



KaozhengPro

IT認證考試題庫 專業平臺

考證寶提供最新考古題與模擬試題
協助您高效通過認證考試

www.kaozhengpro.com

Exam : **1Z0-599**

Title : Oracle WebLogic Server
12c Essentials

Version : Demo

1.You deploy more than one application to the same WebLogic container. The security is set on JavaEE level and all deployed JavaEE applications use the same security roles.

What is your recommendation for an architecture with those requirement

- A. Combine all applications into a single one.
- B. Define global roles on the WebLogic Domain level.
- C. Use Ms Active Directory to keep the roles there.
- D. Use Oracle Identity and Access Management solution to simplify the management.
- E. Keep role mapping in the external WebLogic Role Mapped developed for that solution.

Answer: B

Explanation/Reference

Note:

* Types of Security Roles: Global Roles and Scoped Roles

There are two types of security roles in WebLogic Server:

/ A global security role can be used in any security policy. Oracle provides several default global roles that you can use out of the box to secure your WebLogic resource

/ A scoped role can be used only in policies that are defined for a specific instance of a WebLogic resource (such as a method on an EJB or a branch of a JNDI tree). You might never need to use scoped roles. They are provided for their flexibility and are an extra feature for advanced customers.

Incorrect:

Not E: Role mapping is the process whereby principals (users or groups) are dynamically mapped to security roles at runtime. In WebLogic Server, a Role Mapping provider determines what security roles apply to the principals stored a subject when the subject is attempting to perform an operation on a WebLogic resource. Because this operation usually involves gaining access to the WebLogic resource, Role Mapping providers are typically used with Authorization providers.

2.When your WebLogic Server solution needs to be scaled out with additional capacity and you don't want to add additional hardware, which three techniques should you us?

- A. Assign more than one managed server to a physical hardware that allows better CPU utilization.
- B. Assign more than one application to one managed server to better utilize threads within a single JVM process.
- C. Assign the same application to more than one managed server to load balance requests between servers.
- D. Assign the same heap size to the managed server across the cluster for easier control of memory footprint
- E. Create a virtualized environment with hypervisor for an easier solution

Answer: B, D, E

3.A highly available WebLogic cluster in UNIX is configured for automatic server migration. Node Managed is configured on both machines to start managed servers.

How should you simulate a managed server failure to test whether automatic server migration is working?

- A. Shut down the managed server from the WebLogic console.
- B. Shut down the managed server using the WLST command through Node Manager.
- C. Run "kill -9" once to kill the managed server process.
- D. Run "kill -9" to kill the managed server process, and run "kill -9" one more time if the managed server is

restarting.

Answer: A

Explanation/Reference

Note:

- * It is recommended that you shutdown WebLogic Server instances through the Administration Console.
- * If automatic server migration is enabled, the servers are required to contact the cluster leader and renew their leases periodically. Servers will shut themselves down if they are unable to renew their leases. The failed servers will then be automatically migrated to the machines in the majority partition.

4.A customer claims that while redeploying a web application in the production system all their customers are having to log in again.

What do you recommend?

- A. Sessions can't be preserved when redeploying applications. The customer needs to consider redeployment during late nights when the traffic is low.
- B. Change the flag responsible for the development mode of their environment. In the production mode, all sessions are preserved while redeploying application.
- C. Change Hotspot to JRockit. Sessions can't be preserved on HotSpot when redeploying application.
- D. Use flag -version when redeploying the application. This will switch on the Side By Side deployment feature and preserve existing sessions.
- E. Open a service request with Oracle Support. This is unexpected behavior. Sessions are preserved without any extra settings.

Answer: D

Explanation/Reference

Restrictions on Production Redeployment Updates

WebLogic Server can host a maximum of two different versions of an application at one time.

Note:

- * When you redeploy a new version of an application, you cannot change:

An application's deployment targets

An application's security model

A Web application's persistent store settings

To change any of the above features, you must first undeploy the active version of the application.

