

Riverbed

Exam 499-01

Riverbed Certified Solutions Professional – Application Performance Management

Version: 6.0

[Total Questions: 196]

Topic break down

Topic	No. of Questions
Topic 1: AppInternals Xpert	19
Topic 2: BrowserMetrix	6
Topic 3: AppResponse Xpert	79
Topic 4: AppTransaction Xpert	48
Topic 5: RPM Dashboards	21
Topic 6: General Knowledge	23

Topic 1, ApplInternals Xpert**Question No : 1 - (Topic 1)**

What is the maximum number of managed nodes indicated per symmetric multiprocessor (SMP)?

- A. 10
- B. 20
- C. 50
- D. 100
- E. 200

Answer: C

Question No : 2 - (Topic 1)

The ApplInternals Xpert Transaction Trace Warehouse can extract usernames from transactions:

- A. Through parsers that extract information from packets
- B. Through an integration with Lightweight Directory Access Protocol (LDAP)
- C. From an application's config file
- D. From the method getRemoteUser() of the HttpServletRequest object

Answer: D

Question No : 3 - (Topic 1)

What useful information can you get from AppResponse Xpert (ARX) database browsing insights? (Select 3)

- A. Response time for SQL statements
- B. Database usage by application and user
- C. It tracks throughput in terms of database calls, transactions, rows, bytes, and active sessions.
- D. A link to "View Database SQL queries" insights

Answer: A,B,C

Question No : 4 - (Topic 1)

Which of the following best describes ApplInternals Xpert heterogeneous transaction tracing?

- A. Tiers with identical code are mapped and correlated for each transaction at run time.
- B. Tiers with identical code are deduplicated and correlated for each transaction at run time.
- C. Each trace includes all the tiers necessary to execute the transaction at run time.
- D. Tiers that are identical in the same transaction are excluded at run time.

Answer: C

Question No : 5 - (Topic 1)

The Transaction Trace Warehouse collects data from:

- A. Symmetric multiprocessor (SMP) managed nodes on the network
- B. Virtual machines specified in the Trace Transaction Warehouse interface
- C. Managed nodes specified by SMP
- D. Any server that is network accessible

Answer: C

Question No : 6 - (Topic 1)

Which tab in the ApplInternals Xpert interface do you use to enable continuous tracing for JIDA and .Net data adapter instances?

- A. Metric Filters
- B. Schedule
- C. Users
- D. Configure Groups

Answer: D

Question No : 7 - (Topic 1)

Which metrics help the user perform heap analysis? (Select 3)

- A. Percent of free heap
- B. Percent of time garbage collecting
- C. Percent of usage
- D. Average number of live objects

Answer: A,B,D

Question No : 8 - (Topic 1)

What is the maximum number of transactions stored by Transaction Trace Warehouse (TTW)?

- A. Infinite
- B. 1 billion
- C. Last 3 years' worth
- D. Infinite, but dependent on disk size
- E. Last years' worth
- F. Previous weeks' worth

Answer: D

Question No : 9 - (Topic 1)

How should you instrument a Java application when a command line to start the application is not readily accessible?

- A. This is not feasible currently.
- B. Recompile or modify the code.
- C. Set the JAVA_TOOL_OPTIONS environment variable to add the argument into the application start up.
- D. We should not instrument applications that use embedded virtual machines.

Answer: C

Question No : 10 - (Topic 1)

In ApplInternals Xpert consoles, you can click on a transaction trace or metric to open the line of source code that defines the generic class.method that generated that transaction or metric when applicable. Which statements are true regarding linking between source code

and AppInternals Xpert data? (Select 3)

- A. You can toggle back and forth between the source code and the AppInternals Xpert performance data to examine selected class/method transactions that exceed a particular duration.
- B. You can open the source code directly into Transaction Trace Warehouse.
- C. In Visual Studio or Eclipse source code, you can click on the name of a generic method or click inside the code for the method and link to traces on transactions that the .NET or Java code generated, and display those transactions in Transaction Trace Warehouse.
- D. You can click on the name of a generic method or click inside the code for the method and link to AppInternals Xpert metrics that the method generated, displaying in the desktop console.
- E. You can open the source code directly into Transaction Trace Analysis.

Answer: A,C,D

Question No : 11 - (Topic 1)

This change will not require a restart in the application when using AppInternals Xpert:

- A. Changes in jida.config
- B. Enabling monitoring of generic classes with config.package metrics
- C. Changes in the -agentpath parameter
- D. Enabling cross-tier transactions

Answer: B

Question No : 12 - (Topic 1)

By default, the analysis engine uses a "moving-window" algorithm to determine dynamic limits. You can change the algorithm by changing the DynamicLimitsAlgorithm attribute to any of these values, except:

- A. MovingWindowPercentile
- B. MovingWindowSigma
- C. HourlyLimits
- D. HistoricalBaseline
- E. HardLimits

Answer: E

Question No : 13 - (Topic 1)

What is the most recommended network time protocol (NTP) configuration for ApplInternal Xpert?

