Dynatrace Associate Certification Practice Exam (Sample)

Study Guide

Everything you need from our exam experts!

Questions

1. **For how long is 1-minute interval granularity request history available for a key request?**

   A. 7 days

   B. 14 days

   C. 30 days

   D. 60 days

2. **What type of data is sent to mission control?**

   A. Server logs

   B. Usage and billing information

   C. User credentials

   D. Backup files

3. **How frequently is issue detection reset in Dynatrace?**

   A. Every month

   B. Every week

   C. Every day

   D. Every hour

4. **What purpose does the retention period serve for OneAgent Diagnostics?**

   A. It determines how long the logs are visible

   B. It defines the amount of storage allocated

   C. It sets guidelines for data privacy

   D. It impacts performance analytics accuracy

5. **In Dynatrace, what is considered an "Application"?**

   A. Tools used for application deployment

   B. Logical constructs onto which customer applications are mapped for monitoring

   C. Only web-based applications

   D. Server configurations for application hosting

6. **How can you change the host group of a host?**

    A. By deleting the host

    B. By re-installing the agent

    C. By using the control panel

    D. By system auto-update

7. **How can Real User Monitoring (RUM) function on an unmonitored host?**

    A. Using a dedicated VPN

    B. Through manual code integrations

    C. Agentless and RUM browser extension

    D. By installing a monitoring cap

8. **Which feature allows users to customize viewing frequency in dashboards?**

    A. Timeframe Selector

    B. Edit Tile Options

    C. Dashboard Filters

    D. User Preferences

9. **In order to maintain functionality, a cluster must have what minimum number of nodes if it can't lose 3 or more nodes?**

    A. (# of nodes / 2) + 1

    B. (# of nodes / 3) + 1

    C. (# of nodes x 2) - 1

    D. # of nodes - 3

10. **What is the command to launch Dynatrace's CLI tool?**

    A. ./dynatracectl

    B. ./oneagentctl

    C. ./agentctl

    D. ./cli_tool

Answers

SAMPLE

1. B
2. B
3. B
4. A
5. B
6. B
7. C
8. B
9. A
10. B

Explanations

1. For how long is 1-minute interval granularity request history available for a key request?

A. 7 days

B. 14 days

C. 30 days

D. 60 days

The availability of 1-minute interval granularity request history for a key request extends for 14 days. This duration is important as it allows users to analyze request performance and behavior with high resolution, enabling them to identify trends, diagnose issues, and optimize performance over a significant time frame. The 14-day window is designed to balance the need for detailed historical data with system performance and storage considerations. This granularity is crucial for monitoring and troubleshooting applications, as it provides insights into the requests' characteristics during a defined time period, thereby facilitating more effective decision-making based on this analysis.

2. What type of data is sent to mission control?

A. Server logs

B. Usage and billing information

C. User credentials

D. Backup files

The correct choice is usage and billing information. This data is critical for the operation of Dynatrace's services, as it allows the company to track resource utilization and manage subscriptions effectively. This information ensures that customers are billed accurately based on their usage, and it helps Dynatrace analyze service performance and resource demands across their user base. Usage and billing details are essential for providing insights into how products are being utilized, which can inform future improvements and feature developments. By sending this type of information to mission control, Dynatrace can optimize its offerings and ensure customers receive appropriate support based on their usage patterns. Other types of data like server logs, user credentials, and backup files are typically sensitive and would not be directed to mission control for broader analysis and operational purposes. Server logs are often used for diagnostics and troubleshooting, user credentials are kept private for security reasons, and backup files are used to restore data rather than being shared for operational assessments.

3. How frequently is issue detection reset in Dynatrace?

A. Every month

B. Every week

C. Every day

D. Every hour

The frequency of issue detection reset in Dynatrace is typically set to occur every week. This means that any issues that have been detected and reported will be reviewed and potentially reset at the end of each week, helping to maintain a clear and manageable set of current issues that require attention. This regular reset ensures that the issue monitoring remains relevant and up to date, reflecting changes in application performance, infrastructure, and user behavior. This weekly cycle allows teams to periodically reassess the status of detected issues, prioritize them according to their impact on the overall health of the system, and address them as necessary. Allowing detection to reset too frequently, such as daily or hourly, could result in noise and confusion, making it harder to focus on genuine, persistent issues that need resolution. By using a weekly reset, Dynatrace helps organizations maintain effectiveness in monitoring and managing issues without overwhelming users with constant changes.

4. What purpose does the retention period serve for OneAgent Diagnostics?

A. It determines how long the logs are visible

B. It defines the amount of storage allocated

C. It sets guidelines for data privacy

D. It impacts performance analytics accuracy

The retention period for OneAgent Diagnostics primarily determines how long the logs generated by the OneAgent will be visible and accessible for analysis. This is critical for organizations to manage their log data efficiently and ensure that they can troubleshoot issues, analyze performance metrics, and maintain the necessary compliance with data retention policies. By setting a defined retention period, users can maintain control over log data, ensuring that they have the necessary information available for a certain timeframe while also managing storage resources effectively. As log data becomes less relevant over time, an appropriately set retention period helps avoid unnecessary consumption of storage space, and it can also contribute to maintaining the overall performance of the monitoring system. Properly managing log visibility through the retention period allows teams to balance the need for historical data analysis with considerations for data management and compliance standards.

