

Zabbix Certified Specialist Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

SAMPLE

- 1. At what levels can macros be defined in Zabbix, and what is their precedence order?**
 - A. 1. Global 2. Host 3. Template**
 - B. 1. Host 2. Template 3. Global**
 - C. 1. Template 2. Host 3. Global**
 - D. 1. Host 2. Global 3. Template**
- 2. What needs to be constructed before deploying Zabbix?**
 - A. Database with InnoDB engine**
 - B. Database with Unicode character set and UTF8 national character set**
 - C. Database on cloud storage**
 - D. Database with MyISAM storage engine**
- 3. What two options exist for installing Zabbix using Docker containers?**
 - A. Zabbix Docker Hub or GitHub**
 - B. Zabbix GitHub or Bitbucket**
 - C. Zabbix Hub or Docker Compose**
 - D. Zabbix Azure or AWS**
- 4. What role does the "discovery rule" serve in Zabbix?**
 - A. Manages user access and permissions.**
 - B. Automatically identifies and manages networked devices and their items.**
 - C. Generates performance alerts based on predefined thresholds.**
 - D. Creates backup schedules for system data.**
- 5. Which feature allows for automated resource monitoring in Zabbix?**
 - A. Event triggers.**
 - B. Low-Level Discovery (LLD).**
 - C. User access settings.**
 - D. Notification mediums.**

- 6. What package is essential for installing MySQL on a Zabbix server?**
- A. zabbix-server-mysql**
 - B. mysql-server**
 - C. zabbix-agent-mysql**
 - D. mysql-client**
- 7. In Zabbix, if an item's key is not unique, what is the consequence?**
- A. The item cannot be created**
 - B. Data for that item may be overwritten**
 - C. The value will default to Zero**
 - D. The host will stop monitoring**
- 8. How can you identify which tabs have settings enabled when editing a host?**
- A. Red dot**
 - B. Green dot**
 - C. Blue dot**
 - D. Yellow dot**
- 9. How frequently is data collected via the Zabbix proxy sent to the Zabbix server by default?**
- A. Every second**
 - B. Every minute**
 - C. Every five minutes**
 - D. Once every hour**
- 10. What command enables the Zabbix server service to start on every boot?**
- A. systemctl enable zabbix-server --now**
 - B. systemctl start zabbix-server --enable**
 - C. systemctl enable zabbix-server**
 - D. start zabbix-server on boot**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. At what levels can macros be defined in Zabbix, and what is their precedence order?

- A. 1. Global 2. Host 3. Template
- B. 1. Host 2. Template 3. Global**
- C. 1. Template 2. Host 3. Global
- D. 1. Host 2. Global 3. Template

Macros in Zabbix are defined at different levels, and understanding their precedence is essential for effective monitoring setup. The correct approach to this involves recognizing that macros can be defined at the global level, the template level, and the host level, each with specific characteristics and scopes. When a macro is defined at multiple levels, the order of precedence determines which macro value will be utilized when referenced. The correct precedence order is that host-level macros take priority over template-level macros, which in turn take precedence over global macros. This means if a host has a specific value set for a macro, that value will be used instead of any corresponding value defined at the template or global levels. This hierarchy allows for flexible customization of monitoring configurations. For instance, if a template is used for multiple hosts, the global macro can provide a default value, while individual hosts can override this with their specific macro value as necessary, ensuring that monitoring remains relevant and tailored to the particular needs of each host. Understanding this precedence is crucial for Zabbix users to avoid conflicts and ensure that the correct values are applied in their monitoring configurations.

2. What needs to be constructed before deploying Zabbix?

- A. Database with InnoDB engine
- B. Database with Unicode character set and UTF8 national character set**
- C. Database on cloud storage
- D. Database with MyISAM storage engine

To successfully deploy Zabbix, it is essential to set up a database that utilizes a Unicode character set along with the UTF8 national character set. This configuration is crucial because Zabbix may need to handle diverse data, including international characters and various encoding scenarios. By using a Unicode-compatible character set, you ensure that the database can correctly store and retrieve any type of textual data, which is particularly important for environments that may include various language inputs or special characters. Implementing the UTF8 national character set further enhances compatibility, especially for non-Latin characters that might be used in monitoring environments across different regions. This foundational setup elevates data integrity and consistency across the platform, enabling Zabbix to operate smoothly without issues related to encoding. In contrast, other database configurations, such as those focused on specific storage engines or cloud environments, do not directly address the critical aspect of character encoding support required for Zabbix. Establishing a robust character set framework takes precedence over considerations of storage engines or cloud storage options in ensuring optimal performance and compatibility.