Incorrect:

- * (not A)

Production redeployment enables you to update and redeploy an application in a production environment without stopping the application or otherwise interrupting the application's availability to clients.

Production redeployment saves you the trouble of scheduling application downtime, setting up redundant servers to host new application versions, manually managing client access to multiple application versions, and manually retiring older versions of an application.

- * (not C) Not dependant on whether the application is JRockit or Hotspot.

- * (not E)

The production redeployment strategy is supported for:

Standalone Web Application (WAR) modules and enterprise applications (EARs) whose clients access the application via a Web application (HTTP).

Enterprise applications that are accessed by inbound JMS messages from a global JMS destination, or from inbound JCA requests.

All types of Web Services, including conversational and reliable Web Services, but not 8.x Web Services. Production redeployment is not supported for:

Standalone EJB or RAR modules. If you attempt to use production redeployment with such modules, WebLogic Server rejects the redeployment request. To redeploy such modules, remove their version identifiers and explicitly redeploy the modules.

Applications that use JTS drivers. For more information on JDBC application module limitations, see JDBC Application Module Limitations in Configuring and Managing JDBC Data Sources for Oracle WebLogic Server.

Applications that obtain JDBC data sources via the DriverManager API; in order to use production redeployment, an application must instead use JNDI to look up data sources.

Applications that include EJB 1.1 container-managed persistence (CMP) EJBs. To use production redeployment with applications that include CMP EJBs, use EJB 2.x CMP instead of EJB 1.1 CMP.

Reference: Reference; Deploying Applications to Oracle WebLogic Server 12c, Redeploying Applications in a Production Environment

5.A common approach to solve application performance issues is to add more hardware.

Which two reasons describe why this alone does not always solve performance problems?

- A. There may be a bottleneck in the application that additional hardware would not solve.
- B. Adding more hardware is costly and time consuming.
- C. Adding hardware at the application layer may place additional load on an already overloaded database, network appliance, or storage system.
- D. Adding more hardware may introduce new configuration requirements that need to be tested.

Answer: A,C

6.A customer has a critical, performance-sensitive web application that connects to a multinode Oracle RAC database.

Which feature of WebLogic can provide significant performance benefit?

- A. The Web Session Affinity feature of Active GridLink for RAC
- B. WebLogic Clustering
- C. The Transaction Affinity feature of Active GridLink for RAC
- D. Coherence*Web Session Replication

Answer: C

Explanation/Reference

XA Affinity and Failover

When accessed within a global transaction, the member data source from which the JDBC connection was obtained is pinned to the global transaction for the life of the transaction. This ensures that all database operations performed on connections obtained from the Multi Data Source, for a particular transaction, all execute on the same RAC instance. XA affinity results in improved performance and is even a requirement for older versions of RAC, such as prior to 11g.

The XA failover is also supported by the Multi Data Source and transaction manager implementations. If a pinned RAC instance suffers a failure, then a global transaction can complete utilizing a different RAC instance using a connection obtained one of the other member data sources.

Reference: Oracle WebLogic Server Active GridLink for Oracle Real Application Clusters(RAC)

7.A customer needs to implement a Highly Available solution for JMS that has a primary data center and a backup.

Which three steps would you perform when designing your solution?

- A. Store Transaction Logs in a database and use Database stores for JMS to make replication between sites easier.
- B. Use file based Transaction Logs and JMS stores and implement a separate replication solution for files in addition to database in case database replication fails.
- C. Implement Oracle RAC at each site to provide a highly available solution within each datacenter.
- D. Configure Whole Server Migration to migrate WebLogic Managed Servers from the primary to the secondary site.
- E. Configure Automatic Service Migration for JMS high availability within a datacenter.

