KillTest

Higher Quality, Better Service!

Q&A

http://www.killtest.com

We offer free update service for one year.

Exam : Oracle 1Z0-108

Title : Oracle WebLogic Server 10gSystem Administration

Version : DEMO

1. Scenario : A single tier WebLogic cluster is configured with six Managed Servers. An Enterprise application is deployed to the cluster with a web application and EJBs packaged in the same EAR file.

An EAR file is being deployed to the cluster using two-phase deployment. Which of the following is true if one of the Managed Servers is NOT started in the cluster?

A. First and second phase deployment will be completed on all remaining instances in the cluster.

B. First phase deployment will not be completed on the remaining instances in the cluster.

C. First phase deployment will be completed but the second phase will not be completed on the remaining instances in the cluster.

D. First and second phase deployment will not be completed if enforceClusterConstraints is set to true for the Deployer.

E. Both a and d

Answer: E

2. Which of the following statements are true for built-in roles in the WebLogic Server Administration console? Choose two.

A. Users who belong to the Operator role can deploy the applications.

B. Users who belong to the Monitor role can view the server configuration, except for the encrypted attributes.

C. Users who belong to the Admin role cannot deploy applications.

D. Only users who belong to the Admin role can start, stop, and resume the server.

E. Users who belong to the Admin role can modify the entire server configuration.

Answer: BE

3. Scenario : Consider an upgrade requirement from pre-WebLogic Server 10.x to the WebLogic Server 10.x. Since the administration scripts in the prior version used WebLogic Ant tasks for creation and deployment, the team decides to use the Ant build.xml approach and extend this to the current version of the server.

Which of the following about the Ant task wlserver is INCORRECT?

