

EXAM PREPARATION MATERIAL

Practice files designed to help you understand exam patterns, question styles and the overall assessment format.

IMPORTANT NOTE



Please note that these files may not be up to date.
However, the questions will help you understand the exam format and typical question patterns.



Stay connected with our support team
for the latest file updates and exam-related information.

Question: 1

Which two methods are used for confirming the agent is communicating properly to the controller? (Choose two.)

- A. Check if data is being written to the agent log in /logs
- B. Log in to the Controller UI > Settings cog icon > AppDynamics Agents In the list look for the agent in the list by machine hostname.
- C. Select Application > Tiers and Nodes > Look for the Node.
- D. Use Grep for the agent process on the application server using the specific agent version name

Answer: AC

Explanation:

To confirm that the agent is communicating properly to the controller, you can use the following methods:

Check if data is being written to the agent log in /logs. This indicates that the agent is able to send metrics and events to the controller. [You can also look for any errors or warnings in the log file that might indicate a connection issue1](#)

Select Application > Tiers and Nodes > Look for the Node. This shows you the list of nodes that are registered with the controller and their status. [You can see if the node is up or down, the agent version, the last reported time, and the machine name2](#)

[Reference: 1: Troubleshoot Agent Connectivity 2](#)

Question: 2

The application agent starts but cannot connect to the controller. Which two things should you check to solve this problem? (Choose two.)

- A. Check if the controller is using SSL or not.
- B. Check if the agent is using the correct application name.
- C. Check if the agent has enough memory to run.

D. Check that there are licenses available on the controller

Answer: AD

Explanation:

[According to the Cisco AppDynamics Associate Administrator Certification document1](#), one of the topics covered in the exam is "Monitor agent performance and connectivity". [To do this, you need to ensure that the application agent can communicate with the controller, which may require SSL configuration if the controller is using SSL2. You also need to check that there are enough licenses available on the controller for the application agent to register and report data3.](#) If the agent is using the wrong application name or does not have enough memory, it may not start at all or may crash, which are different problems from not being able to connect to the controller. Reference:

[Cisco AppDynamics Associate Administrator Certification](#)

[Configure SSL for the Controller and the App Server Agent](#)

[License Rules and Restrictions](#)

Question: 3

Which two stats are available in Database Visibility? (Choose two.)

- A. Time Spent in Database
- B. Average Number of Slow Connections
- C. Top Leaked Connections
- D. Top 10 Query Wait States
- E. Top 10 Connections by Application

Answer: BC

Explanation:

Database Visibility provides metrics on the performance of your database and helps troubleshoot

performance-related issues. Database Visibility consists of four main components: Database Agent, Collector, Controller, and Events Service (on-premises only). The Database Agent is a standalone Java program that collects performance metrics about your database instances and database servers. The Collector is the process that runs within the Database Agent to collect performance metrics about your database instances and database servers. The Controller is the central interface where you can see all your database instances and database server performance metrics. The Events Service stores high volumes of metric data.

According to the [Overview of Database Visibility](#), the following types of information are sent to the Controller:

Database-level metrics, such as the number of queries processed and other database statistics

Names and attributes of all sessions, clients, queries, and other objects on the monitored system

The following types of information are sent to the Events Service:

Time that each query spends at each wait state

Individual query statistics for databases that support it

Information about individual execution plans in databases that support it

Therefore, B (Average Number of Slow Connections) and C (Top Leaked Connections) are two stats that are available in Database Visibility. Reference:

[Cisco AppDynamics Associate Administrator Certification](#)

[Overview of Database Visibility - AppDynamics](#)

[Database Visibility - AppDynamics](#)

Question: 4

A customer wants to monitor 50th and 99th percentile response time in the Metric Browser for a particular Business Transaction Where do you configure and enable percentile metrics?

- A. Configuration > Baselines > Configure Percentile Metrics
- B. Configuration > Instrumentation > Configure Percentile Metrics
- C. Configuration > Slow Transaction Thresholds > Configure Percentile Metrics
- D. Configuration > Instrumentation > Transaction Detection > Configure Percentile Metrics

Answer: C

Explanation:

Percentile metrics are a way to measure the distribution of response times for a Business Transaction. They show the value below which a certain percentage of measurements fall. For example, the 50th percentile metric shows the median response time, and the 99th percentile metric shows the worst-case response time for 99% of the transactions. To enable percentile metrics, you need to go to Configuration > Slow Transaction Thresholds and indicate 5 whole numbers between 1 and 99 as Percentiles to Collect. You can then view the percentile metrics in the Metric Browser under Business Transaction Performance > Business Transactions > <Application Name> > <Business Transaction Name> > Percentiles.

<https://docs.appdynamics.com/display/PRO45X/Percentile+Metrics>

Question: 5

What gives administrators the ability to test and debug policy execution?

- A. Event Simulation Tool
- B. Policy Test Bench
- C. Action and Policy Execution Tool

D. Alert Simulation Environment

Answer: B

Explanation:

The Policy Test Bench gives administrators the ability to test and debug policy execution. It allows you to simulate events and see how the policies and actions respond to them. [You can also view the policy execution logs and troubleshoot any issues](#)¹

Reference: ¹: Policy Test Bench

Question: 6

Which three databases are supported by Database Visibility? (Choose three.)

- A. PostgreSQL
- B. Cassandra
- C. DB2
- D. MongoDB
- E. Derby

Answer: ABD

Explanation:

[According to the Database Visibility Supported Environments document](#)¹, Database Visibility can monitor the following databases: PostgreSQL, Cassandra, MongoDB, MySQL, SQL Server, Oracle, Sybase ASE, Sybase IQ, and Couchbase. Therefore, the correct answer is A, B, and D. Derby and DB2 are not supported by Database Visibility. Reference:

[Database Visibility Supported Environments](#)

Question: 7

What are three advantages of the custom dashboard feature? (Choose three.)

- A. It makes drill down across tiers seamless
- B. It turns data sharing on/off on the fly.
- C. It monitors metrics of interest.
- D. It finds the line of code having a performance issue
- E. It schedules dashboard as a report

Answer: ACE

Explanation:

The custom dashboard feature in AppDynamics allows you to display a specific set of metrics and data points on one screen. You can use custom dashboards to present selected metrics for a user who only needs a relatively narrow or focused view of the data. You can also share custom dashboards with other users and stakeholders.

According to the [Custom Dashboards - AppDynamics], some of the advantages of the custom dashboard feature are:

It makes drill down across tiers seamless: You can use the grid layout type to create a flexible layout that is easy to rearrange on the canvas. The grid layout also scales in size when viewed on mobile devices. You can also use the absolute layout type to control width and height and the exact placement of widgets on the canvas.

It monitors metrics of interest: You can add widgets from various categories, such as application, server, database, performance, events, and so on. You can also customize the widget properties panel to configure the widget settings, such as title, description, color scheme, time range, autorefresh interval, and so on.

It schedules dashboard as a report: You can export your custom dashboard as a PDF or HTML report that you can send by email or save as a file. You can also schedule your custom dashboard to run periodically using cron expressions.

Therefore, A (It makes drill down across tiers seamless), C (It monitors metrics of interest), and E (It schedules dashboard as a report) are three advantages of the custom dashboard feature. Reference:

[Custom Dashboards - AppDynamics]

[Create and Manage Custom Dashboards and Templates - AppDynamics]

[Custom Dashboard Permissions - docs.appdynamics.com](https://docs.appdynamics.com/custom-dashboard-permissions)

Question: 8

An HTTP Data Collector can capture which three kinds of data? (Choose three.)

- A. Parameter Values
- B. Headers
- C. Cookies
- D. Method Invocation
- E. Variables

Answer: ABC

Explanation:

An HTTP Data Collector is a type of data collector that captures data from HTTP requests and responses. It can capture three kinds of data: parameter values, headers, and cookies. Parameter values are the values of the query string or form parameters in the HTTP request. Headers are the key-value pairs that are sent or received as part of the HTTP request or response. Cookies are the small pieces of data that are stored by the browser and sent to the server with each request. An HTTP Data Collector can be configured to capture any of these data types by specifying the name of the parameter, header, or cookie in the data collector settings. The captured data can be used for various purposes, such as adding context to business transactions, creating custom metrics, or triggering health rules.

<https://community.appdynamics.com/t5/NET-Agent-Installation/HTTP-Data-Collector-How-to-know-the-HTTP-Params-name/td-p/24186>

Question: 9

Which three Operating Systems can be enabled in the Database Collector's Hardware Monitoring configuration? (Choose three.)

