IBM C1000-174 Dumps

Shared by Fuller on 20-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An administrator receives the following error while upgrading WebSphere Application Server ND from 9.0.5.2 to 9.0.5.14:

CRIMA1020E ERROR: Error during ''install'' phase:

The administrator determines that the installation is corrupt. How can the corrupt installation be resolved with minimum effort?

Options:

A- Run backupConfig, uninstall and re-install WebSphere, and run restoreConfig.
B- Restore the file system to last working backup, apply the upgrade, and re-configure the environment.
C- Run backupConfig, restore the file system, apply the upgrade, and run restoreConfig.
D- Restore the VM snapshot, apply the upgrade, and re-configure WebSphere environment.

Answer:

C

Explanation:

The error CRIMA1020E during the install phase tells the administrator that Installation Manager could not complete the update, and in this case the installation files themselves are damaged. The aim is to return to a good installation, apply the fix pack, and keep the existing configuration, with as little work as possible. The most economical approach combines the product's backup and restore tools with a file-system restore. The administrator first runs backupConfig, which saves the configuration of the profile to an archive. The installation is then repaired by restoring the file system from the last working backup, which brings back a consistent set of product files. The upgrade to the target level is applied next, and finally restoreConfig reloads the saved configuration. Option C describes this sequence. Option A, uninstalling and re-installing the product, takes more effort because it needs a fresh installation, with the fixes reapplied to reach the required level, before the configuration can be restored. Option B restores the file system and then reconfigures the environment by hand, which is slow and error-prone, and the configuration in the file-system backup may be out of date. Option D, restoring a virtual machine snapshot and reconfiguring, has the same drawback, since the manual reconfiguration is the expensive part. Before any restore, the administrator should stop all the processes of the installation, and should verify the result with imcl listInstalledPackages and the versionInfo command.

Question 2

Question Type: MultipleChoice

Which statement is correct about Cross Component Trace (XCT)?

Options:

A- It is a tool to analyze log and trace files
B- It is used to capture data snapshots
C- It is best suited in a production environment
D- It is used to reduce the log and trace files with correlation information

Answer:

D

Explanation:

Cross Component Trace is a facility of the product's logging framework that records, in the log and trace output, the information needed to trace a request as it passes through different components, threads, and even separate servers. It does this by adding correlation data, such as a request identifier, to the records. Because every record belonging to the same request carries the same identifier, an analyst can reconstruct the path of the request through the log and trace files from different sources. The correlation information is thereby added to log and trace data; it does not reduce the volume. Option D, as worded, describes the use of correlation information with the log and trace files, which is the closest to the facility's purpose among the choices, although the verb 'reduce' applies more to the effort of analysis than to the size of the files. Option A is incorrect because XCT is not a tool for analysing logs; analysis is carried out with other tools, such as the log viewers, and with XCT the data is only prepared. Option B refers to capturing data snapshots, which is the effect of the highest correlation level only, and not the general purpose. Option C is incorrect because the tracing adds overhead, so it is best used for troubleshooting, in test or targeted production investigation, and not as a permanent production setting. The level can be set to request ID only, for the least overhead, up to full snapshots.

Question 3

Question Type: MultipleChoice

Which two methods can be used to generate migration reports for modernizing enterprise applications in WebSphere Application Server?

Choose two

Options:

A- Run the wsadmin AdminTask.createReport command.
B- Run the wsadmin AdminTask.createMigrationReport command.
C- Run the wsadmin AdminTask.createApplicationMigrationReport command.
D- Use the Analyze buttons on the administrative console enterprise application collection page.
E- Run the application code scanner and document potential bottlenecks.

Answer:

B, D

Explanation:

Migration reports can be generated for installed enterprise applications by two supported methods. In scripts, the wsadmin AdminTask.createMigrationReport command analyses a named application, or several, and stores the report with it, which is option B. In the administrative console, the enterprise application collection page, which lists the installed applications, includes an Analyze button. Selecting one or more applications and clicking it launches the same analysis, which is option D. The other options are not valid. AdminTask.createReport, option A, and AdminTask.createApplicationMigrationReport, option C, are not tasks in the product, since the correct name carries the words Migration and Report in that order without Application. Running a code scanner to document potential bottlenecks, option E, is a description of performance analysis, which is not what the migration reports do; they evaluate the use of APIs and technologies against the target platform and not the application's performance characteristics. After generation, the report can be reviewed in the console, where its severe, warning, and recommendation findings are listed, or retrieved with the getMigrationReport task. As the report is deleted when the application is uninstalled or redeployed, it should be generated again whenever the application is changed, so that the findings correspond to the code that is currently installed.