A. wlserver can be used to start, reboot or shut down server instances.

B. wlserver can be used to connect to existing server instances.

C. wlserver can be used to create a multiple server domain configuration.

D. wlserver can be used to create a single server domain and to connect to the server.

Answer: C

4. Click on the Exhibit button.

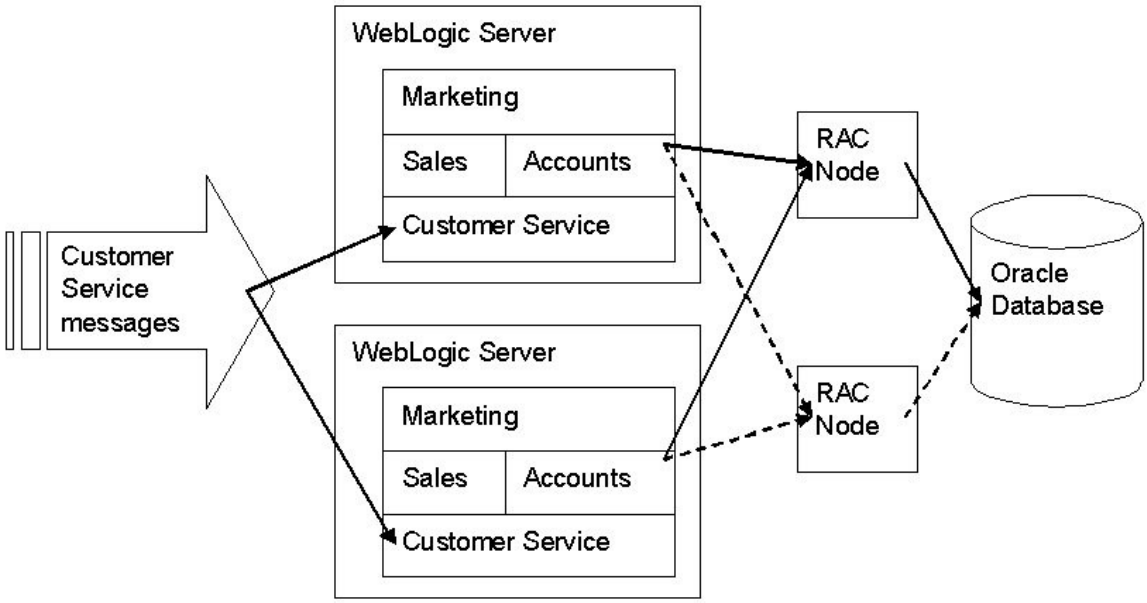

Scenario : You have made the decision to migrate all your operations onto WebLogic Server. An application has been chosen from each department for the initial phase of migration ?one from each of Sales, Marketing, Accounts, and Customer Service. The applications will be deployed as individual EAR files to a cluster of WebLogic Server 10.x instances, each with a pool of connections to an Oracle 10g database that provides corporate and departmental data services. The database team hopes to upgrade the database to use Oracle RAC (Real Application Clusters) shortly, which will require you to configure multiple sets of connections. Connectivity to the existing Service Engineer Dispatch system is via a JMS message bridge. When a service engineer calls in with a report, a message is sent via JMS to the Customer Service application deployed in WebLogic Server. The customer record in the Customer Service database must be updated with details from the message. It is important that the message is not lost because the call management system is not able to resend messages. To address this, the Customer Service application is accessing a persistent JMS Queue using an XA compatible connection factory.

The database security team insists that any use of the corporate Oracle database can be traced to the individual user that made the request. How can this be achieved when using WebLogic JDBC Data

Sources?

A. Set the Enable Credential Mapping property for the Data Source.

B. Set Initial Capacity to zero, so each user creates a new connection.

C. Set the value of user to {WLS_USER} in the Properties field.

D. This can only be done using application code.

E. This cannot be done because Data Sources connect using a single username.

Answer: A

5. In a multi-server WebLogic domain, monitoring the JVM statistics showed frequent full garbage collections (gc) on one of the server instances. Which one of the following would greatly reduce the frequency of full gcs?

A. Increasing the young object space in the heap

B. Turning the -verbosegc flag on

C. Adding more servers to the domain

D. Running more servers per machine

Answer: A

6. What are the four basic components of a JMS system-resource module in the config.xml in WebLogic 10.x?

A. Name, target, subdeployment, descriptor-file-name

B. Domain, jms-server, connection factory, destination

C. Config.xml, jms-module, subdeployment, descriptor-file-name

D. Name, jms-module, subdeployment, descriptor-file-name

Answer: A

7. The deployment configuration for an application or module is NOT stored in which type of the following XML document:

A. J2EE deployment descriptors

B. weblogic-*.xml descriptors

C. WebLogic Server deployment plans

D. WebLogic config.xml

Answer: D

8. Scenario : Consider a b2b commerce application hosted on a WebLogic cluster with JMS messages being received, processed, and forwarded to external servers. The server log indicates a potential stuck thread problem where threads are continually getting stuck and JMX monitoring indicates a continuous backlog of pending messages that could be due to any of the in-flight messages, delayed messages, or messages being inhibited from delivery.

The recommended approach for setting the number of execute threads would be:

A. Set the default thread pool max value to match load expectations

B. Use an iterative approach to set the max thread count value

C. Set this value the same as the max connections value set on database connections

D. Leave this setting to be handled by the Work Manager by configuring MaxThreadsConstraint

Answer: D

9. An application-wide deployment plan was used to deploy an application. Based on post-performance diagnosis, the HTTP post size needs to be reduced. Following would be a valid approach:

A. Configure the http maximum post size value in the existing deployment plan for redeploying using the weblogic.Deployer utility.

B. Update the http maximum post size for the application using weblogic.Deployer.

C. Use weblogic.PlanGenerator to update http maximum post size configuration and redeploy the application.

D. Use weblogic.Admin to set auto-tune-http-setting to be true and redeploy the application.

Answer: A

10. Scenario : The Stock Control Application you are going to deploy to a cluster of WebLogic Managed Servers requires a JMS destination to asynchronously deliver the messages to the back-end Warehousing Application. To distribute the messages across the cluster, you decided to configure the distributed destination. Also the Warehousing Application is a legacy application written to implement a proprietary MQ API connecting an MQ-Series server. The Stock Control Application is implemented as a web application on

a WebLogic cluster delivering the message to the distributed destination. The cluster has three Managed Servers hosted on two high powered machines and one low range/powered machine.