- A. OSX
- B. Linux

C. Windows

D. Solaris

E. HPUX

Answer: BCD

Explanation:

The Database Collector's Hardware Monitoring configuration allows you to enable the Database Agent to monitor the server hardware in addition to the database. [You can choose from the following operating systems: Linux, Windows, and Solaris](#)¹ You need to provide the credentials and the connection details for the server that hosts the database¹ Reference: 1: [Configure the Database Agent to Monitor Server Hardware](#)

Question: 10

When creating a scheduled report which field needs to be changed so the desired information is available in the report?

A. Recipients

B. Report Title

C. Report Type

D. Schedule

Answer: C

Explanation:

[According to the Reports document](#)¹, the Report Type field determines what kind of information is captured in the report. There are different report types available, such as Application Health Report, Dashboard Report, Controller Audit Report, and so on. Each report type has different fields in the Report Data tab that can be customized. Therefore, to get the desired information in the report, you need to select the appropriate Report Type from the dropdown menu when creating a scheduled report. Reference:

[Reports](#)

Question: 11

Which two options can be excluded using error configuration? (Choose two.)

- A. Database error return codes
- B. Uncaught exceptions
- C. JavaScript errors
- D. HTTP errors

Answer: AC

Explanation:

Error configuration is a feature in AppDynamics that allows you to exclude certain types of errors and exceptions from being reported on the dashboard. You can use error configuration to filter out noise and focus on the most relevant and actionable issues.

According to the [Error Configuration - AppDynamics](#), the following types of errors and exceptions can be excluded using error configuration:

Database error return codes: These are codes that indicate a problem with the database server, such as 0x80004005 (access denied) or 0x8000005E (access violation). You can exclude these errors from being reported on the dashboard by adding them to the error configuration list.

JavaScript errors: These are errors that occur in the browser due to invalid or malformed JavaScript code, such as syntax errors or reference errors. You can exclude these errors from being reported on the dashboard by adding them to the error configuration list.

Therefore, A (Database error return codes) and C (JavaScript errors) are two options that can be excluded using error configuration. Reference:

[Error Configuration - AppDynamics](#)

[Cisco AppDynamics Associate Administrator Certification](#)

[Create and Manage Error and Exception Configurations - AppDynamics]

Question: 12

Which two locations does an AppDynamics administrator use to view Remote Services metrics? (Choose two.)

- A. Tier Dashboard
- B. Business Transaction Dashboard
- C. Exit Point Metrics
- D. Tiers and Nodes Dashboard

Answer: AC

Explanation:

Remote Services metrics are the metrics that show the performance and behavior of the external systems that an application interacts with, such as web services, databases, message queues, etc. An AppDynamics administrator can use two locations to view Remote Services metrics: the Tier Dashboard and the Exit Point Metrics. The Tier Dashboard shows the flow map of a tier and its downstream dependencies, including the remote services that the tier calls. The administrator can click on any remote service node on the flow map to see the metrics such as average response time, calls per minute, errors per minute, etc. The Exit Point Metrics show the metrics for all the exit points (remote service calls) of an application, grouped by type, such as HTTP, JDBC, JMS, etc. The administrator can drill down to see the metrics for each exit point, such as the backend name, the tier name, the business transaction name, etc.

Question: 13

You need to examine the Java App agent logs on a host, but you do not have login access to the relevant host. How do you accomplish this via the Controller User Interface (UI)?

- A. Controller UI > Node Dashboard > Agents tab > App Server Agent tab > Agent Operations > Request Agent Logs
- B. Controller UI > application > Transaction Snapshots > Periodic Collection, then wait for the log to download
- C. Controller UI > Configuration > Instrumentation > Data Collectors > and then add a new collector for

class *.*

D. Controller UI > application > Alert Respond > Create Action then issue an HTTP request to request agent logs

Answer: A

Explanation:

To examine the Java App agent logs on a host without login access, you can use the Controller UI to request the agent logs. This feature allows you to download the agent logs from the Controller UI without having to access the host machine. [You can specify the log level, the time range, and the file size limit for the logs](#)

[Reference: 1: Request Agent Logs](#) [2: Java Agent Logging](#)

Question: 14

What is the result of starting a Diagnostic Session?

- A. A single snapshot is captured for the next Business Transaction that occurs.
- B. Snapshots are captured at an accelerated rate for the Business Transaction specified
- C. A snapshot is captured for every transaction flowing through the application until the Diagnostic Session is ended
- D. Snapshots are captured at an accelerated rate for all Business Transactions configured

Answer: B

Explanation:

[According to the Diagnostic Sessions document](#)¹, a diagnostic session is a timed session in which AppDynamics captures transaction snapshots for a business transaction at a more frequent interval that you set. In the session, AppDynamics captures full call graphs for the transactions. A diagnostic session can be triggered manually or automatically based on a health rule violation or a threshold of slow or error transactions. A diagnostic session is always associated with a specific business transaction, not all business transactions configured. Therefore, the correct answer is B. Reference:

Question: 15

You have a custom dashboard that has only one widget. Which statement about deleting that widget is true?

- A. When the last widget is deleted, the dashboard is also deleted.
- B. The widget is deleted, but an empty custom dashboard is maintained.
- C. When the widget is deleted, AppDynamics creates a blank widget because empty dashboards are not allowed.
- D. The delete widget option is not enabled when there is only one widget left.

Answer: B

Explanation:

When you have a custom dashboard that has only one widget, you can delete that widget by clicking the Delete Widget button on the widget properties panel. This will remove the widget from the

dashboard, but it will not delete the dashboard itself. You can still view and edit the dashboard from the Controller UI or the Events Service.

According to the [Create and Manage Custom Dashboards and Templates - AppDynamics](#), when you delete a widget from a custom dashboard, AppDynamics does not create a blank widget because empty dashboards are not allowed. Instead, it preserves the existing widgets on the dashboard and displays them in a single row.

Therefore, B (The widget is deleted, but an empty custom dashboard is maintained) is true. Reference:

[Create and Manage Custom Dashboards and Templates - AppDynamics](#)

[How to bulk delete custom dashboards using API - Cisco ... - AppDynamics](#)

[Custom Dashboard Permissions - docs.appdynamics.com](#)

Question: 16

What is the default auto-refresh interval for custom dashboards?

- A. 60 sees
- B. 120 sees
- C. 180 sees
- D. 240 sees

Answer: B

Explanation:

The default auto-refresh interval for custom dashboards is 120 seconds. This means that the data displayed on the dashboard will be updated every two minutes. The auto-refresh interval can be configured by the dashboard owner or administrator from the dashboard settings. The auto-refresh interval can be set to any value between 15 and 300 seconds, or disabled altogether. The autorefresh interval applies to all the widgets on the dashboard, unless a widget has its own refresh interval specified.

<https://docs.appdynamics.com/appd/23.x/latest/en/appdynamics-essentials/dashboards-and-reports/custom-dashboards/create-and-manage-custom-dashboards-and-templates>

Question: 17

The App Server agent configuration GUI sets the configuration for which two agent properties? (Choose two.)

- A. Auto-naming
- B. Capture-raw-sql
- C. Node-name
- D. Max-business-transactions per node

E. use-encrypted-credentials

Answer: CD

Explanation:

The App Server agent configuration GUI allows you to set the configuration for the following agent properties:

Node-name: This property defines the name of the node that appears in the Controller UI. [You can use the GUI to change the node name or use auto-naming rules based on environment variables or system properties1](#)

Max-business-transactions per node: This property sets the maximum number of business transactions that can be registered on a node. If the number of business transactions exceeds this limit, the agent stops registering new business transactions and reports them as "All Other Traffic". [You can use the GUI to increase or decrease this limit2](#)

[Reference: 1: Node Naming 2: Configure Business Transaction Limits](#)

Question: 18

A customer complains that their dashboard is broken They tell you that every time they change the time range using the time picker, the data on several of their dashboard widgets do not change to match the new time range What is happening in this situation?

- A. the dashboard is too large and needs to be smaller
- B. the widgets are the wrong type for the dashboard
- C. the widgets are set to use their own time range
- D. the data in those widgets is old and needs to be updated

Answer: C

Explanation:

[According to the Time Range Comparisons document1](#), you can use the Compare Time Range feature in Dash

Studio to compare two time periods on a widget. However, this feature also means that the widget may not use the global time range set for the entire dashboard by default. Instead, the widget may use its own time range (denoted as T1) and a second time range (denoted as T2) for comparison. Therefore, if the customer changes the time range using the time picker, the data on the widget may not change accordingly, unless they also update the widget's time range settings. To fix this problem, the customer can either disable the Compare Time Range feature or adjust the widget's time range to match the global time range. Reference:

[Time Range Comparisons](#)

Question: 19

If you clear the "Mark Business Transaction as error" checkbox under Configuration > Instrumentation > Error Detection, which two statements are true? (Choose two)

- A. The transaction is included in Response Time metrics.
- B. The transaction is added to error count metrics
- C. The transaction does not add to error count metrics.
- D. The transaction does not add to call count metrics.

