DUMPSBOSS.

Dynatrace Associate Certification Exam

Dynatrace Dynatrace-Associate

Version Demo

Total Demo Questions: 9

Total Premium Questions: 94

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Dynatrace Platform	17
Topic 2, Deployment	14
Topic 3, Configuration	29
Topic 4, Monitoring	18
Topic 5, Analysis	16
Total	94

QUESTION NO: 1

Which metrics are presented within the Databases app? (Select all that apply)

A. Code-level insights

B. User Calls

C. Database vulnerabilities

D. Memory Consumption

E. Active Sessions

ANSWER: D E

Explanation:

Memory Consumption and **Active Sessions** are database-health and workload metrics that can be presented through Dynatrace database monitoring. Memory consumption helps teams understand the database process's resource usage and identify conditions in which memory pressure may affect database stability or performance. Active sessions reflects concurrent database activity, helping operators assess current workload, connection demand, and potential contention. Together, these metrics provide useful operational context: an increase in session activity can be correlated with resource consumption and performance behavior to investigate database load. Dynatrace collects and presents database-related observability data from monitored database technologies and their supporting processes, allowing teams to analyze database behavior alongside the wider environment. The precise metric availability and dimensions depend on the database technology, monitoring configuration, and whether the relevant database extension or built-in monitoring capability supplies the metric. Dynatrace's database monitoring documentation describes the supported monitoring approach and the database observability data available in the platform. See Dynatrace database monitoring documentation and Dynatrace Extensions 2.0 documentation.

QUESTION NO: 2

Select the topology levels that support the viewing of Problems. (Select all that apply)

A. Synthetics

B. Process

C. Host

D. Application

E. Service

F. Data Center

ANSWER: B C D E F

Explanation:

Dynatrace Smartscape represents monitored environments through a layered topology that includes data centers, hosts, processes, services, and applications. Each of these entity layers provides context for detected issues and enables users to inspect associated problem information. Consequently, **Process**, **Host**, **Application**, **Service**, and **Data Center** all support viewing Problems. Selecting an entity at one of these topology levels lets users understand its health state, identify impacted components, and navigate to correlated Davis AI problem details. This topology-driven investigation is central to Dynatrace: Davis AI evaluates dependencies between infrastructure and application entities so that a problem can be viewed from the relevant layer of the environment rather than as an isolated event. A data center is also a Smartscape topology level, grouping hosts and their dependent entities; its status and related problem context can therefore be examined from the data-center perspective. Dynatrace documents Smartscape as the topology visualization for relationships among data centers, hosts, processes, services, and applications. See Dynatrace Smartscape documentation and Davis AI problem detection documentation.

QUESTION NO: 3

I want to reduce the mean-time-to-identify critical security vulnerabilities in my environment by sending a notification to the correct team's Slack channel when detected. How can I achieve this?

A. Application Security app

B. Manually check the Security Vulnerability page

C. Workflow with an Event Trigger, Ownership and Slack integration

D. Manually check the Problems page

ANSWER: C

Explanation:

Workflow with an Event Trigger, Ownership and Slack integration is the appropriate solution because it automates the complete notification path from detection through delivery to the responsible team. A workflow can start when Dynatrace emits the relevant security event, allowing critical vulnerabilities to be handled as soon as they are detected rather than waiting for a person to review a page. Ownership information supplies the contextual routing needed to associate the affected entity or application with its accountable team. The workflow can then use the Slack integration to post the notification to that team's designated channel, including relevant vulnerability and affected-service context. This creates a consistent, actionable alerting process that shortens the time between a vulnerability being identified and the team that can triage or remediate it being informed. It also scales across many services and teams without requiring separate manual monitoring routines. Dynatrace documents workflow triggers as the mechanism for starting automated workflows and provides Slack actions for sending messages from those workflows. See the Dynatrace workflow triggers documentation and the Dynatrace Slack workflow action documentation.

QUESTION NO: 4

What types of notifications can you set up for Security Problems? (Select all that apply)

A. Jira

B. Custom Integrations

C. Emails

D. PagerDuty

E. ServiceNow

ANSWER: A B C D E

Explanation:

Dynatrace Security notifications can route security-problem information to several operational and collaboration channels, so all listed notification types apply. Email notifications provide direct delivery to recipients who need awareness of newly detected or updated security problems. Jira and ServiceNow integrations support ticket-based response processes by creating and updating work items in the team's established issue-management or ITSM platform. PagerDuty can be used to engage on-call incident responders when a security problem requires immediate attention. Custom integrations are supported through webhook-based notification delivery, enabling an organization to send Dynatrace security-problem events to its own automation, chat, SIEM, or incident-response tooling. These notification configurations help ensure that vulnerability and runtime security findings reach the appropriate teams and can be incorporated into existing remediation workflows. Dynatrace documents security notifications as supporting integrations such as email, Jira, ServiceNow, PagerDuty, and webhooks. See the Dynatrace Security notifications documentation and the Dynatrace workflow actions and integrations documentation.