To distribute the physical destinations unevenly across the clustered WebLogic Managed Servers so that the high-powered machines can receive more messages than the low-powered machines, configure the distributed destination with:

A. "Allocate Members Uniformly" set to false and manually select more physical destinations from the high-powered machines

B. "Uneven Distribution" set to true and select all the physical destinations

C. "Uniform Distribute Destination" set to true and select all the physical destinations from the high-powered and low- powered machines

D. "Load Balancing Policy" set to Random and select only the destination from the high-powered machine

Answer: A

11. JMS interop modules are different from JMS system resource modules because (Choose two)

A. JMS interop modules cannot use "Unit-of-Order".

B. JMS interop modules can be targeted to multiple WebLogic Servers at the same time in a domain.

C. JMS interop module configuration is always stored as interop-jms.xml.

D. JMS interop module cannot use "Quota."

Answer: BC

12. Which of the following is true about JMS Quota resources in WebLogic Server 10.x?

A. Quota settings can be shared across multiple JMS destinations so that they can compete.

B. Quota settings are defined on the JMS connection factory level.

C. Quota settings can be used to limit the number of JMS Connections to a JMS Server.

D. Quota settings can be shared across multiple WebLogic domains.

Answer: A

13. The feature to group messages into a single unit that is processed sequentially in the order that the messages were created is called

A. Unit-or-Sort

B. Message Sorting

C. Destination Keys

D. Unit-of-Order

Answer: D

14. When a message is being processed by a consumer, the other unprocessed messages which belong to the same Unit-of-Order will be:

A. Expired

B. Blocked

C. Delivered to the consumers

D. Rolled back

Answer: B

15. Click on the Exhibit button.

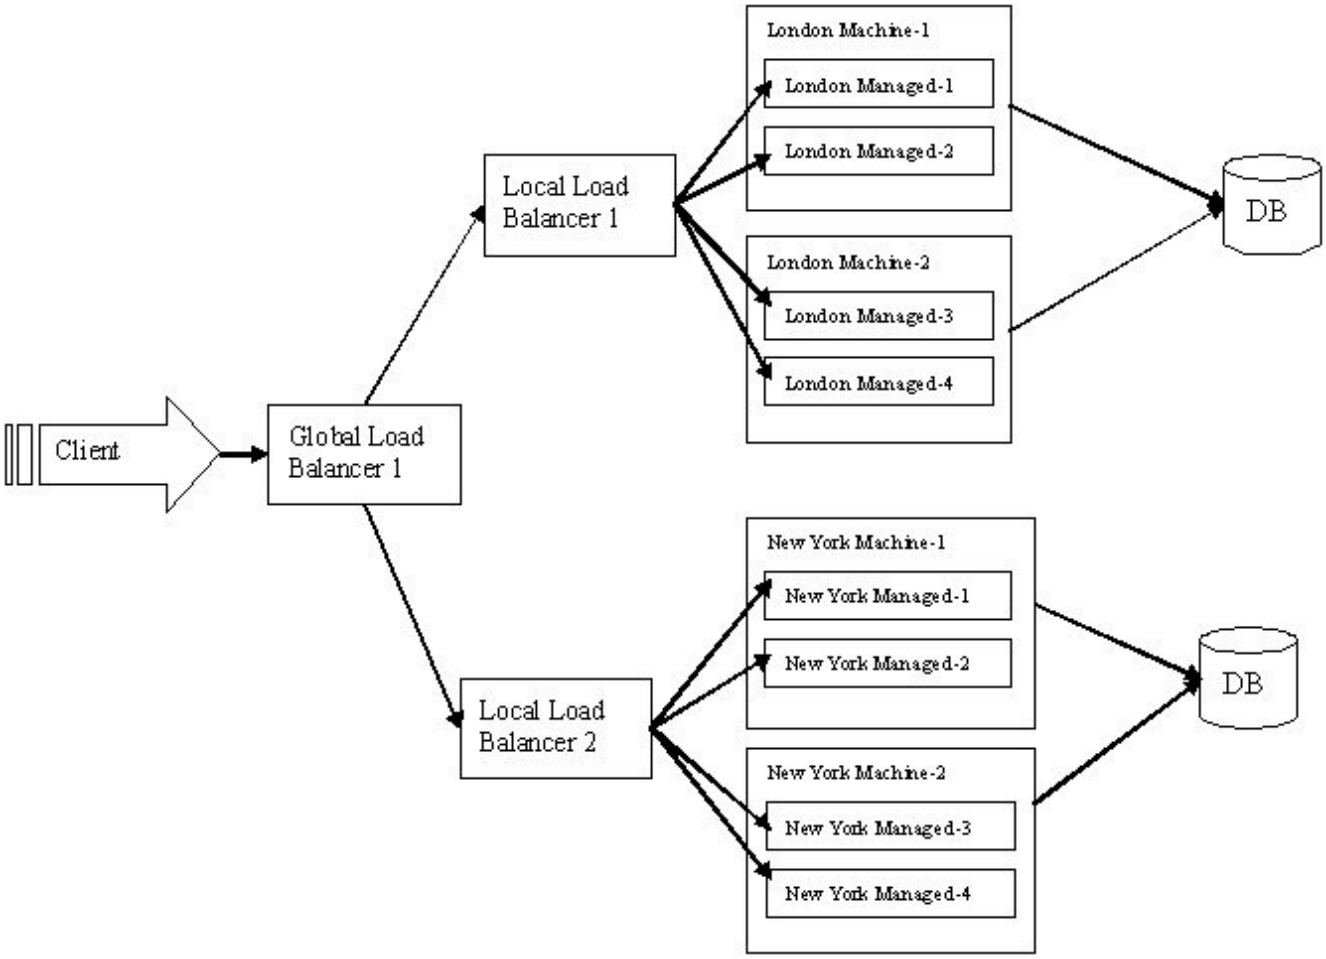

Scenario : A multi-cluster architecture, with two geographically separated clusters (WAN-1 and WAN-2), is

being created. It will provide cross-cluster WAN HTTP session failover. The WAN-1 cluster is located in London and the WAN-2 cluster is located in New York and is connected via a slow interconnect. The client request passes through a Global load balancer which load balances the request to local load balancers which eventually pass the load to WebLogic Server instances in the cluster. Four WebLogic Server instances will be running on each cluster distributed on two physical machines, two WebLogic instances per physical machine. A database is configured on each cluster for HTTP Session state failover between the clusters.

What would be the characteristics of the above WAN Cluster for optimal performance and scalability?

A. High network latency, asynchronous persistent replication between WAN-1 and WAN-2, synchronous intra-cluster replication and cross-cluster replication via in-memory, JDBC or file persistence.

B. Fast network, synchronous persistent replication between WAN-1 and WAN-2, asynchronous replication intra-cluster and cross-cluster replication via file persistence.

C. High network latency, asynchronous persistent replication between WAN-1 and WAN-2, synchronous intra-cluster replication and cross-cluster replication via JDBC

D. Low network latency, asynchronous persistent replication between WAN-1 and WAN-2, no intra-cluster replication and cross-cluster replication via JDBC

E. Fast network, synchronous persistent replication between WAN-1 and WAN-2, asynchronous intra-cluster replication and cross-cluster replication via cookie persistence.