Answer: AB

Explanation:

The Mark Business Transaction as Error checkbox under Configuration > Instrumentation > Error Detection allows you to control whether a business transaction is reported as an error or not. If you clear this checkbox, the business transaction will not be marked as an error and will not affect the error count metrics.

According to the [Error Detection: Mark Business Transaction as Error - AppDynamics](#), the following statements are true when you clear this checkbox:

The transaction is included in Response Time metrics: This means that the transaction will still contribute to the response time metric, which measures the average time it takes for a request to be completed by AppDynamics. However, it will not affect other response time metrics, such as First Response Time or Average Time to First Byte.

The transaction is added to error count metrics: This means that the transaction will still count towards the error count metric, which measures the number of transactions that have failed due to errors. However, it will not affect other error count metrics, such as Error Rate or Error Percentage.

Therefore, A (The transaction is included in Response Time metrics) and B (The transaction is added to error

count metrics) are true. Reference:

[Error Detection: Mark Business Transaction as Error - AppDynamics](#)

[Business Transaction - AppDynamics](#)

[Error Detection - AppDynamics](#)

Question: 20

To capture business data for Transaction Analytics what must be configured first?

- A. a Custom Metric
- B. a Data Collector
- C. a Health Rule
- D. an Information Point

Answer: B

Explanation:

To capture business data for Transaction Analytics, you must first configure a data collector. A data collector is a configuration that specifies what data to capture from a business transaction and how to store it in the analytics platform. You can capture data from HTTP requests, method invocations, custom fields, or information points. You can also apply filters and transformations to the captured data. Once you have configured a data collector, you can use it to enrich the transaction analytics data with business context and enable advanced queries and visualizations

<https://docs.appdynamics.com/accounts/en/cisco-appdynamics-saas-deployment-user-management/manage-custom-roles-for-cisco-appdynamics/transaction-analytics-permissions>

Question: 21

What are two situations where the UniqueHostID agent property is recommended to be used? (Choose two.)

- A. when the machine agent is with an app agent on different machines
- B. when the machine agent is on multiple machines
- C. when the machine agent is with multiple app agents on the same machine
- D. when the machine agent is with an app agent on the same machine

Answer: CD

Explanation:

The UniqueHostID agent property is recommended to be used when the machine agent is with multiple app agents on the same machine or when the machine agent is with an app agent on the same machine. This property allows you to specify a unique identifier for the host machine that overrides the default hostname resolution. [This helps to avoid conflicts or confusion when multiple agents are running on the same host or when the hostname is not consistent across different agents](#)

[Reference: 1: How are you able to associate the machine agent with application](#)

Question: 22

Which two conditions would be a reason to update an existing Data Collector? (Choose two.)

- A. The Business Transaction has been deprecated.
- B. A new class and method have been deployed to the application.
- C. A new HTTP parameter has been added to an existing Business Transaction
- D. A new method parameter has been added to an instrumented class

Answer: BD

Explanation:

[According to the Data Collectors document](#), a data collector is a configuration that captures application data from a method invocation or an HTTP request. A data collector is associated with a specific method signature or an HTTP parameter name. Therefore, if a new class and method have been deployed to the application, or a new method parameter has been added to an instrumented class, the existing data collector may not match

the new code and may need to be updated.

However, if the business transaction has been deprecated, or a new HTTP parameter has been added to an existing business transaction, the existing data collector may still work as expected, unless the data collector is specifically configured to filter by the business transaction name or the HTTP parameter value.

Reference:

[Data Collectors](#)

Question: 23

Which three types of performance degradation analysis are used with Business Transaction metrics? (Choose three.)

- A. Correlation Analysis
- B. Remote Services Analysis
- C. Scalability Analysis
- D. JMX Analysis
- E. Compare Releases

Answer: ABE

Explanation:

Business Transaction metrics are the key performance indicators that measure the health and performance of your applications. They include metrics such as average response time, calls per minute, errors per minute, and so on. You can use various types of performance degradation analysis to identify and troubleshoot the root causes of performance issues using Business Transaction metrics.

According to the [Troubleshoot Business Transactions - AppDynamics](#), the following types of performance degradation analysis are used with Business Transaction metrics:

Correlation Analysis: This type of analysis helps you to find the correlation between different metrics and events that may affect the performance of your Business Transactions. For example, you can use the correlation analysis to see how the response time of a Business Transaction is related to the CPU utilization of a node, or how the error rate of a Business Transaction is related to the number of slow database calls.

Remote Services Analysis: This type of analysis helps you to understand how the performance of your Business Transactions is impacted by the calls to remote services, such as databases, web services, message queues, and so on. For example, you can use the remote services analysis to see the response time, call count, error count, and wait time of each remote service call, and drill down to the details of each call.

Compare Releases: This type of analysis helps you to compare the performance of your Business Transactions across different releases of your application. For example, you can use the compare releases analysis to see how the response time, error rate, and throughput of your Business Transactions have changed after a new deployment, and identify any performance regressions or improvements.

Therefore, A (Correlation Analysis), B (Remote Services Analysis), and E (Compare Releases) are three types of performance degradation analysis that are used with Business Transaction metrics. Reference:

[Troubleshoot Business Transactions - AppDynamics](#)

[Business Transaction - AppDynamics](#)

[Cisco AppDynamics Associate Performance Analyst Certification](#)

Question: 24

Which two statements about Health Rules are true? (Choose two.)

- A. It sends an email
- B. It remediates a problem
- C. It changes the status of a KPI
- D. Status changes are represented as an event

Answer: CD

Explanation:

Health Rules are the rules that define the normal or expected behavior of your application, such as the response time, error rate, or CPU utilization of a business transaction, tier, node, or backend. You can create health rules to monitor the performance and availability of your application and its components. When a health rule is violated, it changes the status of the corresponding key performance indicator (KPI) from normal to warning or critical, depending on the severity of the violation. The status change is also represented as an event in the AppDynamics Controller UI, which can trigger alerts, policies, and actions. You can view the health rule events in the Events tab of the application dashboard, or in the Health tab of the entity dashboard.

<https://developer.cisco.com/docs/appdynamics/health-rules/>

<https://university.appdynamics.com/certifications/associate-administrator>

Question: 25

What is the main benefit of defining Database Collectors?

- A. You can see which applications are calling a database
- B. You can view the calls per minute from each instrumented application
- C. You can view session data.
- D. You can clean up stale database connections

Answer: A

Explanation:

The main benefit of defining Database Collectors is that you can see which applications are calling a database and how they impact the database performance. [Database Collectors enable the Database Agent to collect data from your database instances and database servers, such as metrics, queries, sessions, wait states, and more](#)¹ You can also link a database on the application flow maps to a database instance monitored by Database Visibility, which allows you to see the correlation between the application and the database²
[Reference: 1: Add Database Collectors 2:](#) [Link a Database on the Application Flow Map to a Database Instance Monitored by Database Visibility]

Question: 26

Which two reasons would cause an organization to upgrade AppDynamics agents? (Choose two.)

- A. agent issues
- B. database issues
- C. application upgrade
- D. new features

Answer: AD

Explanation:

[According to the Upgrade the Java Agent document](#)¹, one of the reasons to upgrade AppDynamics agents is to resolve agent issues, such as bugs, performance problems, or compatibility issues with the Controller or the monitored application. Another reason to upgrade AppDynamics agents is to take advantage of new features, such as enhanced metrics, dashboards, alerts, or integrations, that are introduced in newer agent versions. Therefore, the correct answer is A and D. Database issues and application upgrade are not directly related to the agent upgrade, unless they affect the agent functionality or compatibility. Reference:

Question: 27

Which statement about assigning users to a group is true?

- A. When you assign a user to a group, they can only have the roles assigned to that group
- B. You can list all users in a group that are logged into AppDynamics.
- C. When the user logs in they can log in with the group credentials
- D. When you add a role to a group, every user in the group is assigned that role

Answer: D

Explanation:

When you assign users to a group in AppDynamics, you can manage their permissions and access levels more easily. A group is a collection of users who share the same roles and privileges. A role is a set of permissions that define what actions a user can perform on the Controller UI or the Events Service.