QUESTION NO: 5

An API must be tested synthetically with a required request header and JSON payload. The monitor must also fail when the endpoint returns an unexpected HTTP status. Which three HTTP monitor configuration capabilities address these requirements? (Select three)

A. Request headers

B. Request body

C. Response code validation

D. User session properties

E. Kubernetes labels

F. Browser clickpath actions

ANSWER: A B C

Explanation:

An HTTP monitor can be configured with **request headers**, a **request body**, and **response code validation**. These settings allow the monitor to send the required header and JSON payload, then verify that the endpoint returns an expected HTTP status.

User session properties belong to Real User Monitoring, Kubernetes labels provide workload metadata, and browser clickpath actions are used for browser-based multi-step journeys rather than constructing an HTTP monitor request.

QUESTION NO: 6

Which language is required to be used to query Business Event data?

A. Grail

B. SQL

C. DQL

D. USQL

ANSWER: C

Explanation:

DQL is the correct language for querying Business Event data in Dynatrace. Dynatrace Query Language is the unified query language used with Grail, the Dynatrace data lakehouse, to retrieve, filter, transform, aggregate, and analyze observability data. Business events are ingested into Grail and can be queried from Dynatrace using DQL, including through the Logs & Events viewer, notebooks, dashboards, workflows, and API-driven query use cases. A typical query starts by selecting the business-event data source with a command such as `fetch bizevents`, followed by pipeline commands that narrow and analyze the returned records. This approach lets teams correlate business activity, such as purchases or customer interactions, with technical observability context held in Grail. DQL is intentionally designed around this pipeline-based querying model and is the documented Dynatrace mechanism for accessing business event records. See the Dynatrace documentation on Dynatrace Query Language and the guide to Business events for the supported data model and querying workflow.

QUESTION NO: 7

Which capability allows me to integrate Synthetic monitors into my CI/CD pipeline to automate the execution?

A. HTTP monitor

B. Private Synthetic location

C. On-demand execution

D. Credential vault

ANSWER: C

Explanation:

On-demand execution is the capability designed for programmatically triggering a Synthetic monitor outside of its regular schedule. Dynatrace provides an API endpoint that starts an on-demand execution for a configured Synthetic monitor, allowing a CI/CD system to invoke the monitor at a defined pipeline stage, such as after deploying to a test, staging, or production environment.

This supports automated release validation: the pipeline can start the monitor, obtain the execution result, and use the result as a quality signal before continuing or promoting a build. The monitor itself retains its configured test steps, locations, frequency settings, and alerting configuration, while on-demand execution supplies the explicit, API-driven trigger needed by automation workflows. This makes it useful for validating availability, key user journeys, and externally observable performance immediately following a deployment rather than waiting for the next scheduled run.

Dynatrace documents the on-demand execution endpoint and its request requirements in the Synthetic on-demand execution API documentation. For broader configuration and operation of Synthetic monitors, see the Dynatrace Synthetic Monitoring documentation.

QUESTION NO: 8

I need to take my live environment offline while I perform an upgrade. However, the upgrade will overlap with scheduled Synthetic tests. What can I configure in Dynatrace to avoid false alerting?

A. Maintenance Window

B. Credential vault

C. Synthetic On-demand execution

D. Delete the Synthetic monitors

ANSWER: A

Explanation:

Maintenance Window is the appropriate Dynatrace configuration for planned upgrade downtime that overlaps with scheduled Synthetic monitoring. A maintenance window defines a time period and scope of monitored entities for which Dynatrace suppresses problem creation and alert notifications. By applying the window to the environment or the affected monitored entities while the upgrade is performed, scheduled Synthetic tests can continue to run but their expected failures do not generate misleading problems or notifications. This lets teams retain the monitor configuration and schedule, while avoiding alert noise and unnecessary incident-response activity during an approved maintenance period. Maintenance windows can be configured as one-time or recurring schedules, making them suitable for both a single upgrade and regularly planned maintenance. After the configured window ends, normal problem detection and alerting resume automatically. Dynatrace documents maintenance windows as the mechanism for suppressing alerts during planned maintenance and describes their scope and scheduling behavior in the Maintenance windows documentation. For the relationship between Synthetic monitoring and alerting/problem detection, see the Synthetic monitoring documentation.

