

Cisco.300-910.v2024-07-17.q105

| | |
|---|---|
| Exam Code: | 300-910 |
| Exam Name: | Implementing DevOps Solutions and Practices using Cisco Platforms |
| Certification Provider: | Cisco |
| Free Question Number: | 105 |
| Version: | v2024-07-17 |
| # of views: | 982 |
| # of Questions views: | 1050 |
| https://www.freeqas.com/qa/Cisco/300-910/Cisco.300-910.v2024-07-17.q105.html | |

NEW QUESTION: 1

Refer to the exhibit.

```
Traceback (most recent call last):
  File "api-call.py", line 1, in <module>
    import requests
  File "/Users/devnet/venv/devops/lib/python3.7/site-packages/requests/_init_.py", line 43, in <module>
    import urllib3
ModuleNotFoundError: No Module named 'urllib3'
```

What is the reason for this error message?

- A. The required dependencies for the urllib3 module are not installed.
- B. The requests module is not installed.
- C. The required dependencies for the requests module are not installed.
- D. The site-packages directory has been corrupted.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 2

An interface on a router as a Layer 3 link has been configured to another device by updating an Ansible playbook that is executed via a CI/CD pipeline. The Ansible playbook was written to push the configuration change using the ios_config module. Which automated test validates that the interface is functioning as it should?

- A. Add a pipeline step that captures and parses the packet flow for the interface.
- B. Add a pipeline step that captures and parses the show ip interface output so that the operational state can be referenced.
- C. Automated testing is already built in because the playbook execution is successful only after Ansible validates that it applied the change.
- D. Add a pipeline step that runs the playbook a second type. If it does not modify the device on the second run, then the interface is correctly configured.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 3

Refer to the exhibit.

```
1 ...
2 >>> response = requests.get(
3 ...     f'http://mydnacenter.local/dna/intent/api/v1/network-device/ip-address/{device_ip}',
4 ...     headers = {
5 ...         'Content-type' = 'application/json'
6 ...     }
7 ... )
8
9 Traceback (most recent call last):
10  File "<stdin>", line 1, in <module>
11 NameError: name 'requests' is not defined
```

What is causing the requests code to fail?

- A. The requests coming into stdin fail because device_ip cannot be parsed.
- B. The requests library is not installed.
- C. Rython3 is not compatible with requests.
- D. The requests library is not imported.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 4

| | |
|----------|--------|
| plan | step 1 |
| optimize | step 2 |
| reduce | step 3 |
| mature | step 4 |
| track | step 5 |

Drag and drop the steps in Gartner's public cloud cost management framework from the left into the order on the right.

Answer:



NEW QUESTION: 5

How does eliminating hardcoded or default passwords help to secure an environment?

- A. helps by removing back doors in your environments
- B. helps by enforcing your password in a repository and storing it in a secure vault
- C. helps to manage passwords centrally
- D. helps penetration testing team to focus on other issues more efficiently

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 6

How long analysis systems such as Elasticsearch, Logstash, and Kibana Stack handle ingesting unstructured logs from different devices in various formats?

- A. A single, comprehensive log format is defined on the ELK Stack. All incoming logs, regardless of format, are transformed to match the comprehensive format, and only applicable fields are populated.
- B. All logs are stored in their unstructured text format, and the ELK Stack performs data analysis by intelligently parsing the logs using machine learning algorithms.
- C. All different message formats are parsed separately using custom filters, and the resulting structured data is stored for later analysis.

D. All devices that generate syslogs must use agents that process the local logs and transmit them in a specific format to the ELK Stack.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 7

What is a practice of infrastructure as code?

- A. Configure the deployment consistently
- B. Test, integrate, and deploy once a day
- C. Use multiple version control systems
- D. Document as much as possible

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 8

What is as responsibility for teams running private cloud servers that is less of a concern when running in the public cloud?

- A. cost
- B. metrics and monitoring
- C. hardware maintenance
- D. credentials management

Answer: ([SHOW ANSWER](#))

The responsibility of credentials management is less of a concern when running in the public cloud. This is because the public cloud provider typically manages the credentials and handles authentication and authorization for the cloud services. In a private cloud environment, the organization must manage the credentials themselves, as well as authentication and authorization for the cloud services.

NEW QUESTION: 9

Drag and drop the commands from the bottom onto the correct Terraform code in the exhibit to push a network object to a Cisco ASA Firewall device.

```

[ ] "ciscoasa" {
  api_url    = "https://10.1.1.1"
  username   = "admin"
  password   = "cisco"
  ssl_no_verify = false
}

[ciscoasa_network_object] "ipv4host" {
  name = "devops_host"
  value = "10.2.3.4"
}

```

task

role

provider

module

firewall

resource

Answer:

```

[provider] "ciscoasa" {
  api_url    = "https://10.1.1.1"
  username   = "admin"
  password   = "cisco"
  ssl_no_verify = false
}

[resource] "ciscoasa_network_object" "ipv4host" {
  name = "devops_host"
  value = "10.2.3.4"
}

```

task

role

provider

module

firewall

resource

```
provider "ciscoasa" {
  api_url = "https://10.1.1.1"
  username = "admin"
  password = "cisco"
  ssl_no_verify = false
}

resource "ciscoasa_network_object" "ipv4host" {
  name = "devops_host"
  value = "10.2.3.4"
}
```



NEW QUESTION: 10

When DevOps practices are integrated into an existing organization, which two characteristics are positive indicators of DevOps maturity? (Choose two.)

- A. mean time to recover
- B. cone testing
- C. mean time between success
- D. age of codebase
- E. change lead time

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 11

A company has a CI/CD pipeline implemented to automate application deployment. An urgent change is required in the production environment. A developer produces and tests the required change in the company's environment, but the pipeline fails when pushed. Which action conforms to the rules of a CI/CD pipeline and enables the pipeline to pass?


- A. Remove the failing test.
- B. Manually deploy the change into the production environment
- C. Use a backup pipeline
- D. Replicate the production environment in the pipeline before applying

Answer: D ([LEAVE A REPLY](#))

This ensures that the change is applied in the same environment that it was tested in, and that the pipeline will pass. This is in accordance with the rule of a CI/CD pipeline that states that the same environment should be used for testing and deploying code into production. As stated in the Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, "the same environment should be used for testing and deploying code into production, or the pipeline will fail."

NEW QUESTION: 12

An application is being built to collect and display telemetry streaming data. Drag and drop the elements of this stack from the left onto the correct functions on the right.



Answer:



NEW QUESTION: 13

Refer to the exhibit.

```

public class RestaurantEntity {
    private String name;
    private String address;
    private int rating;

    public RestaurantEntity(String name,
        String address, int rating)
    {
        this.name=name;
        this.address=address;
        this.rating=rating;
    }
}

```

An application has been developed as a fileserver. To create this application as a fileserver, the class is used to store data in an SQL database.

Which two steps help minimize the risk of attack and compromise of the application? (Choose two.)

- A. Validate the content of the name and address parameters
- B. Use floating point instead of integer for rating to avoid overflow

- C. Reject string values that are greater than maximum size
- D. Compile the application using Java HotSpot
- E. Run the application in a Docker container

Answer: A,C ([LEAVE A REPLY](#))

NEW QUESTION: 14

ConfigMap keys have been mapped to different file names using the volumes.configMap.items field. What is the result if a wrong ConfigMap key is specified?

- A. The volume is not created.
- B. The volume is created with errors.
- C. The default path is not used.
- D. The volume is created.

Answer: A ([LEAVE A REPLY](#))

Explanation

If a wrong ConfigMap key is specified, the volume will not be created since the ConfigMap key must match the name of the item specified in the volumes.configMap.items field.