According to the [System Configuration Guide for Cisco Unified Communications Manager - User Management](#), the following statement is true about assigning users to a group:

When you add a role to a group, every user in the group is assigned that role: This means that the users in the group inherit the permissions of the role that you add to the group. For example, if you add the Administrator role to a group, then all the users in that group can perform administrative tasks on the Controller UI or the Events Service.

The other statements are false because:

When you assign a user to a group, they can only have the roles assigned to that group: This is not true because a user can belong to multiple groups and have multiple roles. For example, a user can be in both the Administrator and the Operator groups and have the permissions of both roles.

You can list all users in a group that are logged into AppDynamics: This is not true because there is no such feature in AppDynamics that allows you to see the login status of the users in a group. You can only see the list of users and groups in the User Management page of the Controller UI or the Events Service.

When the user logs in they can log in with the group credentials: This is not true because a user cannot log in with the group credentials. A user can only log in with their own username and password. The group credentials are not used for authentication, but for authorization.

Therefore, D (When you add a role to a group, every user in the group is assigned that role) is the correct answer.

Reference:

[System Configuration Guide for Cisco Unified Communications Manager - User Management](#)

[Cisco AppDynamics Associate Administrator Certification](#)

[User Management - AppDynamics](#)

Question: 28

When creating a Health Rule condition, which two metrics could baselines be used instead of defining a specific threshold? (Choose two.)

- A. overall average response time
- B. Business Transaction errors per minute
- C. a metric with an average baseline of zero
- D. an expression containing multiple metrics
- E. calls per millisecond

Answer: AD

Explanation:

When creating a Health Rule condition, you can use baselines instead of defining a specific threshold for some metrics. Baselines are the dynamic values that represent the normal or expected behavior of a metric, based on the historical data collected by the AppDynamics Cognition Engine. Using baselines can help you account for the variations and seasonality

of your application performance, and avoid false positives or negatives. You can use baselines for metrics that have a non-zero average value and a meaningful deviation, such as the overall average response time of a business transaction or an expression containing multiple metrics. You cannot use baselines for metrics that have a zero average value or a negligible deviation, such as the Business Transaction errors per minute or the calls per millisecond.

<https://www.appdynamics.com/product/how-it-works/application-performance-baseline>

<https://developer.cisco.com/docs/appdynamics/health-rules/>

Question: 29

Where do you go to create a user group?

- A. Settings > AppDynamics Groups > Config Management
- B. Settings > My AppDynamics Account > Groups
- C. Settings > My Preferences > Groups
- D. Settings > Administration > Groups

Answer: D

Explanation:

To create a user group in AppDynamics, you need to go to Settings > Administration > Groups. A user group is a collection of users who share the same roles and privileges. You can use user groups to manage user access and permissions more easily. You can add, edit, delete, or duplicate user groups from the Groups tab on the Administration page.

According to the [Manage Users and Groups - AppDynamics](#), the following steps are required to create a user group:

Click the Add button on the Groups tab.

Enter a name and description for the group.

Select one or more roles for the group from the Available Roles list. You can also create a new role by clicking the Add Role button.

Click Save.

Therefore, D (Settings > Administration > Groups) is the correct answer. Reference:

[Manage Users and Groups - AppDynamics](#)

Question: 30

To alert the operations team that a cloud-based application has provisioned a new instance of a JVM, which standard action should an administrator take?

- A. Create a Policy trigger for Server Restart > New Node Discovered.
- B. Set up a simple extension that checks the number of nodes in a Tier and insert an event
- C. Create a Policy trigger for Application Change > New Tier
- D. Create a Policy trigger for Discovery > New Node Discovered

Answer: D

Explanation:

To alert the operations team that a cloud-based application has provisioned a new instance of a JVM, an administrator should create a policy trigger for Discovery > New Node Discovered. This policy trigger will fire an event when a new node is discovered by the AppDynamics Controller, either by auto-discovery or manual registration. The event will contain information such as the node name, tier name, application name, and agent type. The administrator can then configure an action to send an email, SMS, or custom notification to the operations team with the event details. Alternatively, the administrator can also use the built-in alerting feature of the cloud provider to notify the operations team of the new instance creation.

<https://docs.appdynamics.com/display/PRO42/Alert+and+Respond>

Question: 31

Before creating a new database collector, which two actions are required? (Choose two.)

- A. Create a new DB Agent for each new database
- B. Verify the database connection details
- C. Create a database user and set user permissions
- D. Verify the Time Spent in Executions.
- E. Verify the collected metrics.

Answer: BC

Explanation:

Before creating a new database collector, you need to perform the following actions:

Verify the database connection details. You need to provide the host, port, database name, and credentials for the database that you want to monitor. [You also need to select the database type and the collector type1](#)

Create a database user and set user permissions. You need to create a database user with the minimum required permissions to access the database and run the queries that the Database Agent needs. [You can use the scripts provided by AppDynamics to create the user and grant the permissions2](#)

Reference: 1: [Add Database Collectors](#) 2: [Database Permissions](#)

Question: 32

The customer has several Business Transactions which take significant time to execute. These transactions are continuously reported as being Slow or Stalled transactions An SLA has been established for these transactions How does the engineer set static thresholds for these Business Transactions?

- A. Set the default thresholds for Slow and Stalled transactions
- B. There is no way to adjust Slow Transaction thresholds on individual Business Transactions
- C. Select Individual Transaction Thresholds and set the thresholds for the long-running transactions.

D. Adjust the Health Rules so these long-running transactions are not considered slow or stalled

Answer: C

Explanation:

[According to the Transaction Thresholds document¹](#), you can set static thresholds for individual business transactions by selecting Individual Transaction Thresholds from the Business Transaction Configuration window. You can then specify the values for Slow, Very Slow, and Stall thresholds for each business transaction. This way, you can customize the thresholds according to the SLA and the expected performance of the long-running transactions. Setting the default thresholds or adjusting the health rules would affect all the business transactions, not just the long-running ones. Therefore, the correct answer is C. Reference:

[Transaction Thresholds](#)

Question: 33

Which statement about assigning roles to a group is true?

- A. A role must be assigned to at least one group
- B. All users in a group have the same roles.
- C. A role can only be assigned to one group
- D. A group can contain multiple roles

Answer: D

Explanation:

A group can contain multiple roles, which means that the users in the group inherit the permissions of all the roles assigned to the group. [For example, you can create a group that has both the Administrator and the Analytics Administrator roles, and the users in that group can perform both the administration and the analytics tasks](#)¹ Reference: [1: Roles and Permissions](#)

Question: 34

What kind of problems are detected with Infrastructure monitoring? (Choose One)

- A. slow method calls
- B. thread pool problems
- C. garbage collection problems
- D. blocked threads
- E. process level problems

Answer: E

Explanation:

Infrastructure monitoring is the feature in AppDynamics that allows you to monitor and troubleshoot the health and availability of the servers that host your applications and their components. Infrastructure monitoring includes metrics such as CPU utilization, memory usage, disk I/O, network traffic, and so on. You can use infrastructure monitoring to detect and resolve the problems that affect the performance and reliability of your servers and applications.

According to the [Server Visibility - AppDynamics](#), the following kind of problem is detected with infrastructure monitoring:

Process level problems: These are the problems that occur at the process level of the server, such as high CPU consumption, memory leaks, excessive disk I/O, or network congestion. You can use infrastructure monitoring to identify the processes that are consuming the most resources, and drill down to the details of each process, such as the command line, the environment variables, the open files, the threads, and the sockets. You can also use infrastructure monitoring to correlate the process level problems with the application performance problems, such as slow response time, high error rate, or low throughput.

Therefore, E (process level problems) is the correct answer. Reference:

[Server Visibility - AppDynamics](#)

[Cisco AppDynamics Associate Administrator Certification](#)

[Server Monitoring - AppDynamics](#)

Question: 35

Your organization has defined a set of health rules that accurately represent the business functions of its applications.

Which two ways should you help the organization see the current status of health rules and events of the business process? (Choose two.)

- A. Set up a policy to send an email or text when a health rule status changes.
- B. Create a report that runs every night and sends out the status.
- C. Leave the main applications page open for all to see
- D. Create a dashboard with health and events widgets displayed.

Answer: AD

Explanation:

Health rules are the criteria that AppDynamics uses to evaluate the performance and availability of your applications and their components. Health rules can trigger alerts and actions when the conditions are violated or cleared. Events are the occurrences that AppDynamics records when something significant happens in your monitored environment, such as a health rule violation, a policy execution, a deployment, or a custom event.

