB Jar module.

An EAR can contain at least any one of these modules. EAR may contain more Modules as well.

- * To bring your portal online in a production environment, it is first necessary to prepare your portal application. Typical preparation steps include modifying deployment descriptors for the product, building the Enterprise archive (EAR) with all its pre-compiled classes, and deciding if you want to compress that EAR into an archive or leave it exploded.

Reference: Preparing and Deploying the EAR File

Question No : 34

Identify three types of JMS resources that can be configured as a part of a JMS module.

- A. JMS Distributed Destination
- B. JMS Connection Factory
- C. JMS Data Source
- D. JMS Destination
- E. JMS Server

Answer: A,B,D

Explanation: Configurable JMS Resources in Modules

The following configuration resources are defined as part of a system module or an application module:

- * Distributed destinations (A, D)
- * Connection factories (B)

- * Queue and topic destinations
- * Templates
- * Quota
- * Foreign servers
- * MS store-and-forward (SAF) configuration items

All other JMS environment-related resources must be configured by the administrator as domain configuration resources. This includes:

- * JMS servers required (not E)
- * Store-and-Forward agents (optional)
- * Path service (optional)
- * Messaging bridges (optional)
- * Persistent stores (optional)

Reference: Configuring And Managing WebLogic JMS, Configurable JMS Resources in Modules

Question No : 35

A managed server, myserver1, has a boot.properties file in the security directory. It was started with the startManageWeblogic.sh script(.cmd in windows) and his boot.properties file was used for its startup credentials.

You just used the administration console to change all administrator passwords. To continue using boot.properties, what can you do?

- A.** This is not possible. A boot.properties file can be used only with the Administration Server.
- B.** Delete boot.properties. In the administration console, under the myserver configuration, select Generate Boot Identity file.
- C.** You need not do anything- The password in boot .properties was automatically updated by administration console when you changed the password.
- D.** Edit boot.propetties. Type over the encrypted password with the new password in clear text. The next time myserver1 is started, it will encrypt the password in the file.
- E.** Delete boot .properties. Use the WLST encrypt () command to create a new boot.properties file containing the new password. Copy that file into the security directory of myserver1

Answer: D

Explanation: If you install the WebLogic Server Examples component, the default user weblogic is created that has permission to start and stop WebLogic Server. The default password is welcome1. If you change the password of the weblogic user, WebLogic Server does not automatically update this password in the boot.properties file, which is located in the DOMAIN_NAME/servers/AdminServer/security directory.

If you change the password for user weblogic, you can use either of the following workarounds so that you can continue to boot a WebLogic Server instance via that username and its new password:

* Remove the boot.properties file. Subsequently each time you start WebLogic Server, you are prompted for the username and password. The changed password for the weblogic user will be accepted.

* Modify the existing boot.properties file, changing the username and password as follows:

username=weblogic

password=welcome1

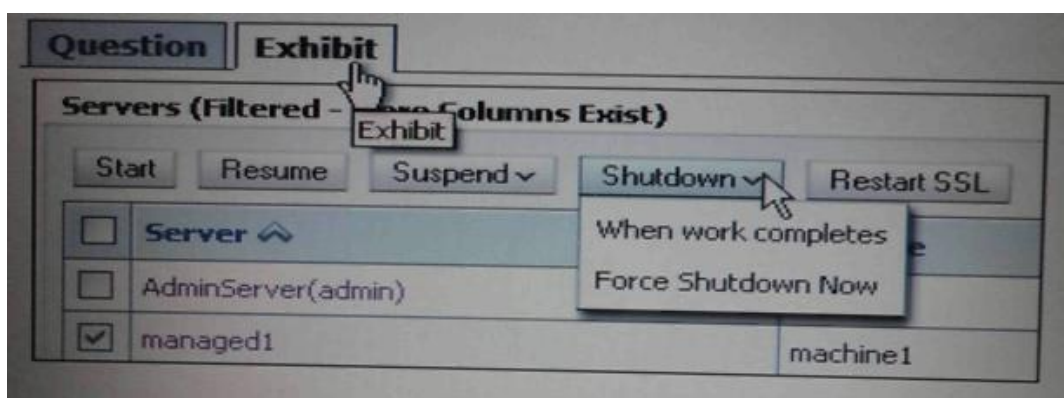
Subsequently during the server startup process, the boot.properties file is encrypted again.

Reference: Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help 11g Release 1, Limitation Regarding User weblogic

http://docs.oracle.com/cd/E15523_01/web.1111/e13708/overview.htm

Question No : 36

Refer to the Exhibit.



Which three statements are true about using the administration console to shut down a server?

- A. The configuration must first be locked.
- B. The choice "Force Shutdown Now" drops in-work requests.
- C. Both the choices, "When work completes" and "Force Shutdown Now", reject any new requests.
- D. The choice "When work completes" allows in work requests to complete before the server down.
- E. The administration console can be used to shut down Managed Servers, such as managed not the Administration Server.

Answer: B,C,D

Explanation: A: See step 3) in the note below.

B: select Force Shutdown Now to stop the server immediately without completing ongoing tasks.

D: Select When Work Completes to gracefully shut down the server

A graceful shutdown gives WebLogic Server subsystems time to complete certain application processing currently in progress.

Question No : 37

Your company is developing batch order system that utilities JMS. Each order message that is created will be broadcasted to and consumed by several order processing and billing applications.

As part of configuring the domain that will support this system, which type of JMS destination resource should you create?

- A. Grid
- B. Queue
- C. Topic
- D. Batch
- E. Client

Answer: C

Explanation: A JMS destination identifies a queue (point-to-point) or topic (publish/subscribe) resource within a JMS module.

A JMS topic identifies a publish/subscribe destination type for a JMS server. Topics are used for asynchronous peer communications. A message delivered to a topic is distributed to all consumers that are subscribed to that topic.

Reference: Configuring Basic JMS System Resources, Queue and Topic Destination Configuration

Question No : 38

You monitor a running JDBC data source with the console and note the following values:

Current Capacity = 10

Current Capacity High Count = 20

Identify two plausible explanations for these metrics.

- A. The pool size has increased.
- B. The pool size has decreased.
- C. The workload was higher than it is now.
- D. The workload was lower than it is now.
- E. The statement cache size is 10.
- F. The statement cache size is 20.

Answer: B,C

Explanation: The pool size has decreased from a maximum 20 to 10. The workload has been 20. Now it is 10. The workload was higher.

Note: Current Capacity: The current count of JDBC connections in the connection pool in the data source.

Current Capacity High Count: Highest number of database connections available or in use (current capacity) in this instance of the data source since the data source was deployed.

Note:

Reference: Administration Console Online Help, JDBC Data Source: Monitoring: Statistics