

Latest Version: 6

Question: 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 requirements?

- A. Combine all applications into a single one.
- B. Define global roles on the WebLogic Domain level.
- C. Use Microsoft 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:

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 as 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.

Question: 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 use?

- 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

Question: 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:

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.

Question: 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:

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

Question: 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

Question: 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: A

Explanation:

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

Question: 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

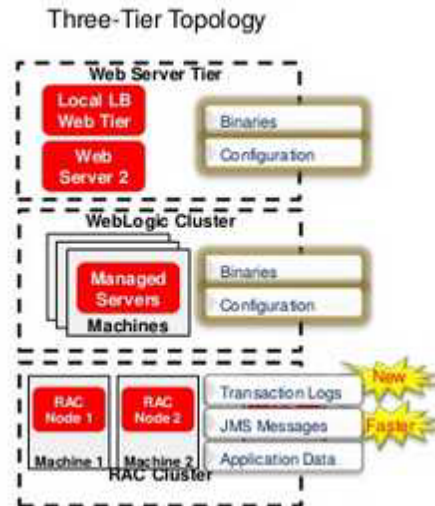
Explanation:

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



Question: 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: B

Explanation:

<http://my-java-planet.blogspot.com/2014/04/which-weblogic-optimization-allows-one.html>

Question: 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:

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.

Question: 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:

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