Question 4

Question Type: MultipleChoice

An administrator wants to use the addNode command to federate an application server installation into a cell. They use the --includeapps option to bring over the applications from the node. One application on that node is the default application, which is already installed on another server in the cell.

What is the outcome?

Options:

A- As the application already exists in the cell, the default application will not be installed.
B- The two nodes will each have their own version of the default application installed.
C- Applications that are shipped with WebSphere Application Server are ignored when running the addNode command.
D- As the binaries of the default application are identical, the application will be merged.

Answer:

A

Explanation:

The addNode command federates a stand-alone application server node into a cell, and its -includeapps option brings the node's installed applications into the cell's master repository. The default application, supplied with the product, is a sample application that is commonly present on every base server. When the cell already contains an application with the same name, installed through another node, the command does not install it again. The existing application is retained and the duplicate from the new node is skipped, with a message in the addNode log noting that it was not installed because an application of that name already exists. Option A is correct. Option B is wrong because two applications with the same name cannot coexist in a cell, since application names must be unique. Option C is wrong because the command does not ignore the product's shipped applications as a class of its own; the outcome results from the name conflict, and applications with unique names would be included. Option D is wrong because no merging of applications takes place, whatever the similarity of their binaries. An administrator who wants the applications of a node preserved should confirm names are unique before the federation. If the application that was skipped needs to run on the new node as well, it can be mapped to the node's server afterwards, through the console's module mapping, or by reinstalling it with that target added.

Question 5

Question Type: MultipleChoice

An administrator has created and configured a managed node for a non-IBM web server. Which two statements are true about this configuration?

Choose two

Options:

A- The administrator can propagate the web server plugin-cfg.xml using the administrator console.
B- The administrator can view web server logs using the administrator console.
C- The administrator can generate plug-in configurations using the administrator console.
D- The administrator can enable web server plug-in trace using the administrator console.
E- The administrator can start and stop the web server using the administrator console.

Answer:

A, C

Explanation:

When a non-IBM web server is represented in the cell as a managed node, a node agent for the web server is set up, which gives the deployment manager a route for managing certain aspects of it. Two capabilities hold for such a node. The administrator can generate the plug-in configuration with the console, and can propagate the resulting plugin-cfg.xml file to the web server through the node. Options A and C are correct. Managed web server nodes are the means by which the console keeps the plug-in configuration current as applications and clusters change. The other capabilities are not available for non-IBM web servers. Viewing the web server's logs through the console, enabling the plug-in trace, and starting or stopping the server rely on the IBM HTTP Server administration server, which is provided for IBM HTTP Server and is not available for other web servers. For a non-IBM server, these actions would have to be done directly on the web server host with its own tools. This explains why options B, D, and E are incorrect. The administrator should therefore treat the console as responsible for plug-in generation and propagation, and manage the lifecycle, logging, and tracing of a non-IBM web server with the vendor's own mechanisms. The plug-in's own log, whose location is set in the plug-in configuration, can be read directly on the host.

Question 6

Question Type: MultipleChoice

What is the proper sequence for setting up LTPA SSO between two WebSphere cells?

Options:

A- Export the LTPA key from CellA and then import the LTPA key to CellB
B- Export the LTPA key and then import the LTPA key to a text file.
C- Generate the LTPA key from CellA, export CellB, and then import the LTPA key to CellA.
D- Export the LTPA key from CellB and then export the LTPA key from CellA.

Answer:

A

Explanation:

Lightweight Third Party Authentication tokens are encrypted and signed with keys held in each cell's configuration. A token generated in one cell can only be accepted by another cell if both cells share the same LTPA keys, along with the same password and a compatible realm configuration. Single sign-on between two cells is therefore set up by taking the keys from one cell and placing them in the other. The sequence is to export the LTPA key from the first cell to a file, using the console's LTPA page or the exportLTPAKeys wsadmin task, and then to import that file into the second cell. Option A describes this: export from Cell A and import to Cell B. The other sequences are not workable. Option B proposes importing the key to a text file, which is not an import into a cell and so does not share the key. Option C generates a key in Cell A, exports from Cell B, and imports to Cell A, which mixes the cells' roles and ends with Cell B's key unchanged, so the cells do not end up holding the same key. Option D exports from both cells, and neither cell receives the other's key. After importing, the administrator saves the configuration, and generally restarts servers so that the new keys take effect. The user registries of the cells should resolve users to the same identities, and the domain name for the SSO cookie should be common to both.