NEW QUESTION: 15

The E-commerce application is being monitored using AppDynamics. AppDynamics has noticed that the application response has degraded and has identified some thread contention that might be the cause of the delays.

Where in the interface does AppDynamics bring this to your attention?

- A. Potential Issues section of an individual transaction
- B. Transaction Score graph of a business transaction
- C. Code Deadlock section of the Event List
- D. Slow Calls and Errors tab of the application

Answer: A ([LEAVE A REPLY](#))

Section: Logging, Monitoring, and Metrics

NEW QUESTION: 16

Which two actions help limit the attack surface of your Docker container? (Choose two.)

- A. Run all services in a single image.
- B. Run only a single service in each container.
- C. Download images over HTTPS supporting sites.
- D. Use version tags for base images and dependencies.
- E. Use Kali Linux as a base image.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 17

What is chaos engineering?

- A. a product engineering approach that lacks a formal methodology
- B. splitting a project between teams that do not communicate with each other
- C. designing an application with fuzzy logic that makes it hard to predict
- D. introducing problems into a system to study behavior and improve reliability

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 18

Refer to the exhibit.

```
TASK [Print the contents of the "show_ip_int_brief" variable]
*****
ok: [ios-xe-mgmt.cisco.com] => {
  "show_ip_int_brief": {
    "changed": false,
    "failed": false,
    "stdout": [
      "Interface          IP-Address      OK? Method Status
Protocol\nGigabitEthernet1  10.10.20.48     YES NVRAM  up    up
\nGigabitEthernet2      10.10.10.18     YES other  up    up    \nGigabitEthernet3
13.13.13.13         Yes other        up    up"
    ],
    "stdout_lines": [
      [
        "Interface          IP-Address      OK? Method  Status      Protocol",
        "GigabitEthernet1  10.10.20.48     YES NVRAM  up          up",
        "GigabitEthernet2  10.10.10.18     YES other  up          up",
        "GigabitEthernet3  13.13.13.13     YES other  up          up"
      ]
    ]
  }
}
```

The exhibit shows the output of an Ansible task that prints the contents of the `show_ip_int_brief` variable that was registered in a different task in the playbook.

Which expression is used to print the output of the command without its header row?

- A. `show_ip_int_brief['stdout_lines'][0][1:]`
- B. `show_ip_int_brief['stdout_lines'][1:]`
- C. `show_ip_int_brief['stdout_lines']`
- D. `show_ip_int_brief['stdout_lines'][0]`

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 19

Which Dockerfile yields the most predictable builds?

```
FROM python:latest
RUN apt update -y
RUN apt install git -y

RUN pip install requests==2.22.0
```

A.

```
FROM python:3.8.0b2-buster
RUN apt update -y
RUN apt install git -y

RUN pip install requests==2.22.0
```

B.

```
FROM python
RUN apt update -y
RUN apt upgrade -y
RUN apt install git -y

RUN pip install requests
```

C.

```
FROM python:3.8.0b2-buster
RUN apt update -y
RUN apt upgrade -y
RUN apt install git -y

RUN pip install requests==2.22.0
```

D.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 20

A team must roll out updates on the production instance of an application. To avoid any service down time, the new version must be released alongside the old version. After the team ensures that the new version operates as it should, traffic will be redirected to the new version. Then the old version will be shut down.

Which deployment strategy must be used?

- A. rolling
- B. blue/green
- C. shadow
- D. canary

Answer: B ([LEAVE A REPLY](#))

The new version of the application is deployed to one of the environments, and then traffic is gradually shifted to the updated environment. This strategy allows the team to test the new version before releasing it to the public, ensuring that it is operating as expected. After the new version is tested and approved, traffic is fully shifted to the new environment and the old environment is shut down.

NEW QUESTION: 21

Which two actions help limit the attack surface of your Docker container? (Choose two.)

- A. Run only a single service in each container.
- B. Run all services in a single image.
- C. Use version tags for base images and dependencies.
- D. Use Kali Linux as a base image.
- E. Download images over HTTPS supporting sites.

Answer: A,C (LEAVE A REPLY)

Running only a single service in each container and using version tags for base images and dependencies helps limit the attack surface of your Docker container. This ensures that only the necessary services are running and that you always have the latest versions of the base images and their dependencies, reducing the risk of malicious code being included in the container image. Reference: Docker Documentation, Security Best Practices.

NEW QUESTION: 22

As a department starts following NetDevOps principles, how should the network monitoring strategy be updated to provide more visibility into the user experience?

- A. Leverage machine learning tools to more quickly identify problems in syslog and SNMP data
- B. Update tools to pull or stream monitoring metrics from APIs instead of using legacy protocols
- C. Identify and capture the key performance indicators that describe the state of the network
- D. Identify and fix network issues faster by leveraging streaming telemetry from network devices

Answer: C (LEAVE A REPLY)

Updating tools to pull or stream monitoring metrics from APIs instead of using legacy protocols is an important part of a NetDevOps strategy, as it allows for more efficient and automated data collection, as well as better visibility into the user experience. APIs for network monitoring are becoming increasingly common, as they allow for a more direct connection to the devices themselves, and can provide access to more detailed and specific metrics. Cisco's DevOps Solutions and Practices Guide recommends leveraging APIs to monitor the user experience, as they can provide access to more specific metrics and can be integrated with other monitoring tools.

NEW QUESTION: 23

Which interface is most commonly used to integrate logging, monitoring, and alerting applications into your CI/ CD pipeline?

- A. AMQP
- B. SNMP
- C. SSH
- D. REST

Answer: (SHOW ANSWER)

The most commonly used interface to integrate logging, monitoring, and alerting applications into your CI/CD pipeline is REST. Representational State Transfer (REST) is an architectural style that defines a set of constraints and properties based on HTTP. It enables

developers to create APIs that are easy to use and access, and is well-suited for integrating logging, monitoring, and alerting applications into your CI/CD pipeline.

NEW QUESTION: 24

Which approach must be used to integrate DevOps practices into the team structure in an existing organization?

- A. Create a team for each stage with structured communication channels
- B. Create a team that can deliver all the stages with active cooperation between the team members
- C. Create a team for each stage with active cooperation between the teams
- D. Create a team that can deliver all the stages with well-defined roles and a communication structure between team members

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 25

A developer wants to deploy a highly available web server cluster and decides to put a load balancer in front of multiple clustered nodes that run the same web service. The goal is for the load balancer to take in users and distribute the load across the whole cluster.

What kind of high-availability configuration is the developer running?

- A. active
- B. passive
- C. active-passive
- D. active-active

Answer: [C \(LEAVE A REPLY\)](#)

Section: Cloud and Multicloud

NEW QUESTION: 26

An application is being developed to be used by the HR department of a company on a global scale. The application will be used as a central repository for employee contracts. For user access, a RADIUS server will be used with authorized user groups.

Which action must be used to prevent developers from accidentally committing secrets in the code?

- A. Add a unit test to block the secrets
- B. Add a precommit Git hook to block the secrets
- C. Add a job in the CI build to block the secrets
- D. Add dast to the repository to block the secrets

Answer: [A \(LEAVE A REPLY\)](#)

NEW QUESTION: 27

A developer is nearing the end of a software development cycle and is ready to deploy the software. The customer wants to continue using a secure software development lifecycle methodology and must plan for postproduction components. Which two areas should be the developer address? (Choose two.)

- A. performing code reviews
- B. requirement gathering
- C. change management process
- D. new code scanning
- E. rollback plans

Answer: ([SHOW ANSWER](#))

Section: Security

NEW QUESTION: 28

What is the impact of using the Drone.io CI/CD tool on the local installation step?

- A. slows down the development
- B. speeds up the procedure
- C. complicates the application process
- D. delays the deployment of components

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 29

What is an advantage of using configuration management tools to automate infrastructure services?