According to the [Health Rules - AppDynamics](#), the following two ways can help you see the current status of health rules and events of the business process:

Set up a policy to send an email or text when a health rule status changes: This way, you can get notified in real time when a health rule condition is met or cleared, and take appropriate actions to resolve or prevent issues. You can configure policies to send emails, texts, or other notifications to

specific recipients or groups, and include relevant information such as the health rule name, the affected entity, the severity level, and the time range.

Create a dashboard with health and events widgets displayed: This way, you can visualize the health and performance of your applications and their components in a single view, and drill down to the details of each health rule or event. You can create custom dashboards with various widgets that show the health status, the event list, the event summary, the event correlation, and the event timeline of your applications and their components.

Therefore, A (Set up a policy to send an email or text when a health rule status changes) and D (Create a dashboard with health and events widgets displayed) are two ways that can help you see the current status of health rules and events of the business process. Reference:

[Cisco AppDynamics Associate Administrator Certification](#)

[Health Rules - AppDynamics](#)

[Policies and Actions - AppDynamics]

[Events - AppDynamics]

[Custom Dashboards - AppDynamics]

Question: 36

Which two actions would trigger a health rule violation canceled warning event or a health rule violation canceled critical event? (Choose two.)

- A. Mark the health rule as canceled
- B. Add a new condition to the health rule.
- C. Disable the health rule
- D. Remove the health rule from a policy

Answer: BC

Explanation:

[A health rule violation canceled warning event or a health rule violation canceled critical event is generated when the health rule processor can no longer accurately assert that the health rule violation continues to violate or that it has ended¹. This can happen when the health rule is edited or disabled, among other reasons¹. Adding a new condition to the health rule or disabling the health rule are two actions that would trigger a health rule violation canceled event, as they would change the evaluation criteria or stop the evaluation altogether¹.](#)

[Marking the health rule as canceled or removing the health rule from a policy are not valid actions, as they do not affect the health rule configuration or status](#)

<https://docs.appdynamics.com/appd/22.x/22.1/en/appdynamics-essentials/alert-and-respond/troubleshoot-health-rule-violations>

Question: 37

A new Java process instrumented with the AppDynamics agent started with no errors in the log file. There is verified network connectivity between the host running Java and the Controller on all ports and load applied to the process.

However, the host does not appear in the Tiers and Nodes view, and there are no metrics in the Metric Browser. What else should you check?

- A. Settings > My Preferences, and Enable Debug Mode to see more detailed error messages
- B. Troubleshoot > Errors, and observe changes as you vary the time window
- C. Configuration > Development Level Monitoring, and enable Development Level Monitoring
- D. Settings > License, and audit the license count for that agent type

Answer: D

Explanation:

One possible reason why the host does not appear in the Tiers and Nodes view and there are no metrics in the Metric Browser is that the license count for that agent type has been exceeded. You can check the license usage and availability by going to Settings > License in the Controller UI. [If the license count is exceeded, you can either purchase more licenses or deactivate some agents to free up the license slots](#)¹² Reference: 1: [License Rules and Restrictions](#) 2: [Manage Licenses](#)

Question: 38

Which two types of standard controller Actions are available to be created? (Choose two.)

- A. Resolve
- B. Close Incident

C. Undeploy

D. Email

E. Thread Dump

Answer: DE

Explanation:

[According to the Actions document1](#), there are four types of standard controller actions that are available to be created: Notification, Diagnostic, Remediation, and JIRA. Each type of action has different subtypes that perform specific tasks. For example, the Notification type includes Email, SMS, and HTTP Request subtypes, while the Diagnostic type includes Thread Dump, Heap Dump, and Diagnostic Session subtypes. Therefore, the correct answer is D and E, as Email and Thread Dump are subtypes of standard controller actions. Resolve, Close Incident, and Undeploy are not standard controller actions, but they may be custom actions that are created by the user or the administrator. Reference:

[Actions](#)

Question: 39

In AppDynamics, what are two places you would find metrics to be added to a custom dashboard? (Choose two.)

A. the Health Rules configuration

B. the Node list

C. the JMX tab on the Node

D. the Metric Browser

Answer: CD

Explanation:

Metrics are the numerical values that AppDynamics collects and reports to measure the performance and availability of your applications and their components. You can use metrics to create custom dashboards that display the data that is relevant to your needs and interests.

According to the [Metrics - AppDynamics](#), the following two places are where you would find metrics to be added to a custom dashboard:

The JMX tab on the Node: This tab shows the Java Management Extensions (JMX) metrics that are exposed by the Java Virtual Machine (JVM) or the application server. JMX metrics provide information about the JVM memory, threads, garbage collection, and other aspects of the Java environment. You can add JMX metrics to a custom dashboard by selecting them from the metric browser tree and dragging them to the dashboard canvas.

The Metric Browser: This is a tool that allows you to browse, search, and view all the metrics that are available in AppDynamics. The metric browser shows the metrics in a hierarchical tree structure that reflects the logical organization of your applications and their components. You can add metrics to a custom dashboard by selecting them from the metric browser tree and dragging them to the dashboard canvas.

Therefore, C (the JMX tab on the Node) and D (the Metric Browser) are two places where you would find metrics to be added to a custom dashboard. Reference:

[Metrics - AppDynamics](#)

[Custom Dashboards - AppDynamics](#)

[Cisco AppDynamics Associate Administrator Certification](#)

Question: 40

What is used to capture application data in a single Business Transaction?

A. Data Collector

- B. Windows Performance Counters
- C. Information Point
- D. MX Rules

Answer: A

Explanation:

A data collector is used to capture application data in a single business transaction. A data collector is a configuration that specifies what data to capture from a business transaction and how to store it in the AppDynamics Controller or the analytics platform. You can capture data from HTTP requests, method invocations, custom fields, or information points. You can also apply filters and transformations to the captured data. The captured data can be used for various purposes, such as adding context to business transactions, creating custom metrics, triggering health rules, enriching transaction analytics, or creating business journeys.

<https://university.appdynamics.com/certifications/associate-administrator>

<https://docs.appdynamics.com/appd/22.x/latest/en/application-monitoring/business-transactions>

Question: 41

You have validated that the agent is able to register and communicate with the controller successfully by reviewing the agent and controller logs. However in the controller UI, you do not see the agent report under Tiers and Nodes where it was expected. Where in the controller UI would you find this information that shows the current Application, Tier, and Node the agent currently reports to?

- A. under View my settings, AppDynamics Agents section
- B. under View my settings, License section
- C. under View my settings, Administration section
- D. under Application: Troubleshoot section

Answer: A

Explanation:

To find the information that shows the current application, tier, and node that the agent currently reports to, you can go to View my settings > AppDynamics Agents section in the Controller UI. This section shows you the list of agents that are connected to the Controller, along with their status, version, application, tier, and node names. [You can also filter the list by agent type, status, or application](#)¹ Reference: [1: AppDynamics Agents](#)

Question: 42

While troubleshooting a performance issue on a Java application the engineer determines there is a possible memory leak in the JVM Using AppDynamics, how would the engineer determine if there is a memory leak?

- A. Examine the values on the Server tab on one of the affected Nodes.
- B. Configure Object Instance Tracking on the Tier in question.
- C. Verify and adjust the Memory Monitoring configuration for the Tier in question
- D. Analyze the information on the Memory tab on one of the affected Nodes

Answer: D

Explanation:

[According to the Java Memory Leaks document](#)¹, one of the ways to detect and troubleshoot Java memory leaks is to use the Automatic Leak Detection feature on the Memory tab of the Node Dashboard. This feature captures and analyzes the collections that are actively used and growing in size over time, indicating a potential memory leak. You can also drill down into the leaking collections to see the content summary and the access traces that show the code path and the business transactions that access the collection. Therefore, the correct answer is D. Examining the values on the Server tab, configuring Object Instance Tracking, or verifying and adjusting the Memory Monitoring configuration may not directly reveal the source of the memory leak, but they may provide some useful information or metrics to support the diagnosis. Reference:

[Java Memory Leaks](#)

Question: 43

After agents are installed, and application data is flowing through a Java application, a user sees that a remote database has not been discovered in a known Business Transaction. To determine and configure the exit call to the database where would you go?

- A. Configuration > Java > Backend Detection > JDBC
- B. Databases > Configuration > Collectors
- C. Configuration > Instrumentation > Backend Detection > Java
- D. Database Calls > Options > Backend Detection > JDBC

Answer: C

Explanation:

To determine and configure the exit call to the database, you need to go to Configuration > Instrumentation > Backend Detection > Java. An exit call is a call from an application component to an external service or system, such as a database, a web service, a message queue, or a cache. A backend is a logical representation of an external service or system that receives exit calls from an application component. AppDynamics automatically detects and identifies common types of backends, such as JDBC, HTTP, JMS, and so on.