Answer: A, C, E

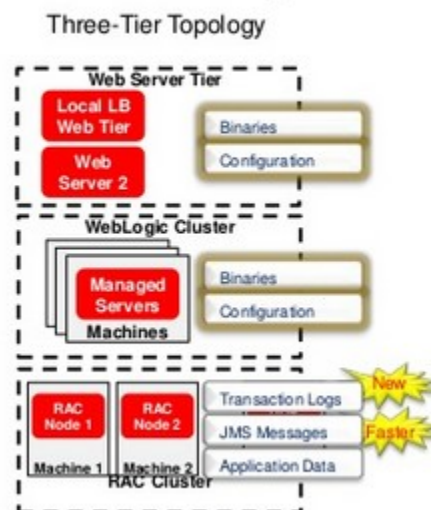
Explanation/Reference

Note:

WebLogic Server 12c Disaster Recovery

Features from WebLogic Server 10.3.6

- Transaction Logs in Database
 - Option to use DB instead of file store
 - Enables common storage, replication with application and JMS data
 - Configuration benefits
 - Supported with RAC, Exalogic, Exadata
 - Maximum performance with GridLink
- JDBC Store Performance
 - Optimizations for DB Store use cases
 - Enabler for new DR architectures
 - Supported with RAC, ExaLogic, Exadata
 - Maximum performance with GridLink



8.Which WebLogic optimization allows one non-XA resource to participate in a distributed transaction?

- A. enabling Pinned to Thread
- B. enabling Logging Last Resource
- C. increasing the Statement cache size
- D. setting the statement cache type to LRU
- E. setting the initial and maximum capacity to the same number

Answer: A

Explanation/Reference

XA Affinity and Failover

When accessed within a global transaction, the member data source from which the JDBC connection was obtained is pinned to the global transaction for the life of the transaction. This ensures that all database operations performed on connections obtained from the Multi Data Source, for a particular

transaction, all execute on the same RAC instance. XA affinity results in improved performance and is even a requirement for older versions of RAC, such as prior to 11g.

The XA failover is also supported by the Multi Data Source and transaction manager implementations. If a pinned RAC instance suffers a failure, then a global transaction can complete utilizing a different RAC instance using a connection obtained one of the other member data sources.

Reference: Oracle WebLogic Server Active GridLink for Oracle Real Application Clusters(RAC)

9.A customer has a Stock Watch application that publishes stock recommendations to different customers and programs. The stock recommendation message should be processed by only one of the JMS Servers in the cluster.

Which JMS model should be used?

- A. JMS Queue
- B. Distributed JMS Queue
- C. Uniform Distributed Topic
- D. Partitioned Distributed Topic

Answer: A

Explanation/Reference

JMS queue

A staging area that contains messages that have been sent and are waiting to be read. Note that, contrary to what the name queue suggests, messages don't have to be delivered in the order sent. A JMS queue only guarantees that each message is processed only once.

Incorrect:

Not B: (Only one queue required)

* A distributed destination is a set of destinations (queues or topics) that are accessible as a single, logical destination to a client. A distributed destination has the following characteristics:

It is referenced by its own JNDI name.

Members of the set are usually distributed across multiple servers within a cluster, with each destination member belonging to a separate JMS server.

* A distributed queue is a set of physical JMS queue members. As such, a distributed queue can be used to create a QueueSender, QueueReceiver, and a QueueBrowser. The fact that a distributed queue represents multiple physical queues is mostly transparent to your application.

Not Topic:

JMS topic A distribution mechanism for publishing messages that are delivered to multiple subscribers.

10.A customer has a critical, performance-sensitive web application that connects to a multimode Oracle RAC database.

Which feature of WebLogic can provide signification benefit?

- A. The Web Session Affinity feature of Active GridLink for RAC.
- B. WebLogic Clustering
- C. The Transaction Affinity feature of Active GridLink for RAC
- D. Coherence*Web Session Replication

Answer: C

Explanation/Reference

Active GridLink for Oracle RAC

In Oracle WebLogic Server 10.3.4, a single data source implementation has been introduced to support an Oracle RAC cluster. It responds to FAN events to provide Fast Connection Failover (FCF), Runtime Connection Load-Balancing (RCLB), and RAC instance graceful shutdown. XA affinity is supported at the global transaction Id level. The new feature is called WebLogic Active GridLink for RAC; which is implemented as the GridLink data source within WebLogic Server.