Question 7

Question Type: MultipleChoice

What is the primary function of the Performance Advisors in WebSphere Application Server?

Options:

A- Design the website layout based on user experience.
B- Provides a starting point for the server configuration optimization.
C- Analyze the application code and error codes on runtime application server logs.
D- Manage the server resources and dynamic runtime configurations.

Answer:

B

Explanation:

Performance advisors in WebSphere Application Server examine runtime data from the Performance Monitoring Infrastructure and compare it against rules that represent good practice for configuring resources. When a counter shows that a setting is likely to be inappropriate, for instance a pool that is too small or a cache that is evicting heavily, the advisor issues a recommendation that the administrator can review. The advice is a starting point for optimising the configuration of the server, and it relies on the data of the load that the server is actually under, so it reflects real behaviour. Option B is correct. The Performance Advisor works within the Tivoli Performance Viewer, and the Runtime Performance Advisor runs within the server and writes messages to the log. The other statements are mistaken. The advisors do not design site layouts for user experience, option A, and they do not analyse application code or interpret error codes in server logs, option C. They also do not manage resources or change the configuration dynamically, option D; they recommend, and the administrator decides whether to apply the change. Because the advice is based on observed data, it is only as good as the load and the sampling period behind it. It should be gathered under representative peak load and applied one change at a time, with the result measured before the next adjustment.

Question 8

Question Type: MultipleChoice

Which enterprise module may be installed in WebSphere Application Server ND?

Options:

A- SAR
B- CAR
C- BAR
D- JAR

Answer:

D

Explanation:

An enterprise application is packaged as an enterprise archive, or EAR, which contains modules of several types. The module types the product recognises are web modules, packaged as WAR files, enterprise bean modules, packaged as JAR files, application client modules, also JAR files, and resource adapter modules, packaged as RAR files. Of the four extensions offered, only JAR corresponds to a module type that can be installed, since it is used for EJB modules and client modules. Option D is correct. The other extensions do not correspond to enterprise modules. SAR, option A, is a name used by some other application servers for service archives, and it is not an enterprise module type in this product. CAR, option C, is not an enterprise module type either; the term client archive does appear in some other contexts, but the product's application client module uses the JAR extension. BAR, option B, is likewise not a module type in this product. When an application is installed, the console or wsadmin reads the modules from the EAR and presents each one for mapping to servers and for binding of references. A standalone WAR or JAR may also be installed, in which case the product wraps it in an enterprise application, so the same module types apply. Understanding the module types helps in planning the mapping of modules to targets.

Question 9

Question Type: MultipleChoice

An administrator configured database session persistence for an application deployed to a WebSphere Application Server ND cluster. Application users have begun reporting data loss in HTTP session objects. How can the administrator resolve the issue with database session persistence?

Options:

A- Write session data at periodic intervals.
B- Select multi-row schema on session datasource: drop the existing session table, and recreate it.
C- Switch to memory to memory session persistence.
D- Create a session database with large BUFFERPOOL SIZE with PAGESIZE 32K.

Answer:

B

Explanation:

Database session persistence stores HTTP session data in a table of a relational database, so that sessions survive a server failure and can be shared among cluster members. The product offers two layouts. In the single-row schema, all of the session's data is held in one row, and the size of that row is limited by the database's row size, so large sessions can exceed the limit and be truncated or lost. In the multi-row schema, a session's data is spread across several rows, which removes the limit on session size and allows large sessions to be saved without loss. When users report data loss in session objects, the likely cause is large session data exceeding the single-row limit. The remedy is to select the multi-row schema on the session data source, drop the existing session table, and let the product recreate it with the new layout. Option B is correct. Option A, periodic writing, would raise the risk of loss between writes because data changed since the last write is lost on failure. Option C, switching to memory-to-memory replication, changes the persistence mechanism altogether and does not correct the cause. Option D, enlarging the buffer pool and page size, may help the database performance but does not address a limit on the session's data in a row. Dropping the table discards existing sessions, so the change should be scheduled for a quiet period, with users informed of the need to log in again.