- A. eliminates the need to run integration tests within the CI/CD pipeline
- B. integrates with container orchestration
- C. eliminates the need for CI/CD tools
- D. provides high native monitoring of services

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 30

Refer to the exhibit.

```
kind: pipeline
name: default

steps:
- name: test
  image: python
  commands:
  - pip install -r requirements.txt
  - pytest
```

Which CI solution uses this file?

- A. Drone
- B. GitLab CI
- C. Travis CI
- D. Jenkins

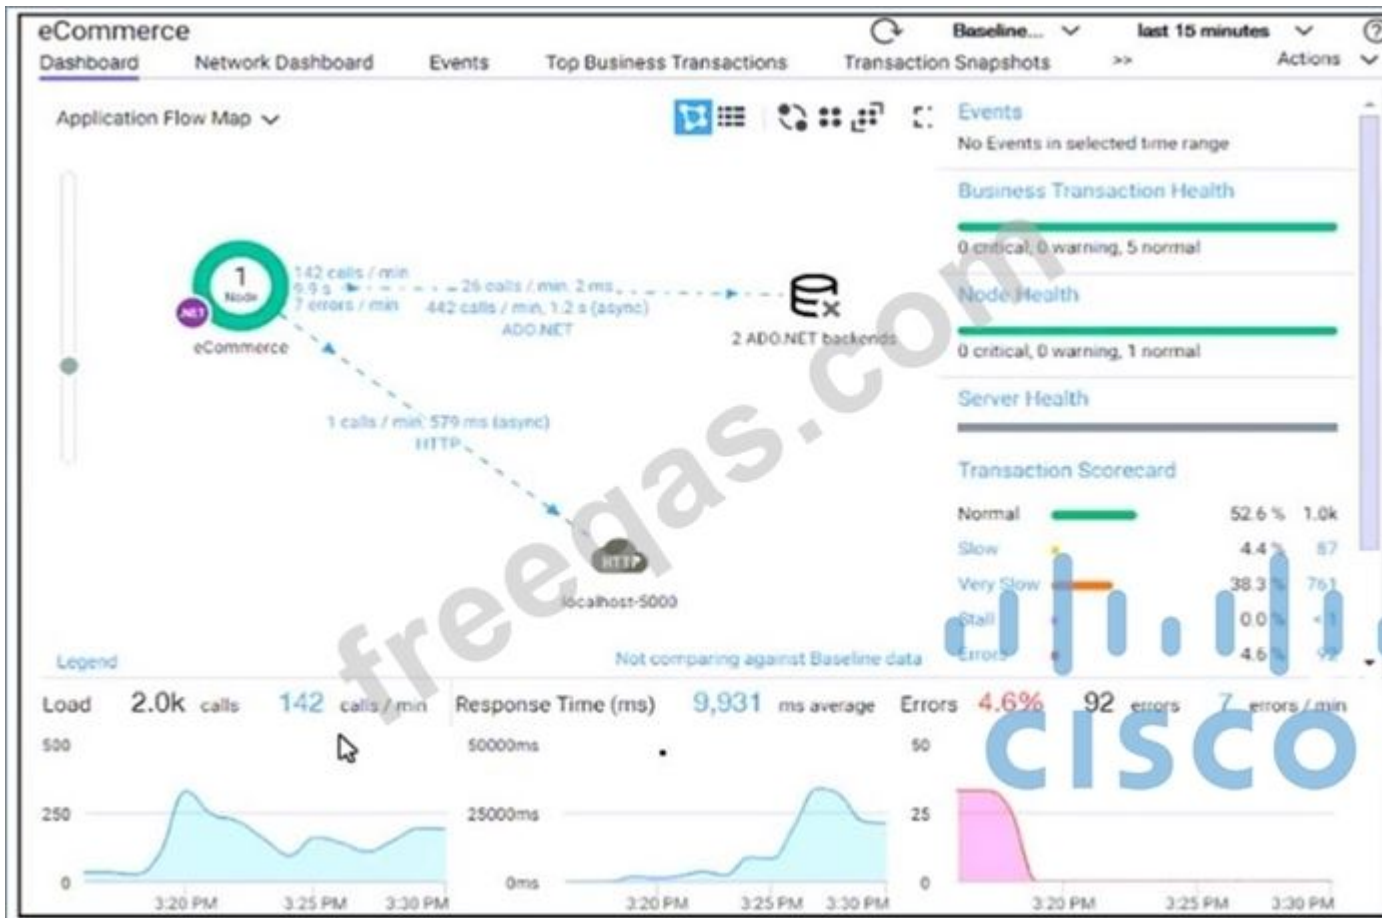
Answer: A ([LEAVE A REPLY](#))

syntax for GitLab CI, Travis CI & Jenkins are "stage...", while Drone use "steps..." Here are the configuration files I referred

<https://docs.drone.io/pipeline/environment/syntax/> <https://docs.gitlab.com/ee/ci/yaml/includes.html> <https://blog.travis-ci.com/2019-05-30-setting-up-a-ci-cd-process-on-github> <https://www.jenkins.io/doc/pipeline/examples/>

NEW QUESTION: 31

Refer to the exhibit.



An application for data storing and processing has been created. The application is time-sensitive and manages confidential data. AppDynamics identifies error transactions within a distributed application.

Which action will help resolve the errors?

- A. Rewrite the SQL query identified in the transaction to be more efficient
- B. Log in to the node directly and investigate the errors by reviewing log files
- C. Drill down to investigate and identify the errors
- D. Scale out the node horizontally

Answer: C (LEAVE A REPLY)

<https://docs.appdynamics.com/appd/23.x/latest/en/application-monitoring/remote-services/monitor-databases>

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the **newest 300-910 exam dumps**, the PrepPdf.com 300-910 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 32

Which step must be taken to enable centralized logging in a Kubernetes environment?

- A. No steps need to be taken. The master node automatically aggregates logs from all worker nodes and stores them on the specified persistent volume.
- B. Create a CustomResourceDefinition in each deployment that specifies the IP or names the log collector.
- C. Deploy a sidecar node that aggregates logs from the entire cluster.
- D. Create a DaemonSet that deploys a container with a logging agent on every node in the cluster.

Answer: C (LEAVE A REPLY)

Section: Cloud and Multicloud


NEW QUESTION: 33

Drag and drop the operations from the left into the correct order on the right when using CI/CD, assuming that there are no failures or conflicts.

| | |
|---------------------------------|--------|
| generating deliverables | step 1 |
| starting of rolling deployment | step 2 |
| developer commits | step 3 |
| integration testing | step 4 |
| unit testing | step 5 |
| merging of commit into mainline | step 6 |

Answer:

| | |
|---------------------------------|---------------------------------|
| generating deliverables | generating deliverables |
| starting of rolling deployment | merging of commit into mainline |
| developer commits | starting of rolling deployment |
| integration testing | unit testing |
| unit testing | integration testing |
| merging of commit into mainline | developer commits |



NEW QUESTION: 34

Which two characteristics of a build tool that is used in an automated build pipeline? (Choose two.)

- A. It builds your software every 500 milliseconds.
- B. It must have a plug-in for Puppet and Maven.
- C. It builds your software on the production environment.
- D. It runs unit tests against code that is committed.
- E. It builds your software when changes are committed.

Answer: ([SHOW ANSWER](#))

Explanation

The two characteristics of a build tool that is used in an automated build pipeline are A) it builds your software when changes are committed and B) it runs unit tests against code that is committed. Automated build tools are used to detect changes to source code and then build the software that resulted from the change.

Additionally, these tools can be configured to run unit tests against the code to ensure that it works as expected.