- A. Managed nodes use the embedded client and synchronize with a common server.
- B. Managed nodes use the embedded NTP client and synchronize with symmetric multiprocessing (SMP).
- C. Managed nodes use their own NTP client.
- D. No specific NTP configuration is needed, as ApplInternal Xpert could handle the correct timing by itself.

Answer: B

Question No : 14 - (Topic 1)

How do you control the overhead of ApplInternals Xpert agents?

- A. Install agents on a limited number of machines.
- B. Restart the server.
- C. Agents self tune so there is no need to control the overhead.
- D. Never install more than 5 agents.
- E. Use multiple consoles.
- F. Unplug the server.

Answer: C

Question No : 15 - (Topic 1)

Which statement is incorrect about how dynamic sampling agents (DSAs) manage 1-second data storage?

- A. By default, the DSA attempts to keep 1-second data for 2 hours. The length of this time period is specified by the RawDataPersistDuration attribute in the dsa.xml file.
- B. The DSA deletes the oldest 1-second data from the database if there is less than 10 percent of free disk space (the default for the RawDataPersistFreeSpaceTargetPercent attribute)
- C. The DSA deletes the oldest 1-second data from the database if the database file is larger than specified in the RawDataPersistMaxDiskTarget attribute (unlimited as default).
- D. If both the RawDataPersistFreeSpaceTargetPercent and RawDataPersistMaxDiskTarget attributes are satisfied, the DSA deletes data, if necessary,

to retain only as much data as specified by RawDataPersistDuration.

E. At startup, and every 2 hours thereafter (or the frequency specified by RawDataPersistPurgeInterval), the DSA deletes the oldest 1-second data from the database.

Answer: A

Question No : 16 - (Topic 1)

What granularity of data is available on the Service Management Platform (SMP)? (Select 2)

- A. 1 minute
- B. 5 minutes
- C. 1 second
- D. 15 minutes

Answer: C,D

Question No : 17 - (Topic 1)

JIDA automatically collects metrics from predefined categories of classes except:

- A. EJBs: classes that implement the javax.ejb.EntityBean or javax.ejb.SessionBean interfaces
- B. JDBC - classes that implement any of several interfaces in the java.sql package
- C. JNDI - classes that implement the javax.naming.Context or javax/naming/directory/DirContext interfaces
- D. Portlets - classes that implement the javax.portlet.Portlet interface
- E. Generic classes

Answer: E

Question No : 18 - (Topic 1)

What issues arise in instrumenting a Java application using IBM SDK Java when the command line to start the application is not readily accessible? (Select 2)

- A. Some code compilation or modification is needed.

- B. You can set the IBM_JAVA_OPTIONS environment variable to add the argument into the application startup.
- C. IBM SDK Java is not supported.
- D. Setting the IBM_JAVA_OPTIONS environment variable will add the same argument into every IBM SDK Java application startup.

Answer: B,D

Question No : 19 - (Topic 1)

With AppInternals Xpert, how are SQL queries captured?

- A. An agent installed on the database
- B. AppResponse Xpert Appliance capturing traffic
- C. Agent installed on application server
- D. Agent installed on web server

Answer: C

Topic 2, BrowserMetrix

Question No : 20 - (Topic 2)

Which of these are not available to instrument a web application with BrowserMetrix?

- A. Use Riverbed Steelhead to inject the script at run time from the web server.
- B. Use Riverbed Stingray to inject the script at run time from the web server.
- C. Manually paste the Javascript in any web page you want to monitor with BrowserMetrix.
- D. Edit a master file for the application by adding the Javascript to the header section of the file.

Answer: A

Question No : 21 - (Topic 2)

Safari seems to outperform other browsers in its response time report. Why is its "time to first byte" so much shorter than other browsers? (Select 2)

- A. Safari does not support the W3C Navigation Timing specification.

- B. Safari is a more responsive browser.
- C. BrowserMetrix does not process "time to first byte" currently.
- D. This is the user's first visit to the site or the previous page was not one of the instrumented pages.

Answer: A,D

Question No : 22 - (Topic 2)

Each of the "dots" you see in the BrowserMetrix user interface represent:

- A. An end-user page view
- B. The web server tier
- C. The web application
- D. The server node for the transaction

Answer: A

Question No : 23 - (Topic 2)

Which of these can be used to inject BMX JavaScript? (Select 3)

- A. AppResponse Xpert
- B. Browser plugin
- C. A load balancer such as Stingray
- D. SysAdmin or webmaster
- E. Proxy server

Answer: B,C,E

Question No : 24 - (Topic 2)

For BrowserMetrix, the "time to first byte" and "download and render" time is considered:

- A. Back-end delay
- B. Front-end delay
- C. Network delay
- D. Server delay