5. In Dynatrace, what is considered an "Application"?

A. Tools used for application deployment

B. Logical constructs onto which customer applications are mapped for monitoring

C. Only web-based applications

D. Server configurations for application hosting

In Dynatrace, an "Application" is defined as a logical construct onto which customer applications are mapped for monitoring. This concept is central to Dynatrace's ability to provide insights and performance analytics for applications within environments. By mapping customer applications to this framework, Dynatrace can collect performance data, user experience metrics, and other relevant information that aids in monitoring the health and efficiency of the application in a real-time context. Mapping applications in this way allows Dynatrace to correlate data from various services, processes, and infrastructure components, offering a comprehensive view of application performance. This logical mapping is essential for understanding how different pieces of an application interact and perform, enabling better troubleshooting and optimization strategies. The other options focus on areas that do not represent the core definition of an application within Dynatrace. Tools for deployment, web-based applications only, and server configurations do not capture the essence of how Dynatrace conceptualizes and monitors applications. The emphasis is on creating a logical representation that enables effective monitoring and performance analysis, rather than on the technical aspects of deploying or hosting applications.

6. How can you change the host group of a host?

A. By deleting the host

B. By re-installing the agent

C. By using the control panel

D. By system auto-update

The ability to change the host group of a host is typically achieved by reinstalling the agent on that host. This process is essential because the configuration settings, including the host group assignment, are determined during the installation of the monitoring agent. By reinstalling the agent, you can specify a different host group in the configuration settings, allowing you to effectively move the host to a new group. It's important to note that simply deleting the host would remove it from monitoring altogether, and thus would not change its group assignment. Using the control panel for such modifications might not offer the capability to change host groups directly, as this usually requires agent-level configuration adjustments. An automatic system update generally pertains to software updates or patches, not to the reassignment of hosts to different groups.

7. How can Real User Monitoring (RUM) function on an unmonitored host?

A. Using a dedicated VPN

B. Through manual code integrations

C. Agentless and RUM browser extension

D. By installing a monitoring cap

Real User Monitoring (RUM) can function on an unmonitored host through an agentless approach, specifically by utilizing RUM browser extensions. These extensions capture user interactions directly from the user's browser without needing any pre-installed agents or configurations on the host machine. This capability allows for monitoring and collecting performance data as the application is used, providing insights into the real user experience regardless of the underlying infrastructure or environment. RUM browser extensions are designed to gather key performance metrics, interaction times, and other relevant data that can be crucial for understanding user behavior and application performance. Since they operate in the user's browser, they can effectively bypass any constraints posed by the lack of installed monitoring agents, making them a flexible solution for monitoring across various environments, including unmonitored hosts.

8. Which feature allows users to customize viewing frequency in dashboards?

A. Timeframe Selector

B. Edit Tile Options

C. Dashboard Filters

D. User Preferences

The feature that enables users to customize viewing frequency in dashboards is the Edit Tile Options. This functionality allows users to modify how data is displayed within a specific tile, including settings related to refresh intervals. By accessing the Edit Tile Options, users can set the desired frequency for data refreshes, allowing them to view the most current information or maintain a particular viewing cadence as per their needs. In this context, the other options do provide valuable functions for managing dashboards, but they do not specifically address the customization of viewing frequency. The Timeframe Selector generally allows users to define which time period they want to view data for, rather than setting how often that data refreshes. Dashboard Filters are used to filter the displayed data based on specific criteria but do not control the refresh frequency. User Preferences may contain various personal configurations, but they typically don't directly affect the frequency at which the data in individual tiles updates.

9. In order to maintain functionality, a cluster must have what minimum number of nodes if it can't lose 3 or more nodes?

A. (# of nodes / 2) + 1

B. (# of nodes / 3) + 1

C. (# of nodes x 2) - 1

D. # of nodes - 3

To determine the minimum number of nodes required for a cluster to maintain its functionality without losing 3 or more nodes, the correct answer is rooted in the concept of redundancy and fault tolerance within clustered systems. When a cluster cannot afford to lose 3 nodes, it's essential to ensure that there is still a majority of nodes available to make decisions and process requests. In general, a cluster operates effectively as long as a majority of its nodes are functioning properly. The formula used to calculate the minimum number of nodes needed in a system capable of withstanding failures involves the requirement to have more than half of the total nodes operational. The provided choice, which states that a cluster must consist of (# of nodes / 2) + 1, aligns precisely with this principle. This formula ensures that, even in the event of losing a certain number of nodes (in this case, up to 3), a majority will still remain, thus maintaining the cluster's functionality. For example, if you have a cluster of 6 nodes, losing up to 3 means you could end up with only 3 nodes remaining. To maintain functionality, you need at least 4 nodes (which is more than half of 6). Applying the formula (# of nodes

10. What is the command to launch Dynatrace's CLI tool?

A. ./dynatracectl

B. ./oneagentctl

C. ./agentctl

D. ./cli_tool

The command to launch Dynatrace's CLI tool is correctly identified as ./oneagentctl. This tool is specifically designed to interact with Dynatrace OneAgent, which is responsible for collecting performance metrics and monitoring application environments. The CLI allows users to manage the OneAgent configuration, perform diagnostics, and troubleshoot issues effectively. Understanding the role of this command within the Dynatrace ecosystem is essential. The oneagentctl command provides various functionalities that help in managing the OneAgent behavior, such as starting, stopping, or updating the OneAgent. This is particularly useful in environments where automated monitoring and management of applications are important. The other options do not reflect the correct command for launching the CLI tool related to Dynatrace's OneAgent. Each has a different context and function, which doesn't align with the OneAgent management utilities provided by Dynatrace. Familiarity with these commands helps in navigating the system efficiently.