NEW QUESTION: 35

Construct an Ansible script to gather information about target routers and then use it to apply no ip redirects to every interface on each device. Drag and drop the code from the bottom onto the correct location in the exhibit to complete the tasks section of the Ansible playbook so that it accomplishes your goal.

```

tasks:
- name: Get info from devices
  
  provider: "{{ credentials }}"
   dev_info
- name: Add 'no ip redirects' to all interfaces
  
  provider: "{{ credentials }}"

  lines:
  - "no ip redirects"
  parents: "interface {{ item.key }}"
  with_items: "{{ dev_info['ansible_facts']['ansible_net_interfaces'] | dict2items }}"

```




Answer:

```

tasks:
- name: Get info from devices
  
  provider: "{{ credentials }}"
   dev_info
- name: Add 'no ip redirects' to all interfaces
  
  provider: "{{ credentials }}"

  lines:
  - "no ip redirects"
  parents: "interface {{ item.key }}"
  with_items: "{{ dev_info['ansible_facts']['ansible_net_interfaces'] | dict2items }}"

```


Explanation

ios_facts

register

ios_config

NEW QUESTION: 36

A DevOps engineer must validate the working state of the network before implementing a CI/CD pipeline model. Which configuration management tool is designed to accomplish this?

- A. Python YAML data libraries
- B. Jenkins
- C. Genie CLI
- D. Travis CI

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 37

Refer to the exhibit.

```
1 ---
2 version: '3'
3 services:
4   database:
5     image: postgres
6     container_name: postgres_data
7     ports:
8       - "5432:5432"
9     environment:
10      POSTGRES_DB: "postgres_database"
11      POSTGRES_USER: "postgres"
12      POSTGRES_PASSWORD: "testpass"
13     networks:
14       - "database_net"
15     healthcheck:
16       test: ['CMD', 'psql', '--username', 'postgres']
17       interval: 10s
18       timeout: 10s
19       retries: 5
20
21 networks:
22   database_net:
23     driver: "bridge"
```

A docker-compose.yml file implements a postgres database container.

Which .gitlab-ci.yml code block checks the health status of the container and stops the pipeline if the container is unhealthy?

```
Validate Application Infrastructure
stage: validate_infrastructure
before_script:
  - apk add --no-cache docker-compose
script:
  - docker-compose up -d
  - sleep 15s
  - health_state=$(docker container inspect
    postgres_data --
    format='{{.State.Health.Status}}')
  - if [[ "$health_state" != "healthy" ]]; then
    exit 1; fi
```

A.

```

Validate Application Infrastructure:
stage: validate_infrastructure
before_script:
- apk add --no-cache docker-compose
script:
- docker-compose up
- sleep 15s
- health_state = $(docker container state
  postgres --health-status)
- if [[ "$health_state" != "running" ]]; then
  exit 1; fi

```

B.

```

Validate Application Infrastructure:
stage: validate_infrastructure
before_script:
- apk add --no-cache docker-compose
script:
- docker-compose up
- sleep 15s
- health_state = $(docker container logs
  postgres_data | grep healthcheck)
- $health_state == 'healthy'

```

C.

```

Validate Application Infrastructure:
stage: validate_infrastructure
before_script:
- apk add --no-cache docker-compose
script:
- docker-compose up
- sleep 15s
- health_state = $(docker container attach
  postgres_data &&
  psql -username postgres status)
- $health_state == 'running'

```

D.

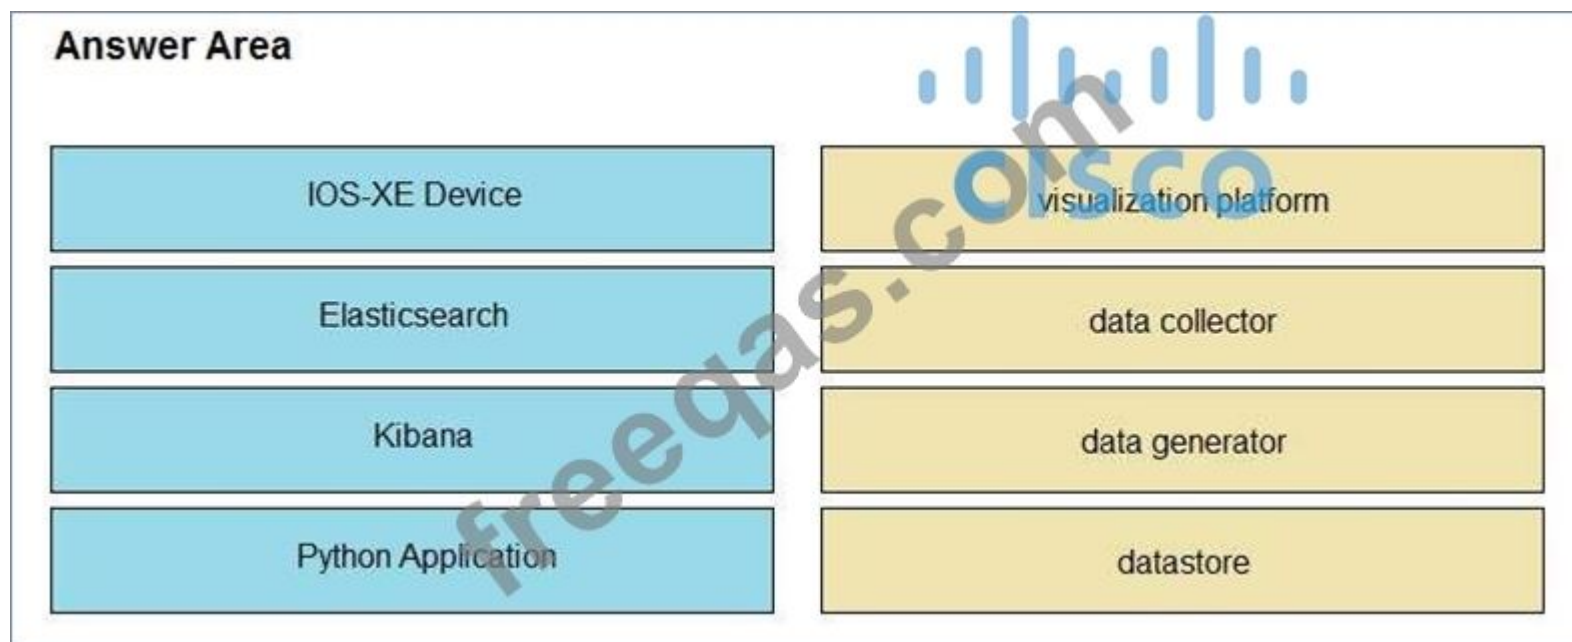
Answer: ([SHOW ANSWER](#))