QUESTION NO: 9 - (SIMULATION)

Welcome to the Practical section of your exam!

This portion of the exam requires you to configure and analyze various components within Dynatrace and is open book using Dynatrace documentation ONLY by opening a new tab. The Dynatrace Operator is already deployed for you in a Kubernetes environment (Linux host) in Cloud-Native Full-Stack. The easyTrade and easyTravel applications are also already deployed

for you. You will most likely only interact with the Dynatrace Environment (SaaS Platform). For a majority of these questions: paste or attach your screen captures directly to the question using the image control.

Information for connecting to exam components is provided in Dynatrace University in your event classroom. Open the Associate Exam Event / Classroom and go to the Environments tab, https://universitv.dvnatrace.com/account/events

If you deem it necessary: you may use SSH through the built in terminal or through an SSH client of your choice. Use the IP address and credentials provided to log in to the Linux virtual machine running your Kubernetes cluster. When you SSH to the public IP and are presented with the credential prompt, enter the provided username and then the provided password.

Review these and additional instructions for connecting to your environment and creating screenshots by clicking 'View all questions' and then the 'Info' icon.

easyTravel Information

easyTravel is a multi-tier application implemented in .Net and Java. It provides a web portal which allows users to log in, search for journeys to various destinations, select promotional journeys that are offered, and to book a journey using credit card details. Additionally, a Business-to-Business (B2B) web portal for travel agencies is provided where travel agencies can manage the journeys that they offer and can review reports about created bookings.

easyTrade Information

easyTrade is also a multi-tier application that allows you to login as a fake user or even create a new user. After which point, you are able to buy and sell stocks or deposit money into your balance using fake information. These interactions will create transactions in Dynatrace for you to analyze. However, there is a load generator running which means that interacting with this page should not be necessary to answer the questions.

Tips for completing tasks:

1. Read each task carefully and fully before starting, and then follow the instructions to complete each task.

2. Be sure your screen captures match the requirements stated and are readable.

3. To create your screen captures, you can zoom your screen, or take multiple screenshots to show all the areas requested. You can attach multiple images to the same question. It is recommended to provide more proof than necessary and document your steps. When in doubt, the old saying "show your work", comes in handy.

4. Use a screen capture tool to create and then copy to your clipboard and paste screen captures directly into the file upload control. To upload a file, allowed file types for the upload are: Jpg, .jpeg, and .png.

5. When saving your images for upload into the test be sure that your image titles do not contain the following characters:!,-,_,..",',(,).

6. To access the web portal, use the links found in the SSH terminal. They will be listed as soon as you login.

Tasks1:

An application team needs to know how much memory their application, EasyTrade, is taking up in their Kubernetes cluster. They have to see the memory usage for all the pods running in the "easytrade' namespace, summed up into a single metric. Use a Dashboard in the new Platform interface (not Dashboards Classic) to show this information to the team and set the timeframe for this data to look at the last 30 minutes. Call this dashboard "EasyTrade memory usage".

Screen capture(s) required:

Upload image(s) that show the following:

1. A dashboard "EasyTrade memory usage".

2. The tile used to display memory usage for all pods in the "easytrade" namespace as well as its configuration.

ANSWER: See the explanation for the answer

Explanation:

Create the dashboard in the new Dynatrace Platform Dashboards app (do not use Dashboards Classic) with the name `EasyTrade memory usage`.

1. Create a new dashboard and set its title to `EasyTrade memory usage`.
2. Set the dashboard time frame selector to **Last 30 minutes**.
3. Add a Data Explorer/chart tile and select the metric **Kubernetes container memory usage**, metric key `builtin:kubernetes.container.memory_usage`.
4. Add a metric filter for **Kubernetes namespace name** / `k8s.namespace.name` equal to `easytrade`.
5. Set the aggregation to **Sum**.
6. Do not configure any **Split by** dimension. In particular, do not split by pod, container, workload, or namespace. This produces one aggregate memory-usage series/value across the containers of every pod in the `easytrade` namespace.
7. Save the tile and dashboard. A useful tile title is `EasyTrade pods total memory usage`.

The effective metric configuration is:

```
Metric: builtin:kubernetes.container.memory_usage
Filter: k8s.namespace.name = easytrade
Aggregation: Sum
Split by: none
Timeframe: Last 30 minutes
```

For the required evidence, capture the dashboard showing its name and the **Last 30 minutes** selector, plus the tile's edit/configuration view showing the selected memory metric, namespace filter, **Sum** aggregation, and no split-by dimension.