3. What two options exist for installing Zabbix using Docker containers?

- A. Zabbix Docker Hub or GitHub**
- B. Zabbix GitHub or Bitbucket
- C. Zabbix Hub or Docker Compose
- D. Zabbix Azure or AWS

Zabbix can indeed be installed using Docker containers from the official Zabbix Docker Hub, which is a repository that contains pre-built images specifically designed for running Zabbix services effectively in a containerized environment. This method simplifies the deployment process, making it easier for users to set up Zabbix without dealing with complex configurations and dependencies. Docker Hub provides a standardized way to pull the necessary images, ensuring that users are working with the latest and most stable versions of Zabbix. The mention of GitHub, while it hosts the source code for Zabbix and many other projects, is not a direct method for installing Zabbix via Docker, as it does not provide pre-built container images. Instead, users may need to build the containers from the source manually, which adds extra steps to the installation process. By utilizing Docker Hub, users can streamline the deployment, allowing for a more efficient installation experience. In summary, the correct answer identifies the official method for obtaining Zabbix Docker images from Docker Hub, while also alluding to the source code repositories for those interested in building or customizing their own installations.

4. What role does the "discovery rule" serve in Zabbix?

- A. Manages user access and permissions.
- B. Automatically identifies and manages networked devices and their items.**
- C. Generates performance alerts based on predefined thresholds.
- D. Creates backup schedules for system data.

The discovery rule in Zabbix plays a crucial role in automatically identifying and managing networked devices and their associated items. This functionality allows Zabbix to dynamically discover hosts and determine their configurations and available metrics without the need for manual intervention. When a discovery rule is defined, it can scan a defined range of IP addresses or specific networks, detect devices, and then auto-provision items related to those devices such as metrics for CPU load, memory usage, and network traffic. This automated process helps maintain an accurate and up-to-date inventory of devices within the monitored environment, ensuring that Zabbix is always aware of the current state of the network. In contrast, user access and permissions management, performance alert generation, and data backup scheduling are facets of Zabbix that relate to operational control and data management but do not pertain directly to the discovery functionality provided by discovery rules. Thus, the specific capacity of the discovery rule to streamline the monitoring process by automatically discovering network devices and configuring them as necessary is what makes this option the correct answer.

5. Which feature allows for automated resource monitoring in Zabbix?

- A. Event triggers.**
- B. Low-Level Discovery (LLD).**
- C. User access settings.**
- D. Notification mediums.**

Low-Level Discovery (LLD) is the feature in Zabbix that enables automated resource monitoring. It functions by automatically discovering items, triggers, and graphs based on predefined rules. This means that if new resources are added to the monitored environment—such as new file systems, network interfaces, or processes—LLD can detect these changes and create the necessary monitoring items without manual intervention from the administrator. This capability significantly reduces the operational overhead associated with manually configuring new resources for monitoring, thereby enhancing the efficiency and scalability of the monitoring setup. By using LLD, Zabbix can adapt to changes in the IT environment in real-time, ensuring that all relevant metrics and alerts are promptly incorporated into the monitoring framework. Other choices, such as event triggers, are important for defining when alerts should be generated based on certain conditions or metrics, while user access settings control permissions for individuals interacting with the Zabbix interface. Notification mediums relate to how alerts are communicated to users but do not contribute to the monitoring setup itself. Only Low-Level Discovery specifically facilitates the automation of resource monitoring.