NEW QUESTION: 38

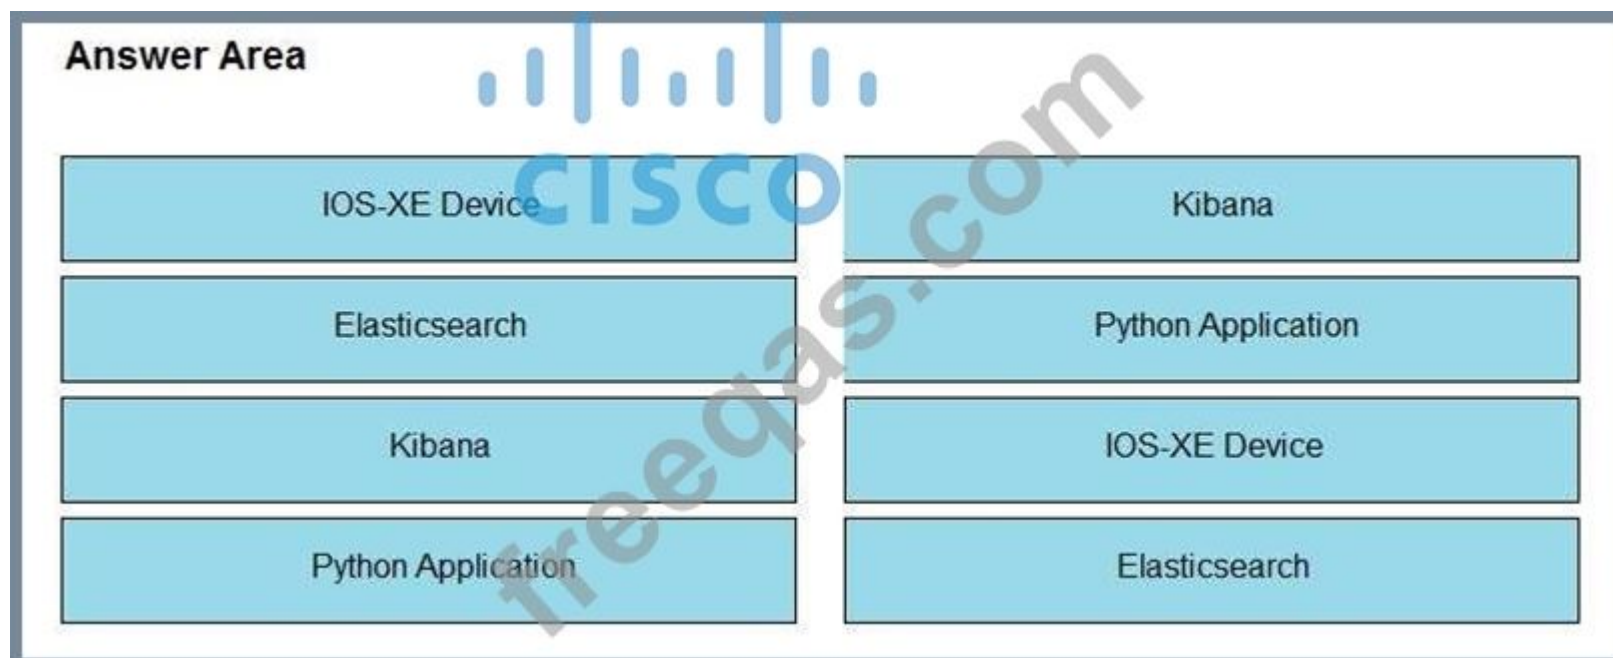
DRAG DROP

An application is being built to collect and display telemetry streaming data. Drag and drop the elements of this stack from the left onto the correct functions on the right.

Select and Place:



Answer:



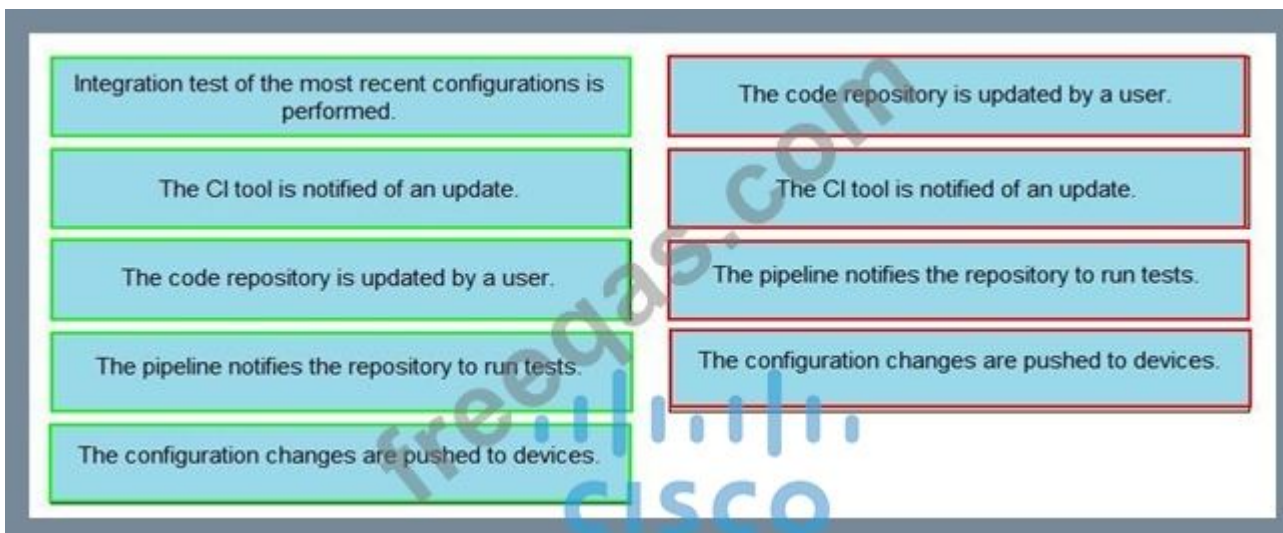
Section: Logging, Monitoring, and Metrics

NEW QUESTION: 39

The IaC configuration for an application is being deployed using a CI/CD pipeline. Drag and drop the steps for this pipeline from the left into the correct order that they would be executed on the right. Not all options are used.



Answer:



NEW QUESTION: 40

What is the effective way to manage computing costs in a public cloud?

- A. Monitor data transfers to minimize cost.
- B. Use dedicated hardware for all instances.
- C. Select the largest instance option available in order to pay for only one instance.
- D. Make use of elastic services and scale on demand.

Answer: [\(SHOW ANSWER\)](#)

Section: Packaging and Delivery of Applications

NEW QUESTION: 41

Which two statements about Infrastructure as Code are true? (Choose two.)

- A. Infrastructure as Code must use the same programming language as the application.

- B. Test-driven development practices make use of Infrastructure as Code.
- C. Infrastructure as Code refers to automated testing libraries.
- D. DevOps builds upon Infrastructure as Code.
- E. Infrastructure as Code is based on practices from software development.

Answer: ([SHOW ANSWER](#))

Explanation

(IaC) is a practice that allows for the automation of infrastructure configuration and deployment, allowing DevOps teams to deploy and manage infrastructure with the same release pipelines and versioning as source code [1]. IaC is based on practices from software development, and promotes the use of automation and repeatable processes to ensure consistency of deployments.

NEW QUESTION: 42

How long analysis systems such as Elasticsearch, Logstash, and Kibana Stack handle ingesting unstructured logs from different devices in various formats?

- A. All devices that generate syslogs must use agents that process the local logs and transmit them in a specific format to the ELK Stack.
- B. All logs are stored in their unstructured text format, and the ELK Stack performs data analysis by intelligently parsing the logs using machine learning algorithms.
- C. All different message formats are parsed separately using custom filters, and the resulting structured data is stored for later analysis.
- D. A single, comprehensive log format is defined on the ELK Stack. All incoming logs, regardless of format, are transformed to match the comprehensive format, and only applicable fields are populated.

Answer: ([SHOW ANSWER](#))

Section: Logging, Monitoring, and Metrics

NEW QUESTION: 43

As a department starts following NetDevOps principles, how should the network monitoring strategy be updated to provide more visibility into the user experience?

- A. Identify and capture the key performance indicators that describe the state of the network
- B. Identify and fix network issues faster by leveraging streaming telemetry from network devices
- C. Update tools to pull or stream monitoring metrics from APIs instead of using legacy protocols
- D. Leverage machine learning tools to more quickly identify problems in syslog and SNMP data

Answer: ([SHOW ANSWER](#))

Explanation

Updating tools to pull or stream monitoring metrics from APIs instead of using legacy protocols is an important part of a NetDevOps strategy, as it allows for more efficient and automated data collection, as well as better visibility into the user experience. APIs for network monitoring are becoming increasingly common, as they allow for a more direct connection to the devices themselves, and can provide access to more detailed and specific metrics. Cisco's DevOps Solutions and Practices Guide recommends leveraging APIs to monitor the user experience, as they can provide access to more specific metrics and can be integrated with other monitoring tools.