According to the [Exit Point Detection Rules - AppDynamics](#), the following steps are required to determine and configure the exit call to the database:

Go to Configuration > Instrumentation > Backend Detection > Java.

Select the application and the tier that contains the exit call to the database.

Click the Add button to create a new exit point detection rule.

Enter a name and description for the rule.

Select the type of exit call from the drop-down list, such as JDBC, HTTP, JMS, and so on.

Specify the class and method that make the exit call to the database.

Specify the match criteria for the exit call, such as the URL, the query, the host, the port, and so on.

Click Save.

Therefore, C (Configuration > Instrumentation > Backend Detection > Java) is the correct answer. Reference:

[Exit Point Detection Rules - AppDynamics](#)

[Cisco AppDynamics Associate Administrator Certification](#)

[Backends - AppDynamics](#)

Question: 44

On a custom dashboard, which three types of information does a non-administrator user view? (Choose three.)

- A. the current weather in Seattle
- B. the number of AppDynamics licenses consumed
- C. the company logo
- D. the version of the AppDynamics Controller
- E. a Timeseries graph

Answer: CDE

Explanation:

On a custom dashboard, a non-administrator user can view the following types of information:

The company logo: A custom dashboard can display the company logo in the upper left corner of the dashboard. This can help to brand the dashboard and make it more recognizable for the intended audience. [The logo can be configured by the dashboard owner or administrator from the dashboard settings1.](#)

The version of the AppDynamics Controller: A custom dashboard can display the version of the AppDynamics Controller in the upper right corner of the dashboard. This can help to identify the Controller that is hosting the dashboard and the features that are available. [The version can be configured by the dashboard owner or administrator from the dashboard settings1.](#)

A Timeseries graph: A custom dashboard can display a Timeseries graph widget that shows the trend of a metric over time. This can help to visualize the performance and behavior of the application, server, or database that is being monitored. [A Timeseries graph can be configured by the dashboard owner or administrator from the widget settings2.](#)

A non-administrator user cannot view the following types of information on a custom dashboard:

The current weather in Seattle: A custom dashboard cannot display the current weather in Seattle or any other location, as this is not a metric or data point that is reported by AppDynamics agents or the analytics platform. [A custom dashboard can only display the data that is collected and stored by AppDynamics2.](#)

The number of AppDynamics licenses consumed: A custom dashboard cannot display the number of AppDynamics licenses consumed, as this is a Controller-level information that is only accessible by the Controller administrator. [A custom dashboard can only display the data that is scoped by the dashboard owner's permissions1.](#)

<https://docs.appdynamics.com/accounts/en/cisco-appdynamics-saas-deployment-user-management/manage-custom-roles-for-cisco-appdynamics/custom-dashboard-permissions>

Question: 45

What minimum details are required to properly configure a database agent from a zip file?

- A. Controller Host Database Agent Name. Account Access Key. Account Name
- B. Database Host Database Port Login Credentials
- C. Controller Host Controller Port, Account Access Key. Account Name
- D. Database Host Login Credentials

Answer: C

Explanation:

<https://docs.appdynamics.com/display/PRO45X/Install+the+Database+Agent>

To properly configure a database agent from a zip file, you need to provide the following minimum details:

Controller Host: This is the hostname or IP address of the machine where the Controller is installed. [You need to specify this so that the Database Agent can communicate with the Controller and send the collected data1](#)

Controller Port: This is the port number that the Controller listens on for incoming connections. You need to specify this so that the Database Agent can connect to the Controller using the correct protocol. [The default ports are 8090 for HTTP and 8181 for HTTPS for on-premise Controllers, and 80 for HTTP and 443 for HTTPS for SaaS Controllers1](#)

Account Access Key: This is the key that authenticates the Database Agent with the Controller. You need to specify this so that the Database Agent can register itself with the Controller and access the license. [You can find the account access key in the Controller UI under Settings > License > Account1](#)

Account Name: This is the name of the account that the Database Agent belongs to. You need to specify this so that the Database Agent can associate itself with the correct application and user permissions. [You can find the account name in the Controller UI under Settings > License > Account1](#)

[Reference: 1: Install the Database Agent](#)

Question: 46

Which three widgets can be on a Custom Dashboard? (Choose three.)

- A. Flow Map
- B. iFrame
- C. Health List
- D. Snapshot
- E. Funnel Analysts

Answer: ABC

Explanation:

[According to the Widgets document1](#), there are various types of widgets that can be used to create visual representations of data in custom dashboards and war rooms. Some of the widget types are:

Flow Map: A flow map widget displays the flow of traffic between different entities, such as applications, tiers, nodes, backends, and so on. You can customize the flow map to show different metrics, filters, and drill-down options.

iFrame: An iFrame widget allows you to embed an external web page or another dashboard within a custom dashboard. You can specify the URL, height, and width of the iFrame widget.

Health List: A health list widget displays the health status of different entities, such as applications, tiers, nodes, business transactions, and so on. You can customize the health list to show different metrics, filters, and drill-down options.

Therefore, the correct answer is A, B, and C. Snapshot and Funnel Analysts are not widget types, but features or tools that can be accessed from other widgets or dashboards. Reference:

[Widgets](#)

Question: 47

Capacity planning for a regionally partitioned application requires that the administrator monitors the number of users and orgs mapped to a database. How does the administrator meet this requirement?

- A. Schedule an analytics query to check and report on the result periodically
- B. Write code in a custom action that alerts when a region has reached capacity
- C. Use a custom metric for databases and alert when it approaches a specific level
- D. Watch the size of the database and alert when it grows too much

Answer: A

Explanation:

Capacity planning is the process of estimating and allocating the resources needed to meet the current and future demands of an application. For a regionally partitioned application, where different users and organizations are assigned to different databases based on their location, capacity planning requires that the administrator monitors the number of users and orgs mapped to each database and ensures that they do not exceed the available resources.

According to the [AppDynamics Analytics - AppDynamics](#), one way to meet this requirement is to schedule an analytics query to check and report on the result periodically. An analytics query is a SQL- like query that allows you to analyze and visualize the data collected by AppDynamics, such as business transactions, logs, metrics, events, and custom data. You can use analytics queries to create dashboards, widgets, reports, and alerts that provide insights into the performance and behavior of your applications and their components.

To schedule an analytics query to check and report on the number of users and orgs mapped to a database, the administrator can follow these steps:

Go to Analytics > Search.

Select the data source that contains the user and org information, such as Business Transactions, Logs, or Custom Data.

Enter the query that counts the number of users and orgs for each database, such as SELECT database, COUNT(DISTINCT user), COUNT(DISTINCT org) FROM data_source GROUP BY database.

Click the Save button and enter a name and description for the query.

Click the Schedule button and select the frequency, time range, and recipients for the query execution and report delivery.

Click Save.

Therefore, A (Schedule an analytics query to check and report on the result periodically) is the correct answer.

Reference:

[AppDynamics Analytics - AppDynamics](#)

Question: 48

Which role would you need to grant to an existing user for them to manage Controller security settings?

- A. Credentials Administrator
- B. Administrator
- C. Security Administrator

D. Analytics Administrator

E. Account Owner

Answer: C

Explanation:

To manage Controller security settings, you would need to grant the Security Administrator role to an existing user. [The Security Administrator role is a built-in role that allows the user to access and modify the security settings of the Controller, such as encryption, certificates, authentication, authorization, and audit¹. The Security Administrator role also inherits the permissions of the Credentials Administrator role, which allows the user to manage the credentials for the Controller and the agents². The other roles listed do not have the permission to manage Controller security settings. The Administrator role can perform most administrative tasks, but not security-related ones². The Analytics Administrator role can manage the analytics platform and data, but not the Controller security³. The Account Owner role can manage the account settings and licenses, but not the Controller security⁴](#)

<https://docs.appdynamics.com/appd/21.x/21.1/en/application-monitoring/install-app-server-agents/agent-to-controller-connections>

Question: 49

Where do you go to create a user in the web user interface?

A. Settings > Administration

B. Settings > AppDynamics Agents

C. Settings > My AppDynamics Account

D. Settings > Groups and Roles > Create User

Answer: D

Explanation:

To create a user in the web user interface, you need to go to Settings > Groups and Roles > Create User. This option allows you to add a new user account and assign it to one or more groups. [You can also specify the user's name, email, password, and authentication method¹](#) Reference: 1: [Create Users](#)

Question: 50

A customer needs to understand how establishing standards for naming Applications, Tiers, and Nodes can help them determine the root of a problem when users are working with AppDynamics or receive an alert Which three statements about establishing naming conventions are true? (Choose three.)