Answer: C

16. Click on the Exhibit button.

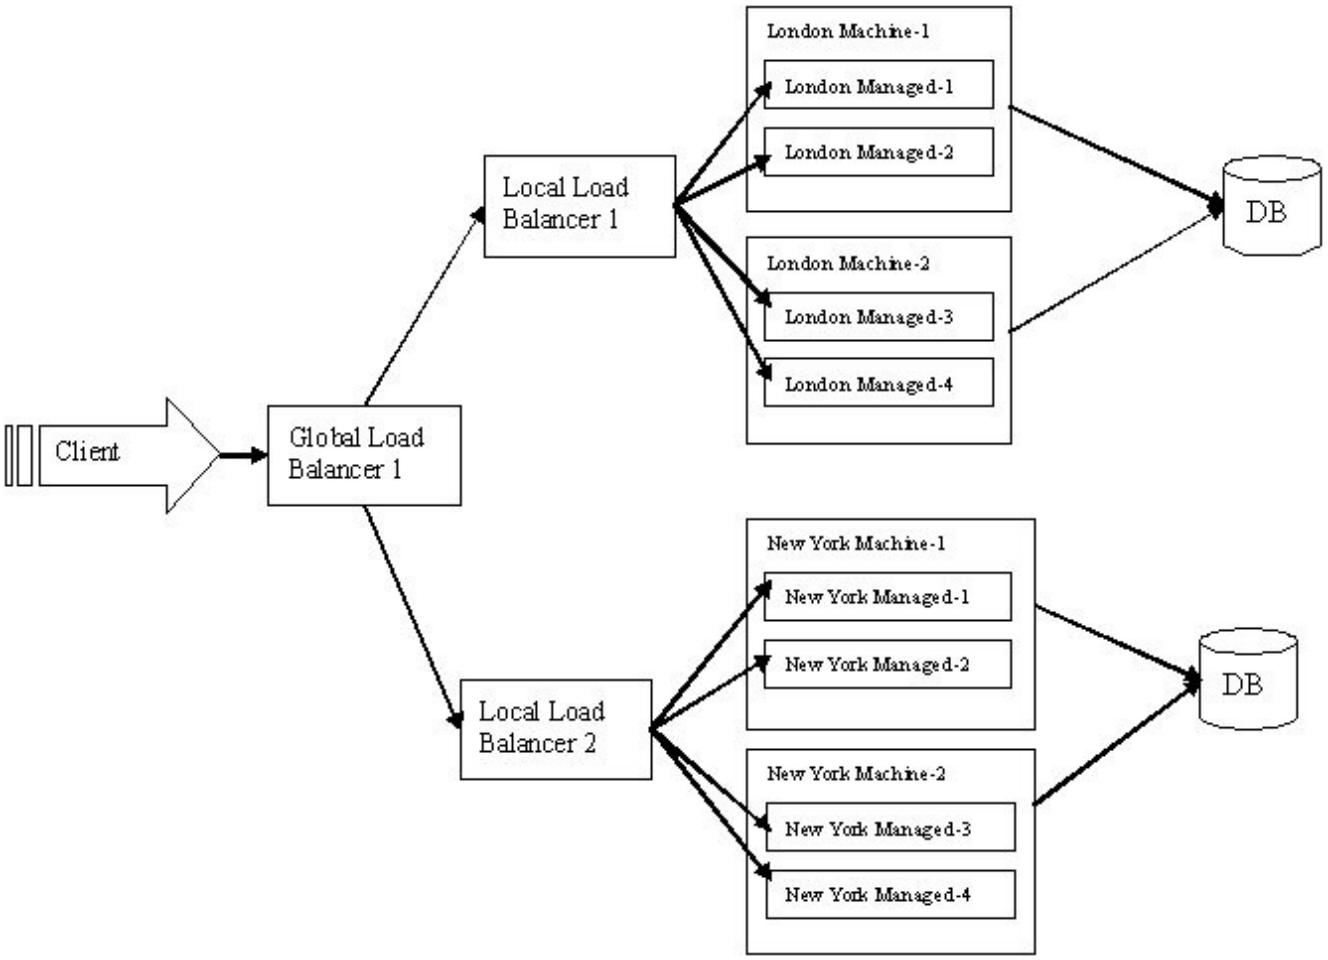

Scenario : A multi-cluster architecture, with two geographically separated clusters (WAN-1 and WAN-2), is being created. It will provide cross-cluster WAN HTTP session failover. The WAN-1 cluster is located in London and the WAN-2 cluster is located in New York and is connected via a slow interconnect. The client request passes through a Global load balancer which load balances the request to local load balancers which eventually pass the load to WebLogic Server instances in the cluster. Four WebLogic Server instances will be running on each cluster distributed on two physical machines, two WebLogic instances per physical machine. A database is configured on each cluster for HTTP Session state failover between the clusters.