NEW QUESTION: 44

A developer has created a deployment that should launch a pod to run their database service. The pod should launch with a metadata name of "Cisco-DB," and the developer has added it to the "Cisco" namespace in their deployment.

Which Kubernetes command confirms that the service is running and usable?

- A. `kubectl -n Cisco get service | grep "Cisco-DB"`
- B. `kubectl -n Cisco get services | grep "Cisco-DB"`
- C. `kubectl get pods | grep "Cisco-DB"`
- D. `kubectl -n Cisco get pods | grep "Cisco-DB"`

Answer: [\(SHOW ANSWER\)](#)

Explanation

The correct answer is B. `kubectl -n Cisco get pods | grep "Cisco-DB"`. This command allows you to view the pods in the Cisco namespace and filter the list for the pod named "Cisco-DB". This will confirm that the pod is running and usable.

References: Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, Chapter 6, Section 6.3: Kubernetes Commands, Page 113.

NEW QUESTION: 45

A new version of an application is being released by creating a separate instance of the application that is running the new code. Only a small portion of the user base will be directed to the new instance until that version has been proven stable. Which deployment strategy is this example of?

- A. recreate
- B. blue/green
- C. canary
- D. rolling

Answer: [C \(LEAVE A REPLY\)](#)

NEW QUESTION: 46

A DevOps engineering wants to build an application implementation based on the CI/CD pipeline model. Which service should be used to provide hosted continuous service for open and private projects?

- A. Ansible
- B. pyATS
- C. Travis CI
- D. Genie CLI

Answer: [C \(LEAVE A REPLY\)](#)

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the **newest 300-910 exam dumps**, the PrepPdf.com 300-910 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 47

A CI/CD pipeline that builds infrastructure components using Terraform must be designed. A step in the pipeline is needed that checks for errors in any of the .tf files in the working directory. It also checks the existing state of the defined infrastructure.

Which command does the pipeline run to accomplish this goal?

- A. terraform plan
- B. terraform validate
- C. terraform check
- D. terraform fmt

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 48

Which two characteristics of a build tool that is used in an automated build pipeline? (Choose two.)

- A. It runs unit tests against code that is committed.
- B. It builds your software on the production environment.
- C. It builds your software when changes are committed.
- D. It builds your software every 500 milliseconds.
- E. It must have a plug-in for Puppet and Maven.

Answer: [A,C \(LEAVE A REPLY\)](#)

Section: Automating Infrastructure

NEW QUESTION: 49

Refer to the exhibit.

```
1  push_configs.yml content:
2  - hosts: "{{ CHANGED_HOST }}"
3    become: yes
4    become_method: enable
5    connection: network_cli
6    gather_facts: no
7    tasks:
8      - name: Push the template
9        ios_config:
10         src: "{{ changed_file }}"
11
12  Command:
13  ansible-playbook push_configs.yml -i
14  * ansible_managed_inventory -e "CHANGED_HOST=${CHANGED_HOST}
15  * CHANGES=${CHANGES}"
16
17  Error Message:
18  "msg": "paramiko: The authenticity of host '[ios-xe-mgmt-
19  * latest.cisco.com]:8181' can't be established.\nThe ssh-rea
20  * key fingerprint is b'b7e974a8cbf96d464f7be3e12a86d265'."
```

The push_configs.yml playbook returns the error shown.

Which action resolves the error?

- A. Install the Paramiko library on the host that runs Ansible
- B. Export the ANSIBLE_HOST_KEY_CHECKING=False variable
- C. Comment out the StrictHostKeyChecking=yes line from ansible.cfg

D. Generate a new SSH key pair and add the public key to the target machine

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 50

A developer is nearing the end of a software development cycle and is ready to deploy the software. The customer wants to continue using a secure software development lifecycle methodology and must plan for postproduction components. Which two areas should be the developer address? (Choose two.)

- A. new code scanning
- B. performing code reviews
- C. requirement gathering
- D. change management process
- E. rollback plans

Answer: B,D [\(LEAVE A REPLY\)](#)

NEW QUESTION: 51

A user wants to deploy a new service to a Kubernetes cluster. Which two commands accomplish this goal? (Choose two.)

- A. Create
- B. Expose
- C. Apply
- D. Install
- E. Deploy

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 52

Refer to the exhibit.

```
TASK [print the contents of the "show_ip_int_brief" variable]
*****
ok: [ios-xe-mgmt.cisco.com] => {
  "show_ip_int_brief": {
    "changed": false,
    "failed": false,
    "stdout": [
      "Interface          IP-Address      OK? Method Status
Protocol\nGigabitEthernet1  10.10.20.48    YES  NVRAM  up      up
\nGigabitEthernet2    10.10.10.18    YES  other  up      up  \nGigabitEthernet3
13.13.13.13          Yes  other    up
",
    ],
    "stdout_lines": [
      [
        "Interface          IP-Address      OK? Method Status Protocol",
        "GigabitEthernet1  10.10.20.48    YES  NVRAM  up      up",
        "GigabitEthernet2  10.10.10.18    YES  other  up      up",
        "GigabitEthernet3  13.13.13.13    YES  other  up      up",
      ]
    ]
  ]
}
```

The exhibit shows the output of an Ansible task that prints the contents of the show_ip_int_brief variable that was registered in a different task in the playbook.

Which expression is used to print the output of the command without its header row?

- A. `show_ip_int_brief['stdout_lines'][0][1:]`
- B. `show_ip_int_brief['stdout_lines']`
- C. `show_ip_int_brief['stdout_lines'][1:]`
- D. `show_ip_int_brief['stdout_lines'][0]`

Answer: A (LEAVE A REPLY)

Explanation

```
show_ip_int_brief[""stdout_lines"][0][1:]
```

`stdout_lines` - nested list,

`[0]`, first element of above nested list (extract of all header + interfaces, coma separated for individual elements),

`[1:]` - 2nd, 3rd, 4th line of the above sub-list.

NEW QUESTION: 53

Fill in the blanks to describe the concepts of extending DevOps practices to the network for NetDevOps.

NetDevOps builds and manages networks as a system that enables network services to be consumed in a DevOps approach.

Organizations practicing NetDevOps see changes in the _____ as routine and expected activities, with a well-defined and practiced process for _____, testing, and _____ network changes. By making them routine, network changes can be small and simple.

Answer:

network, designing, deploying

NEW QUESTION: 54

Drag and drop the code from the bottom onto the box where the code is missing to create a Terraform configuration that builds the network environment for a multitier software application. More EPG, Contract, and Filter definitions have been removed from the code.

```

resource "aci_application_profile" "production_multi_app" {
  tenant_dn = aci_tenant.production_tenant.id
  [redacted] = "multi_app_prod"
  name_alias = "multi_ap_prod"
  prio       = "level1"
}
resource "aci_application_epg" "prod_web" {
  [redacted] =
aci_application_profile.development_multi_app.id
  name           = "web"
  name_alias     = "Nginx"
  relation_fv_rs_bd = [redacted]
}
resource "aci_filter" "db_traffic" {
  tenant_dn = [redacted]
  name      = "db_traffic"
}
resource "aci_filter_entry" "userdb" {
  filter_dn = [redacted]
  name      = "userdb"
  [redacted] = "ip"
  prot      = "tcp"
  d_from_port = "3306"
  d_to_port   = "3306"
}