- A. Application Tier and Node values can be configured in the application agent startup script
- B. Node names must be unique across the entire business application
- C. Nodes that reside on different Tiers and different machines (hosts) can have duplicate Node names.
- D. Nodes that reside on the same Tier but on different machines (hosts) can have duplicate Node names.
- E. Application Tier and Node values can be configured in the controller-info.xml file

Answer: ABE

Explanation:

[According to the Tiers and Nodes document1](#), establishing standards for naming applications, tiers, and nodes can help you identify and organize the components of your monitored environment, as well as troubleshoot performance issues and health rule violations. Some of the statements about establishing naming conventions are true:

Application Tier and Node values can be configured in the application agent startup script: This is true for Java agents, as you can pass values to the agent that will be used as the application, tier, and node name for the system. [For example, you can use the -Dappdynamics.agent.applicationName, -Dappdynamics.agent.tierName, and -Dappdynamics.agent.nodeName properties to specify the names in the JVM startup script2](#).

Node names must be unique across the entire business application: This is true, as node names are used to identify the individual instances of the application components that are monitored by AppDynamics. Having duplicate node names can cause confusion and errors in the data collection and reporting. [Therefore, you should use a naming convention that ensures uniqueness and clarity for each node1](#).

Application Tier and Node values can be configured in the controller-info.xml file: This is true for Java agents, as you can also configure the names in the controller-info.xml file, which is located in the <agent_home>/conf directory. The controller-info.xml file contains the properties that define the agent identity and connection to the Controller. [You can use the <application-name>, <tier-name>, and <node-name> elements to specify the names in the file2](#).

Some of the statements about establishing naming conventions are false:

Nodes that reside on different Tiers and different machines (hosts) can have duplicate Node names: This is false, as node

names must be unique across the entire business application, regardless of the tier or the host they belong to. Having duplicate node names can cause confusion and errors in the data collection and reporting. [Therefore, you should use a naming convention that ensures uniqueness and clarity for each node1.](#)

Nodes that reside on the same Tier but on different machines (hosts) can have duplicate Node names: This is also false, as node names must be unique across the entire business application, regardless of the tier or the host they belong to. Having duplicate node names can cause confusion and errors in the data collection and reporting. [Therefore, you should use a naming convention that ensures uniqueness and clarity for each node1.](#)

Reference:

[Tiers and Nodes](#)

[Name Business Applications, Tiers, and Nodes](#)

Question: 51

The applications and servers are running normally. What are two important reasons for the engineer to view and analyze server performance data at this time? (Choose two.)

- A. Visualizing application performance during normal operations helps identify when the application performance is abnormal.
- B. Graphs and charts collected at this time are useful for showing application owners the stability of their application.
- C. Collecting Server metric performance during normal operations helps determine code problems.
- D. Collecting the errors during this period will identify memory problems in the application.

Answer: AB

Explanation:

Server performance data is the information that AppDynamics collects and reports about the health and availability of the servers that host your applications and their components. Server performance data includes metrics such as CPU utilization, memory usage, disk I/O, network traffic, and so on. You can use server performance data to monitor and troubleshoot the performance issues of your servers and applications.

According to the [Server Visibility - AppDynamics](#), the following two reasons are important for the engineer to view and analyze server performance data at this time:

Visualizing application performance during normal operations helps identify when the application performance is abnormal: This means that by comparing the current server performance data with the historical or baseline data, you can detect any anomalies or deviations that may indicate a potential problem or degradation of your application performance. For example, you can use the Server Dashboard to see the current and historical trends of the server metrics, and use the Anomaly Detection feature to identify any outliers or unusual patterns.

Graphs and charts collected at this time are useful for showing application owners the stability of their application: This means that by using the graphical and visual representations of the server performance data, you can communicate and demonstrate the reliability and efficiency of your application to the stakeholders and customers. For example, you can use the Server Health Report to generate a summary of the server health status, availability, and performance over a specified time period, and share it with the application owners.

Therefore, A (Visualizing application performance during normal operations helps identify when the application performance is abnormal) and B (Graphs and charts collected at this time are useful for showing application owners the stability of their application) are two important reasons for the engineer to view and analyze server performance data at this time. Reference:

[Server Visibility - AppDynamics](#)

[Cisco AppDynamics Associate Administrator Certification](#)

[Server Monitoring - AppDynamics](#)

Question: 52

Users are saying they are experiencing occasional slow response time. You suspect a database issue and want to find snapshots showing slow database response times. Which two statements describe how to locate snapshots showing the slow database response? (Choose two.)

- A. Go to the list of snapshots and apply a filter based on slow database response
- B. Go to the dashboard for the database and select slowest database calls.
- C. From the flow map, right-click the database icon and select view snapshots.
- D. Go to the list of snapshots and select analyze for a group of snapshots.

Answer: AC

Explanation:

Snapshots are the detailed records of the execution of a business transaction, including the call stack, the exit points, the errors, and the metrics. To locate snapshots showing slow database response, you can use the following methods:

Go to the list of snapshots and apply a filter based on slow database response: You can access the list of snapshots from the Business Transaction Dashboard or the Tier Dashboard. You can then apply a filter to show only the snapshots that have a slow database response time, which is the time spent by the business transaction in calling the database. You can specify the minimum and maximum values for the slow database response time, or use the predefined options such as slower than 1 second, slower than 5 seconds, etc. You can also combine this filter with other filters, such as time range, business transaction name, tier name, node name, etc. [to narrow down the results1](#).

From the flow map, right-click the database icon and select view snapshots: You can access the flow map from the Application Dashboard, the Business Transaction Dashboard, or the Tier Dashboard. The flow map shows the graphical representation of the flow of the business transactions, including the tiers, the nodes, the backends, and the remote services. You can right-click on any database icon on the flow map and select view snapshots to see the list of snapshots that have called that database. [You can then sort or filter the snapshots by the database response time column to find the slowest ones2](#).

The other options are not valid methods to locate snapshots showing slow database response. Going to the dashboard for the database and selecting slowest database calls will show the metrics and the flow map for the database, but not the snapshots. Going to the list of snapshots and selecting

analyze for a group of snapshots will show the summary and the comparison of the selected snapshots, but not the database response time.

<https://community.appdynamics.com/t5/Share-a-tip/Snapshots-and-identifying-intermittent-issues/td-p/42896>

Question: 53

What two types of information can be shared without an AppDynamics login? (Choose two.)

- A. Flow Map
- B. Scheduled Reports
- C. Database Queries
- D. Transaction Snapshots
- E. Custom Dashboards

Answer: AE

Explanation:

Two types of information that can be shared without an AppDynamics login are Flow Map and Custom Dashboards. Flow Map is a graphical representation of the application topology, showing the interactions between the tiers, nodes, backends, and remote services. Custom Dashboards are user- defined views of the application performance data, metrics, and events. [You can share these types of information by generating a shareable link from the Controller UI, which can be accessed by anyone who has the link, without requiring a login](#)¹² Reference: 1: [Share a Flow Map](#) 2: [Share a Custom Dashboard](#)

Question: 54

The Default Flow map shows a tier with a full or partial red circle, which indicates one or more nodes have reported Health Rule violations. Which procedure should the engineer take?

- A. Use the metric browser to identify the Health Rule being violated, then notify the application owners of the issue
- B. This is a recurring issue: do nothing as the Health Rule violation will clear without intervention
- C. Disable the Health Rule to prevent further reporting of the issue, then notify the data center team to reboot the server
- D. Click on the Events tab for the application on the Default Row Map then analyze the health rule(s) being violated.

Answer: D

Explanation:

[According to the Troubleshoot Health Rule Violations document](#)¹, one of the ways to troubleshoot health rule violations is to click on the Events tab for the application on the Default Flow Map and then analyze the health rule(s) being violated. This will show you the details of the health rule violation events, such as the affected entity, the severity, the duration, the evaluation time, and the triggered actions. You can also drill down into the health rule violation details to see the health rule configuration, the metric values, and the snapshots. This will help you identify the root cause of the issue and take appropriate remediation actions. Therefore, the correct answer is D. Using the metric browser, disabling the health rule, or rebooting the server may not be the best or the most efficient ways to troubleshoot health rule violations, as they may not reveal the source of the problem, may affect other entities, or may not resolve the issue permanently. Reference:

[Troubleshoot Health Rule Violations](#)

Question: 55

When troubleshooting an application using AppDynamics, which order would help narrow down a code problem with a specific business transaction the fastest?