6. What package is essential for installing MySQL on a Zabbix server?

- A. zabbix-server-mysql**
- B. mysql-server**
- C. zabbix-agent-mysql**
- D. mysql-client**

The choice of "zabbix-server-mysql" is correct because this package specifically includes the necessary components and configurations required for Zabbix to integrate with a MySQL database. When setting up a Zabbix server, it is essential to have this package because it provides the backend capabilities for data storage and retrieval which Zabbix requires to function properly. In contrast, other packages listed serve different purposes. The "mysql-server" package is a fundamental installation for MySQL itself, providing the server component, but it does not include the Zabbix-specific integrations or configurations that "zabbix-server-mysql" offers. The "zabbix-agent-mysql" is not a standard package related to the Zabbix server setup; rather, it pertains to the Zabbix agent that collects metrics from the monitored hosts. Lastly, the "mysql-client" package is used for interacting with a MySQL database from the command line but does not have relevance to the server setup or Zabbix's requirements for MySQL integration. Thus, selecting "zabbix-server-mysql" ensures that the Zabbix server operates efficiently with a MySQL database as its backend.

7. In Zabbix, if an item's key is not unique, what is the consequence?

- A. The item cannot be created**
- B. Data for that item may be overwritten**
- C. The value will default to Zero**
- D. The host will stop monitoring**

In Zabbix, if an item's key is not unique, the data for that item may be overwritten. This means that when multiple items share the same key, each new piece of data collected for that key will replace the previously collected data. This behavior occurs because Zabbix associates the item key with a specific metric for a specific host and does not maintain a history of past values if multiple entries with the same key exist. Having a non-unique item key can lead to loss of valuable monitoring information, as the latest incoming data will simply erase earlier values rather than allowing them to co-exist. This emphasizes the importance of defining unique item keys for each metric you wish to monitor, ensuring that each piece of data is stored and can be accessed later for analysis, reporting, or alerting. Ensuring uniqueness in item keys helps maintain the integrity of your monitoring setup and prevents loss of data, which is crucial for effective problem detection and performance monitoring in IT environments.

8. How can you identify which tabs have settings enabled when editing a host?

- A. Red dot**
- B. Green dot**
- C. Blue dot**
- D. Yellow dot**

When editing a host in Zabbix, a green dot is used to signify that specific tabs or settings are enabled. This visual indicator allows users to quickly identify which features are currently active for a particular host. The presence of the green dot means that the associated tab's settings are not only active but have also been configured. This can help in managing and organizing hosts effectively, as it provides immediate visual feedback about the status of different monitoring options, such as applications, items, triggers, and others. The other color indicators—red, blue, and yellow—typically represent different states or types of configurations, but the green dot specifically denotes enabled settings that a user can review or modify. This consistency in visual representation aids users in efficiently navigating the Zabbix interface while making configuration changes.

9. How frequently is data collected via the Zabbix proxy sent to the Zabbix server by default?

- A. Every second**
- B. Every minute**
- C. Every five minutes**
- D. Once every hour**

Data collected via the Zabbix proxy is sent to the Zabbix server every 30 seconds by default. This frequent polling allows for near real-time monitoring and ensures that the Zabbix server has up-to-date information from its proxies. The ability to have this data transmitted every 30 seconds is particularly beneficial in environments where rapid changes can occur, enabling quick detection of issues. Understanding the default values in Zabbix is crucial for optimal configuration, as these values can impact the performance and responsiveness of your monitoring setup. Adjustments to this frequency can be made in the Zabbix proxy configuration files, depending on the specific requirements of the environment.

10. What command enables the Zabbix server service to start on every boot?

- A. systemctl enable zabbix-server --now**
- B. systemctl start zabbix-server --enable**
- C. systemctl enable zabbix-server**
- D. start zabbix-server on boot**

The command that correctly enables the Zabbix server service to start on every boot is focused on using the systemctl utility, which is a central management tool for controlling the systemd system and service manager. The core function of the command "systemctl enable zabbix-server" is to create the necessary symbolic links to ensure that the Zabbix server service starts automatically as part of the bootup process. This command configures the service to be included in the default target, which is executed during system startup. The additional flag "--now" is also significant because it not only enables the service to start on boot but also immediately starts the Zabbix server service at that moment. This combination of actions is particularly useful for users who want to enable and initiate the service in one command, ensuring that the service is active right away. In contrast, commands that either do not utilize the correct syntax of systemctl or are lacking the appropriate options to achieve the desired outcome would not effectively enable the service to start on boot. Therefore, identifying such commands helps clarify why "systemctl enable zabbix-server --now" is the fitting choice for ensuring the Zabbix server service is both enabled to start on boot and started immediately.