```

[redacted] aci_filter.db_traffic.id

application_profile_dn

aci_tenant.production_tenant.id

ether_t

aci_bridge_domain.production_bd.id

name

Answer:

```

resource "aci_application_profile" "production_multi_app" {
  tenant_dn = aci_tenant.production_tenant.id
  name      = "multi_app_prod"
  name_alias = "multi_ap_prod"
  prio     = "level1"
}
resource "aci_application_epg" "prod_web" {
  application_profile_dn =
aci_application_profile.development_multi_app.id
  name                  = "web"
  name_alias            = "Nginx"
  relation_fv_rs_bd     = aci_bridge_domain.production_bd.id
}
resource "aci_filter" "db_traffic" {
  tenant_dn = aci_tenant.production_tenant.id
  name      = "db_traffic"
}
resource "aci_filter_entry" "userdb" {
  filter_dn = aci_filter.db_traffic.id
  name      = "userdb"
  ether_t   = "ip"
  prot     = "tcp"
  d_from_port = "3306"
  d_to_port  = "3306"
}

```

| | |
|------------------------------------|------------------------|
| aci_filter.db_traffic.id | application_profile_dn |
| aci_tenant.production_tenant.id | ether_t |
| aci_bridge_domain.production_bd.id | name |

Explanation

```

resource "aci_application_profile" "production_multi_app" {
  tenant_dn = aci_tenant.production_tenant.id
  name      = "multi_app_prod"
  name_alias = "multi_ap_prod"
  prio     = "level1"
}
resource "aci_application_epg" "prod_web" {
  application_profile_dn =
aci_application_profile.development_multi_app.id
  name                   = "web"
  name_alias             = "Nginx"
  relation_fv_rs_bd     = aci_bridge_domain.production_bd.id
}
resource "aci_filter" "db_traffic" {
  tenant_dn = aci_tenant.production_tenant.id
  name      = "db_traffic"
}
resource "aci_filter_entry" "userdb" {
  filter_dn = aci_filter.db_traffic.id
  name      = "userdb"
  ether_t   = "ip"
  prot      = "tcp"
  d_from_port = "3306"
  d_to_port  = "3306"
}

```

Correct Answer:

| | |
|------------------------------------|------------------------|
| aci_filter.db_traffic.id | application_profile_dn |
| aci_tenant.production_tenant.id | ether_t |
| aci_bridge_domain.production_bd.id | name |

NEW QUESTION: 55

Fill in the blank to complete the statement.

A user wants a Kubernetes deployment to run three separate pods of a web application at one time. In the deployment YAML, the user must configure the _____ field in the _____ subsection.

A. replicasspec

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 56

Refer to the Exhibit.

Drag and drop the code snippets from the bottom onto the boxes in the code in the Ansible playbook to restart multiple services in the correct sequence Not all options are used

```

- name: 'Restart web and db servers at once'
hosts: CHINESEDUMPS
tasks: 通过测试
  - name: 'Restart web and db servers at once'
    command:
  - name: 'Start the database services on db server nodes'
    service:
      name: mysql
      state:
# Inventory file
[db_nodes]
sql_db1
sql_db2
[web_nodes]
web_node2
[monitoring]
mnt_node
[web_service:children]
web_nodes CHINESEDUMPS
db_nodes 通过测试

```

| | | |
|--------|---------------------|-------------|
| reboot | started | web_service |
| all | '/sbin/shutdown -r' | reloaded |

Answer:

```

- name: 'Restart web and db servers at once'
hosts: all
tasks:
  - name: 'Restart web and db servers at once'
    command: reboot
  - name: 'Start the database services on db server nodes'
    service:
      name: mysql
      state: reloaded
# Inventory file
[db_nodes]
sql_db1
sql_db2
[web_nodes]
web_node2
[monitoring]
mnt_node
[web_service:children]
web_nodes
db_nodes

```

| | | |
|--------|-----------------------|-------------|
| reboot | started | web_service |
| all | ['/sbin/shutdown -r'] | reloaded |

NEW QUESTION: 57

An IT department needs to deploy a new application named 'Entfin434772390' across the entire enterprise. The deployment must gradually transfer user traffic from an older and nearly identical version of the application named 'Entfin027549321'. Both versions of the application are running in the production environment.

Which release strategy should be used by the IT department?

- A. agile
- B. rollbacks
- C. canary
- D. blue/green

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 58

Which step must be taken to enable centralized logging in a Kubernetes environment?

- A. No steps need to be taken. The master node automatically aggregates logs from all worker nodes and stores them on the specified persistent volume.
- B. Create a CustomResourceDefinition in each deployment that specifies the IP or names the log collector.
- C. Deploy a sidecar node that aggregates logs from the entire cluster.

D. Create a DaemonSet that deploys a container with a logging agent on every node in the cluster.

Answer: [\(SHOW ANSWER\)](#)

To enable centralized logging in a Kubernetes environment, you must create a DaemonSet that deploys a container with a logging agent on every node in the cluster. This allows all of the logs from each node to be collected in one place, allowing for easier analysis and management. Reference: Kubernetes Documentation, Logging Architecture.

NEW QUESTION: 59

An end user is seeing long web page load times on the internal business application that they are trying to view. The user is seeing this issue across multiple web browsers, and other users encounter the same issue.

Which action should the system administrator take to start looking for the cause of this issue?

- A. Check to make sure Nginx is running.
- B. Check for response times in Nginx logs.
- C. Check to make sure the web API response is coming back in JSON.
- D. Check the size of the database that the application is using.

Answer: [B \(LEAVE A REPLY\)](#)

Section: Logging, Monitoring, and Metrics

NEW QUESTION: 60

Microservices architecture pattern has been applied and the system has been architected as a set of services. Each service is deployed as a set of instances for throughput and availability. In which two ways are these services packaged and deployed? (Choose two.)

- A. Service instances must be isolated from one another.
- B. Service are written using the same languages, frameworks, and framework versions.
- C. Service must be dependent, deployable, and scalable.
- D. Service instances do not need to be isolated from one another.
- E. Service must be independently deployable and scalable.

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 61

Refer to the exhibit.

```
Traceback (most recent call last):
  File "api-call.py", line 1, in <module>
    import requests
  File "/Users/devnet/venv/devops/lib/python3.7/site-packages/requests/_init_.py", line 43, in <module>
    import urllib3
ModuleNotFoundError: No Module named 'urllib3'
```