- A. Review Application Dashboard > Check transaction scorecard > Check Transactions snapshots for problems
- B. Review Exceptions > Look at Errors > Use Compare Releases
- C. Review Application Dashboard > Look at Events > Check Service Endpoints
- D. Review tier and node metrics > Check Database Calls and metrics > Review Application Dashboard

Answer: A

Explanation:

When troubleshooting an application using AppDynamics, the fastest way to narrow down a code problem with a specific business transaction is to follow this order:

Review Application Dashboard: This is the first step to get an overview of the health and performance of your application and its components. You can see the key metrics, such as response time, throughput, error rate, and load, for each tier, node, and business transaction. You can also see the health rule violations, events, and alerts that indicate any issues or anomalies in your application.

Check transaction scorecard: This is the second step to drill down to the details of a specific business transaction that you want to troubleshoot. You can see the transaction scorecard by clicking on the business transaction name on the application dashboard. The transaction scorecard shows the metrics, such as average response time, calls per minute, errors per minute, and slow calls, for the selected business transaction. You can also see the breakdown of the response time by tier, node, and remote service, and the distribution of the response time by percentile.

Check transaction snapshots for problems: This is the third step to identify the root cause of the code problem for the selected business transaction. You can see the transaction snapshots by clicking on the View Snapshots button on the transaction scorecard. Transaction snapshots are the detailed records of the execution of a business transaction, including the call graph, the exit calls, the errors, the exceptions, the logs, the HTTP parameters, the cookies, the session attributes, and the custom data. You can use the transaction snapshots to analyze the code execution, the database queries, the web service calls, the message queues, the errors, and the exceptions that affect the performance of the business transaction.

Therefore, A (Review Application Dashboard > Check transaction scorecard > Check Transactions snapshots for problems) is the correct answer. Reference:

[Cisco AppDynamics Associate Administrator Certification](#)

[Troubleshoot Business Transactions - AppDynamics](#)

[Business Transaction - AppDynamics](#)

[Transaction Snapshots - AppDynamics](#)

Question: 56

Java APM agents used in a production environment are consistently found to be disabled while no development environments exhibit this behavior. Given that agents in both environments use the default configuration, what solution should be implemented to resolve this issue?

- A. Enable agent auto-restart in the controller UI.
- B. Increase heap space for agent
- C. Decrease worker thread count
- D. Enable heap-storage-monitor-enabled setting on agent.

Answer: D

Explanation:

One possible reason why Java APM agents used in a production environment are consistently found to be disabled while no development environments exhibit this behavior is that the production agents are running out of heap memory and crashing. This can happen if the production agents are monitoring a high volume of business transactions, metrics, or snapshots, and the default heap size is not sufficient to handle the load. To resolve this issue, you can enable the heap-storage-monitor-enabled setting on the agent, which allows the agent to monitor its own heap usage and automatically reduce the amount of data it collects when the heap usage reaches a certain threshold. This can help to prevent the agent from running out of memory and crashing. [You can also adjust the heap-storage-monitor-threshold setting to specify the percentage of heap usage that triggers the data reduction¹. Alternatively, you can also increase the heap space for the agent by using the -Xmx option in the agent startup script, but this may require more resources from the host machine².](#)

The other options are not valid solutions to resolve this issue. Enabling agent auto-restart in the controller UI will not

prevent the agent from crashing due to heap exhaustion, but only restart the agent after it crashes. [This may cause data loss and performance degradation](#)³. Decreasing the worker thread count will not affect the heap usage of the agent, but only limit the number of concurrent tasks that the agent can execute. This may reduce the agent throughput and responsiveness.

<https://docs.appdynamics.com/display/PRMY22/Java+Agent>

<https://docs.appdynamics.com/appd/23.x/latest/en/application-security-monitoring/cisco-secure-application-requirements>

Question: 57

By default, which two Sensitive Data Filters substring does the Java Agent enable? (Choose two.)

- A. substring "ssn"
- B. substring "credit card"
- C. substring "key"
- D. substring "account"
- E. substring "password"

Answer: CE

Explanation:

By default, the Java Agent enables two Sensitive Data Filters substring: substring "key" and substring "password". These filters prevent the agent from capturing and sending any data that contains these substrings to the Controller, such as query parameters, HTTP headers, cookies, environment variables, or system properties. This helps to protect sensitive information from being exposed in the Controller UI or reports. [You can also add more filters or modify the existing ones by editing the agent configuration file](#)¹ Reference: [1: Filter Sensitive Data](#)

Question: 58

After installing the standalone machine agent, how do you enable the agent to start collecting metrics?

- A. The machine agent requires an app agent to collect metrics
- B. The machine agent runs with the monitored application
- C. The machine agent starts collecting metrics as soon as it is installed

D. The machine agent must be started and running

Answer: D

Explanation:

[According to the Install the Standalone Machine Agent document1](#), the machine agent is a standalone Java program that collects performance statistics about your environment. It can be deployed on any machine that hosts application servers, database servers, messaging servers, Web servers, etc. However, the machine agent does not start collecting metrics as soon as it is installed. You need to start the machine agent manually or configure it to run as a service or a daemon. Therefore, the correct answer is D. The machine agent does not require an app agent to collect metrics, nor does it run with the monitored application. These are false statements. Reference:

[Install the Standalone Machine Agent](#)

Question: 59

Which two would be marked as an exception in AppDynamics? (Choose two)

- A. HTTP errors that do not occur in the context of a Business Transaction
- B. a custom-configured error method and error message
- C. any message that is logged with a severity of Error or Fatal
- D. a handled exception after the business transaction terminates

Answer: BC

Explanation:

An exception in AppDynamics is a code-level event that indicates a problem or abnormality in the execution of a business transaction. AppDynamics captures and presents information on exceptions in the monitored environment, and allows you to configure the rules and criteria for detecting and reporting exceptions.

According to the [Errors and Exceptions - AppDynamics](#), the following two types of events would be marked as an exception in AppDynamics:

A custom-configured error method and error message: This is a type of event that you can define as an exception based on a specific method and message that indicate an error condition in your application code. You can configure the error

detection rules to mark a business transaction as an error and capture the exception information when such a method and message are encountered. For example, you can configure a rule to mark a business transaction as an error and capture the exception when the method com.example.MyClass.myMethod returns the message Invalid input.

Any message that is logged with a severity of Error or Fatal: This is a type of event that AppDynamics automatically detects as an exception based on the log messages that are generated by your application code. AppDynamics scans the log messages for the keywords Error or Fatal, and marks the business transaction as an error and captures the exception information when such a message is found. For example, AppDynamics will mark a business transaction as an error and capture the exception when the log message contains Error: Null pointer exception.

Therefore, B (a custom-configured error method and error message) and C (any message that is logged with a severity of Error or Fatal) are two types of events that would be marked as an exception in AppDynamics. Reference:

[Cisco AppDynamics Associate Administrator Certification](#)

[Errors and Exceptions - AppDynamics](#)

[Error Detection - AppDynamics](#)

Question: 60

A customer wants a dashboard that will show them the number of current database connections on their application in a Timeseries graph, and compare it to past averages for the same time What option would solve this the fastest for the customer?

- A. Create another dashboard and put them side by side.
- B. Create a report that shows the historical data they are looking for
- C. Include baseline data within widgets on the dashboard
- D. Open up the application and change the time range to show the time they want

Answer: C

Explanation:

The option that would solve this the fastest for the customer is to include baseline data within widgets on the dashboard.

Baseline data is the data that represents the normal or expected behavior of a metric, based on the historical data collected by the AppDynamics Cognition Engine. By including baseline data within widgets on the dashboard, the customer can easily compare the current value of the metric with the past average value for the same time. For example,

if the customer wants to see the number of current database connections on their application in a Timeseries graph, and compare it to the past averages for the same time, they can create a Timeseries graph widget that shows the metric “Database Connections” and enable the option “Show Baseline Data” in the widget settings. [This will display the current and the baseline values of the metric in the same graph, and allow the customer to see the trends and deviations over time](#)¹.

Question: 61

Unless archived, what is the default length of time events can be viewed with full detail before being purged?

- A. 7 days
- B. 14 days
- C. 21 days
- D. 1 month

Answer: A

Explanation:

The default length of time events can be viewed with full detail before being purged is 7 days, unless archived. Events are records of significant occurrences in the monitored environment, such as errors, exceptions, policy violations, health rule violations, and more. You can view the events in the Controller UI under Events > All Events. [You can also archive the events to retain them for longer periods or export them to external systems](#)¹ Reference: [1: Events](#)