Question 10

Question Type: MultipleChoice

Which wsadmin command generates migration reports for one or more enterprise applications in WebSphere Application Server ND?

Options:

A- AdminTask.createMigrationReport
B- AdminConfig.getMigrationReport
C- AdminConfig.getMigrationSummary
D- AdminTask.createMigrationRecord

Answer:

A

Explanation:

Migration reports for enterprise applications are produced through administrative tasks. In wsadmin the AdminTask object provides createMigrationReport, which analyses one or more installed applications against a target, identifies issues, and stores the report with the application. Option A is correct. The task accepts parameters that identify the application or applications and the target of the analysis, and the resulting report can then be viewed in the console or with getMigrationReport. The matching console action is the Analyze button on the enterprise applications page. The other commands do not exist or belong to a different object. AdminConfig.getMigrationReport, option B, combines a real task name with the wrong object, since AdminConfig deals with configuration objects and does not provide migration report functions. AdminConfig.getMigrationSummary, option C, is not a command in the product. AdminTask.createMigrationRecord, option D, uses a made-up command name, since the term used is 'report'. Because the report depends on the application files, administrators must regenerate it after any redeployment, which removes the earlier version automatically. When many applications are to be assessed, a script that loops through the installed applications and calls createMigrationReport for each is an effective way to produce the whole set, and the reports can be examined together to plan the order of migration.

Question 11

Question Type: MultipleChoice

An administrator has deployed several enterprise applications in WebSphere Application Server ND. How can the administrator configure the applications to start in a particular order?

Options:

A- WebSphere does not support ordered starting of applications.
B- Configure the Launch application before server completes startup parameter.
C- The applications are started in the order that they are defined.
D- Configure the Startup order parameter.

Answer:

D

Explanation:

Applications deployed to an application server are started during server startup, and the product provides a setting to control the order. On the application's configuration page in the console, and as an attribute in wsadmin, there is a Startup order parameter, which accepts an integer. Applications with lower values start before those with higher values, so an administrator can start

a shared service application before the applications that depend on it. Option D is correct. Option A is wrong because the ordering is supported. Option C is wrong because the applications are not started in the order in which they were defined or installed; without ordering values, the sequence is not something that can be relied on. Option B concerns a different parameter, 'Launch application before server completes startup', which determines whether the server waits for the application to start before it declares itself started and accepts requests. That is related to startup behaviour but does not set the order among applications. Administrators typically combine the two settings: the startup order sets the sequence, and the launch-before-completion option makes sure a critical application is fully up before the server opens for traffic. The values should be checked after changes, since the product starts applications in parallel by default when ordering is not specified, and the server log records the start messages for confirmation.

P2P exams

Question 12

Question Type: MultipleChoice

What is the supported runtime environment for a stand-alone thin client?

Options:

A- Apache Geronimc
B- Java EE client container
C- Java SE
D- JBoss

Answer:

C

P2P exams

Explanation:

A stand-alone thin client is a Java program that uses the thin client runtime libraries of WebSphere Application Server to invoke enterprise beans or other remote services. It is called stand-alone because it does not run in an application client container; it is launched like any other Java application, with its own main method. The platform that supports it is therefore Java SE, the standard Java runtime, with the thin client classes on the class path, which is option C. The client carries out its own initialisation of JNDI, security, and the ORB, without the services that a container provides. The Java EE client container, option B, is the runtime for a different client type, the application client, which is packaged as a module and run by the product's application client launcher. Apache Geronimo and JBoss, options A and D, are other vendors' Java EE application servers, and they are not runtime environments for the product's thin client. When

selecting the thin client, the developer should supply the thin client jar for the product and a JRE of a supported level, and should handle configuration, such as the provider URL and the security properties, in the code or in system properties. Where declarative resource references and container-managed security are wanted, the Java EE application client is the more suitable choice.

P2P exams

P2P exams

To Get Premium Files for C1000-174 Visit
https://www.p2pexams.com/products/c1000-174

For More Free Questions Visit
https://www.p2pexams.com/ibm/pdf/c1000-174