What happens if all the servers hosted in the London cluster crashes in the middle of an http session update?

A. The session will be regenerated on the New York cluster without any session data loss, and a new primary server will be created using the secondary session data.

B. The session will be regenerated on the New York cluster but there could be a loss of session data; a new primary server will be created using the last known secondary session data.

C. The session will not be regenerated on any cluster; the user will be redirected to the New York cluster to

create a new session.

D. The session will be regenerated if the primary server on the London cluster is restarted.

E. Both b and d

Answer: B

17. Click on the Exhibit button.

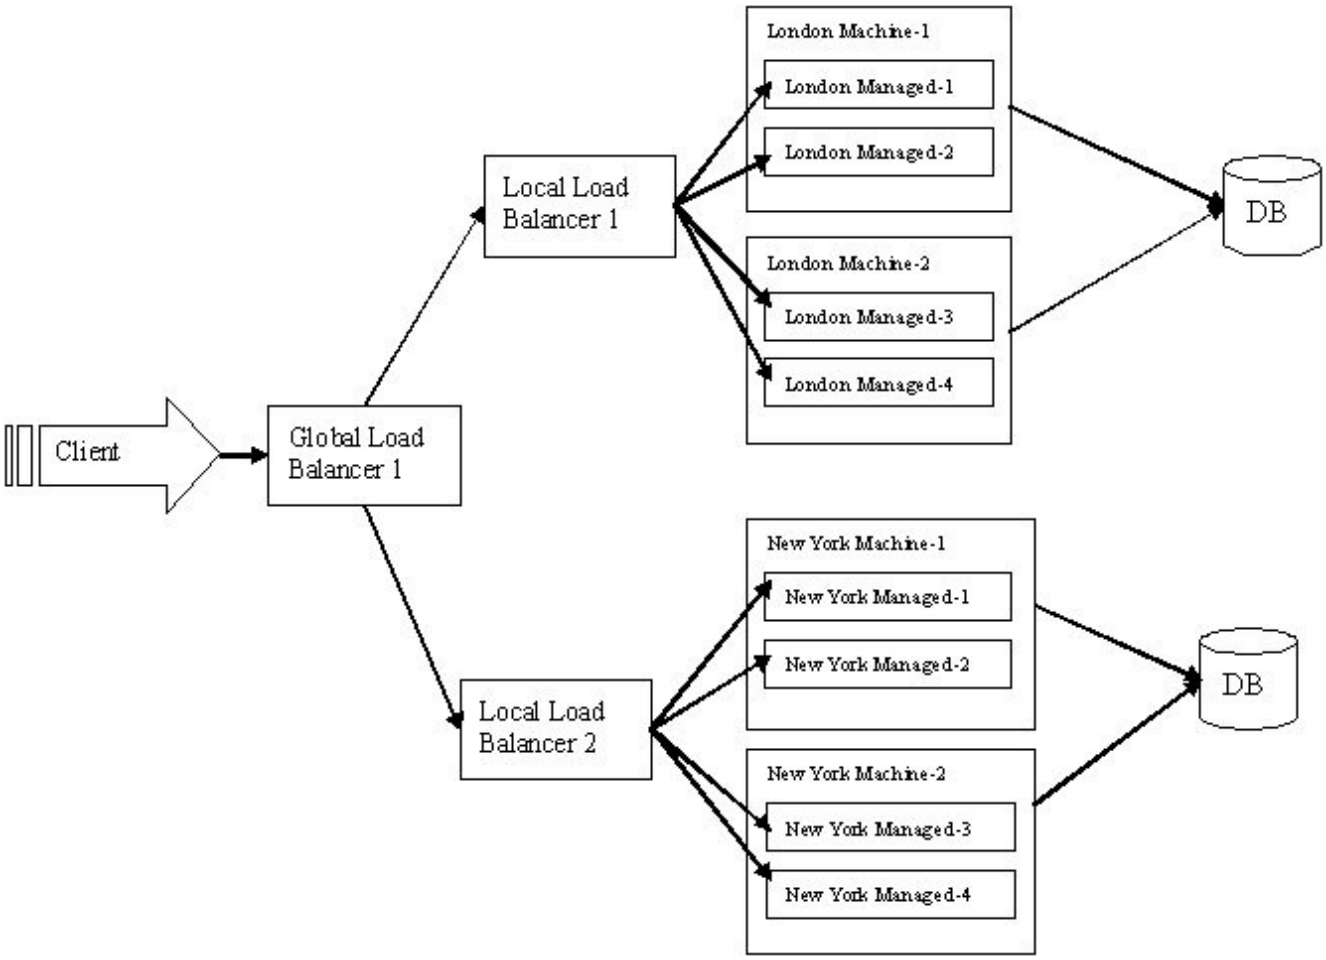

Scenario : A multi-cluster architecture, with two geographically separated clusters (WAN-1 and WAN-2), is being created. It will provide cross-cluster WAN HTTP session failover. The WAN-1 cluster is located in London and the WAN-2 cluster is located in New York and is connected via a slow interconnect. The client request passes through a Global load balancer which load balances the request to local load balancers which eventually pass the load to WebLogic Server instances in the cluster. Four WebLogic Server instances will be running on each cluster distributed on two physical machines, two WebLogic instances per physical machine. A database is configured on each cluster for HTTP Session state failover between the clusters.

If the admin server for New York cluster crashes, what will be the effect on inter-cluster communication

between the two sites?