Note:

* The WebLogic Server JDBC subsystem has supported Oracle RAC since WLS version 9.0, originally developed for Oracle9i RAC. This support is based on a particular type of data source configuration, called a multi data source. A multi data source is a data source abstraction over one or more individual data sources. It serves JDBC connections from each of the member data sources according to a specified policy. A RAC multi data source configuration requires that each member data source obtain connections to a particular RAC instance.

Reference: How-To: Use Oracle WebLogic Server with a JDBC GridLink Data Source

11.A customer needs to ensure that the number of threads servicing an application does not exceed the number of database connections available to the application.

What step must you take to address this situation?

- A. Configure a Max Threads Constraint and add your application to the list of applications for the Constraint.
- B. Configure a Work Manager with a Maximum Threads Constraint tied to the Connection Pool and configuration your application to use the Work Manager.
- C. Configure a Work Manager with a Minimum Threads Constraint tied to the Connection Pool and configure your application to use the Work Manager.
- D. Configure a global MaxThreads constraint and target it to the server or clusters where your application is deployed.
- E. Configure the startup parameter "-Dwls-maxThreads" to be the same as the number of database connections configured.

Answer: B

Explanation/Reference

To manage work in your applications, you define one or more of the following Work Manager components:

Fair Share Request Class:

Response Time Request Class:

Min Threads Constraint:

Max Threads Constraint:

Capacity Constraint

Context Request Class:

Note:

* Max-threads-constraint—This constraint limits the number of concurrent threads executing requests from the constrained work set. The default is unlimited. For example, consider a constraint defined with maximum threads of 10 and shared by 3 entry points. The scheduling logic ensures that not more than 10 threads are executing requests from the three entry points combined.

A max-threads-constraint can be defined in terms of the availability of resource that requests depend upon, such as a connection pool.

A max-threads-constraint might, but does not necessarily, prevent a request class from taking its fair share of threads or meeting its response time goal. Once the constraint is reached the server does not schedule requests of this type until the number of concurrent executions falls below the limit. The server then schedules work based on the fair share or response time goal.

* WebLogic Server prioritizes work and allocates threads based on an execution model that takes into account administrator-defined parameters and actual run-time performance and throughput.

Administrators can configure a set of scheduling guidelines and associate them with one or more applications, or with particular application components.

* WebLogic Server uses a single thread pool, in which all types of work are executed. WebLogic Server prioritizes work based on rules you define, and run-time metrics, including the actual time it takes to execute a request and the rate at which requests are entering and leaving the pool.

The common thread pool changes its size automatically to maximize throughput. The queue monitors throughput over time and based on history, determines whether to adjust the thread count. For example, if historical throughput statistics indicate that a higher thread count increased throughput, WebLogic increases the thread count. Similarly, if statistics indicate that fewer threads did not reduce throughput, WebLogic decreases the thread count. This new strategy makes it easier for administrators to allocate processing resources and manage performance, avoiding the effort and complexity involved in configuring, monitoring, and tuning custom executes queues.

Reference: Using Work Managers to Optimize Scheduled Work

12. What are the three steps you should take to tune a JDBC Connection pool in WebLogic from the initial settings in a production environment?

- A. Ensure the maximum size is increased to an appropriate setting.
- B. Set the minimum and maximum size of the connection pool to the same value.
- C. Increase the statement cache size.
- D. Add more heap to the JVM.
- E. Add more nodes to the cluster.

Answer: A, C, E

Explanation/Reference

A:

* Troubleshooting Slow Response Time from the Client and Low Database Usage

These symptoms are usually caused by a bottleneck upstream of the database, perhaps in the JDBC connectionpooling. Monitor the active JDBC connections in the WebLogic Console and watch for excessive waiters and wait times; increase the pool size, if necessary.

* Attribute: Maximum Capacity

Maximum number of physical database connections that this connection pool can contain. Different JDBC Drivers and database servers may limit the number of possible physical connections.

C: Attribute: Statement Cache Size

The algorithm used to maintain the statement cache:

LRU - After the statementCacheSize is met, the Least Recently Used statement is removed when a new statement is used.

Fixed - The first statementCacheSize number of statements is stored and stay fixed in the cache. No new statements are cached unless the cache is manually cleared.

E: If the queue appears starved but adding execute threads does not improve performance, there may be

resource contention. Because CPU utilization is low, the threads are probably spending much of their time waiting for some resource, quite often a database connection.

Use the JDBC monitoring facilities in the console to check for high levels of waiters or long wait times. Adding connections to the JDBC connection pool may be all that is required to fix the problem.

Note:

* If you had a JDBC connection pool where the Initial Capacity and Maximum Capacity attributes were different, you might want to create a gauge monitor to monitor the maximum and minimum number of connections.

By setting the Threshold Low value to be one less than the Initial Capacity, your gauge monitor trap could monitor the ActiveConnectionsCurrentCount attribute of the JDBCDataSourceRuntime MBean and alert you whenever the number of active connections are less than the Initial Capacity (which might indicate database connectivity problems).

13. You have a durable subscriber, and the subscriber is down or not reachable when the message is produced.

Which two options regarding the expiry of these messages are true?

- A. after the subscriber is unavailable for 10 minutes
- B. when the subscriber is available
- C. after the subscriber is unavailable for after an hour
- D. are available until the specified time elapses
- E. are expired instantly

Answer: B, D

Explanation/Reference

By default, JMS messages never expire. When applications send messages to queues or topics with durable subscribers, WebLogic must retain the message until it is consumed. This is fine in most point-to-point messaging applications because consumers are constantly consuming messages. Any message sent to a queue will typically be consumed in a relatively short period of time. If the consumers get disconnected, they will usually reconnect as soon as possible and start processing any messages that might have built up in the queue.

D: For durable subscribers to a topic, this is not necessarily true. The messaging system is forced to retain any message that has not been consumed by a durable subscriber, regardless of whether that durable subscriber will ever return. In this case, WebLogic is at the mercy of the durable subscriber to unsubscribe when it no longer wishes to receive the messages. If the durable subscriber logic is flawed in such a way that the subscribers do not unsubscribe properly, the messaging system will start to fill up with messages that may never be delivered. This calls for real caution in using durable subscribers.

Fortunately, there is another way to help deal with this problem. Message expiration can be set at the connection factory level. Using a connection factory's default time-to-live attribute, we can specify the number of milliseconds that WebLogic should retain an undelivered message after it is sent.

14. Identify the two options that can be used to patch WebLogic Server without incurring application downtime.

- A. automated process using the Admin Server

- B. automated using JDeveloper
- C. manual process with or without Smart Update
- D. scripted with WLST
- E. automated process using Enterprise Manager 12c

Answer: C, D

Explanation/Reference

Note:

* Rolling Upgrade is the process of upgrading a running WebLogic Server cluster with a patch, maintenance pack, or minor release without shutting down the entire cluster or domain. During the rolling upgrade of a cluster, each server in the cluster is individually upgraded and restarted while the other servers in the cluster continue to host your application.

15. WebLogic 12c, the Maven plug-in has been enhanced with which Maven goal?

- A. wls:unzip, wls:install-domain, wls:start-domain, wls:wlst, wls:appc
- B. wls:install, wls:install-domain, wls:start-domain, wls:wlst, wls:appc
- C. wls:unzip, wls:create-domain, wls:start-domain, wls:wlst, wls:appc
- D. wls:install, wls:create-domain, wls:start-server, wls:wlst, wls:appc:

Answer: D

Explanation/Reference

wls-maven-plugin—Delivered in WebLogic Server 12c, provides enhanced functionality to install, start and stop servers, create domains, execute WLST scripts, and compile and deploy applications.

Reference: Developing Applications for Oracle WebLogic Server, Using the WebLogic Development Maven Plug-In