A. There's no effect on the inter-cluster communication between the two sites

B. JDBC session flushing between two sites will be affected and will only be restarted on the Admin server reboot.

C. Intra-cluster multicasting will be affected and inter-cluster communication will slow down between the sites.

D. The Managed Servers will lose session stickiness and the server Health state will be changed to suspend.

E. Both c and d

Answer: A

18. Click on the Exhibit button.

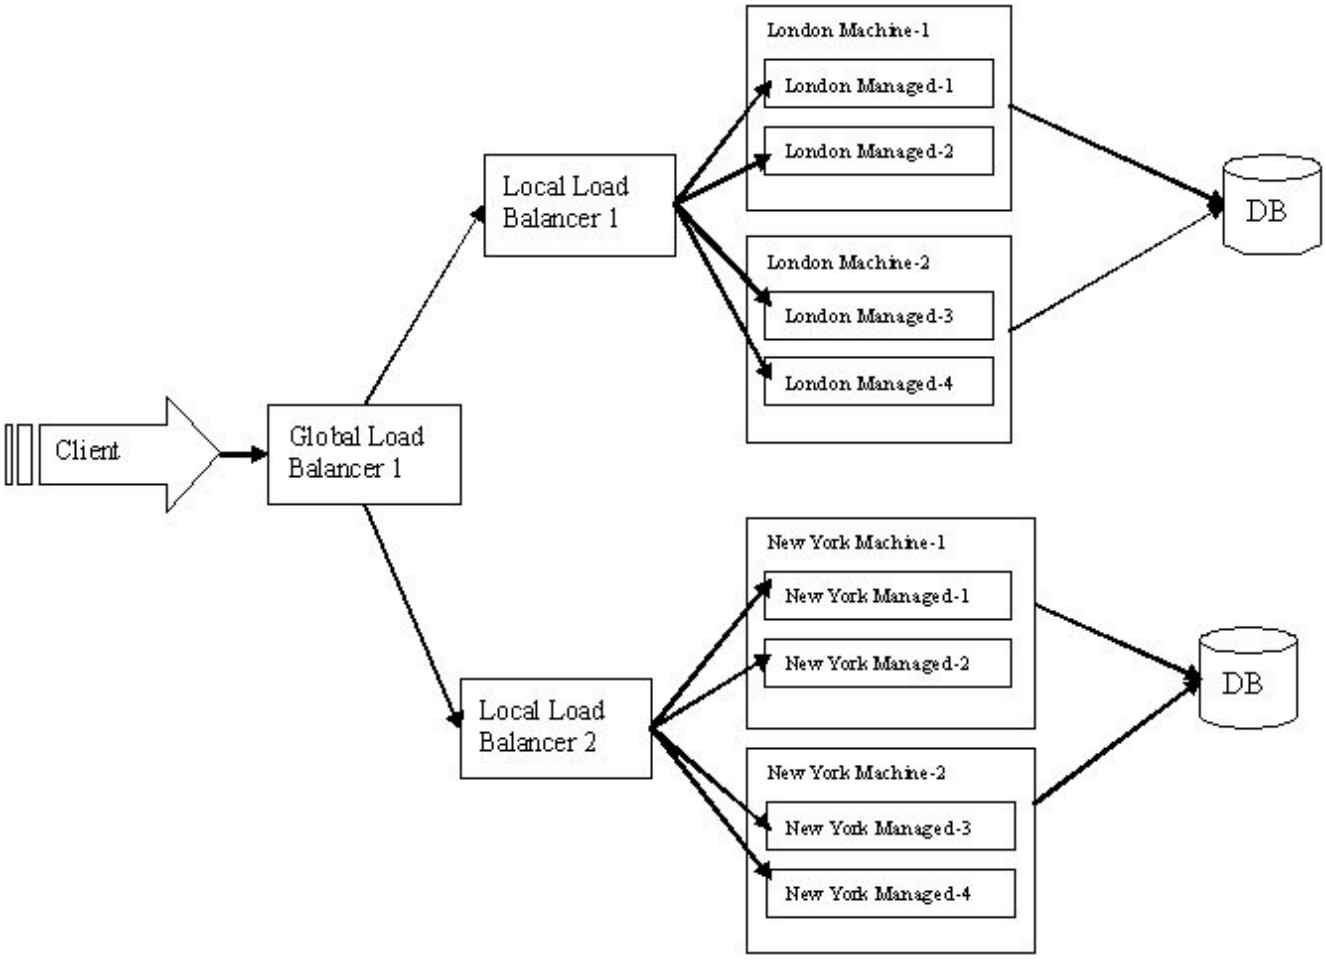

Scenario : A multi-network cluster, say WAN-1 and WAN-2, is being created that will provide WAN HTTP session failover. The WAN-1 cluster is located in London and the WAN-2 cluster is located in New York and is connected via a slow interconnect. The client request passes through a Global load balancer which load

balances the request to local load balancers which eventually pass the load to WebLogic Server instances in the cluster. Four WebLogic Server instances will be running on each cluster distributed on two physical machines, two WebLogic instances per physical machine. A database is configured on each cluster for HTTP Session state failover between the clusters.

How can you configure in-memory session replication on the London cluster?

A. The web application deployment descriptor, weblogic.xml needs to have the persistent-store-type set to memory.

B. The web application deployment descriptor, weblogic.xml needs to have the persistent-store-type set to replicated.

C. The web application deployment descriptor, weblogic.xml needs to have the persistent-store-type set to cookie.

D. The web application deployment descriptor, web.xml needs to have the persistent-store-type set to memory.

E. The web application deployment descriptor, web.xml needs to have the persistent-store-type set to replicated.

Answer: B

19. Scenario : A single tier WebLogic cluster is configured with six Managed Servers. An Enterprise application is deployed to the cluster with a web application and EJBs packaged in the same EAR file.

A JSP in the web application makes an initial Context lookup on the stateless EJB and makes multiple method calls to the EJB. Which of the following is true?

A. Each method call is invoked on an EJB on the same server as the web application as long as the server is not overloaded.

B. Each method call is invoked on the same EJB in the same server as the web application as long as the server is not overloaded.

C. Each method call round robins between the servers.

D. The behavior depends on the load algorithm on the replica-aware stub.

E. Every third call goes to a random server.

Answer: A

20. Requirement demands that an incoming request from a portfolio manager has to have a higher priority compared to other users of the application. As an administrator, which one of the following approaches would you adopt?

A. Use Work Manager with the execute thread model to set up rules to associate a priority to the user.

B. Use Work Manager fair-share-request-class to associate a priority to the user.

C. Use Work Manager response-time-request-class to associate a priority to the user.

D. Use Work Manager context-request-class to associate a priority to the user.

Answer: D