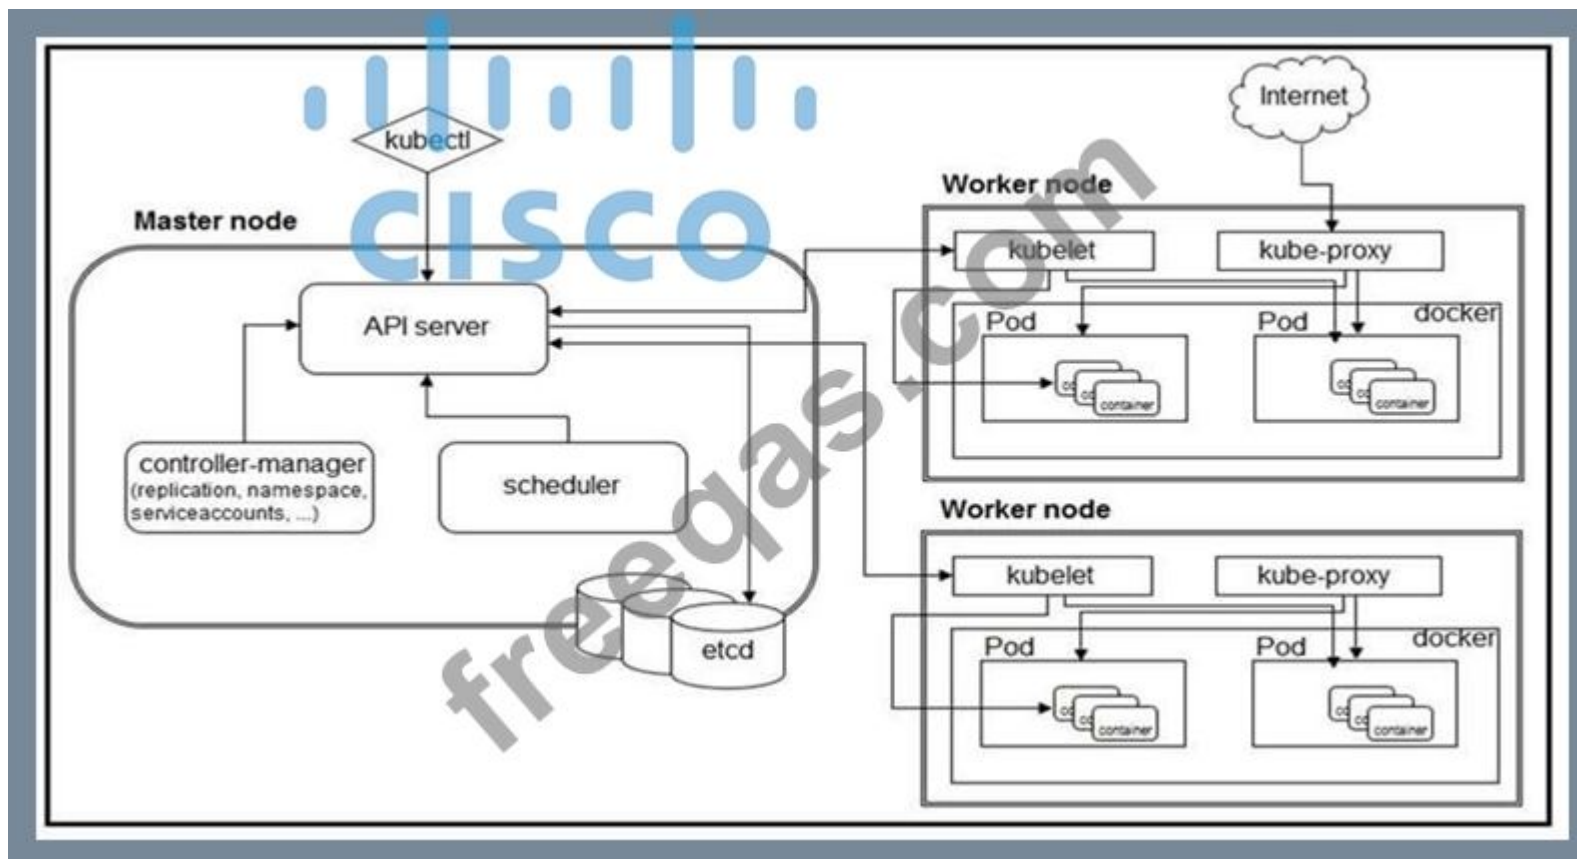
What is the reason for this error message?

- A. The site-packages directory has been corrupted.
- B. The required dependencies for the urllib3 module are not installed.
- C. The required dependencies for the requests module are not installed.
- D. The requests module is not installed.

Answer: [C \(LEAVE A REPLY\)](#)

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the newest 300-910 exam dumps, the PrepPdf.com 300-910 exam questions have been updated and answers have been corrected get the newest PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 62



Refer to the exhibit. A developer needs to scale the existing pods within the worked nodes. Which object should be edited to achieve this goal?

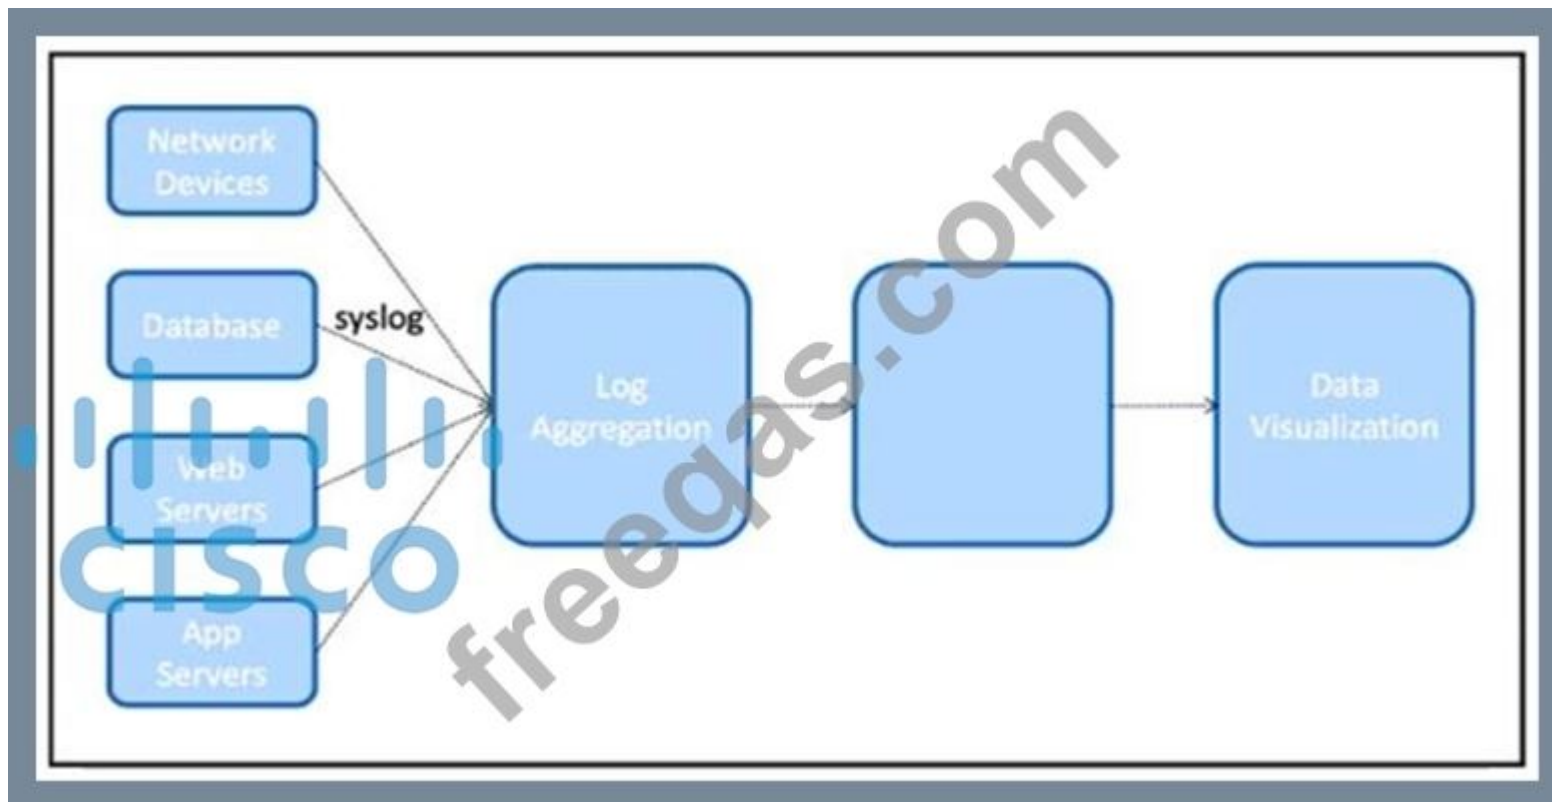
- A. ReplicaSet
- B. PriorityClass
- C. Deployment
- D. Pod

Answer: (SHOW ANSWER)

Section: Automating Infrastructure

NEW QUESTION: 63

Refer to the exhibit.



The IT team is creating a new design for a logging system. The system must be able to collect logs from different components of the infrastructure using the SNMP protocol. When the data is collected it will need to be presented in a graphical UI to the NOC team.

What is the architectural component that needs to be placed in the unlabeled box to complete this design?

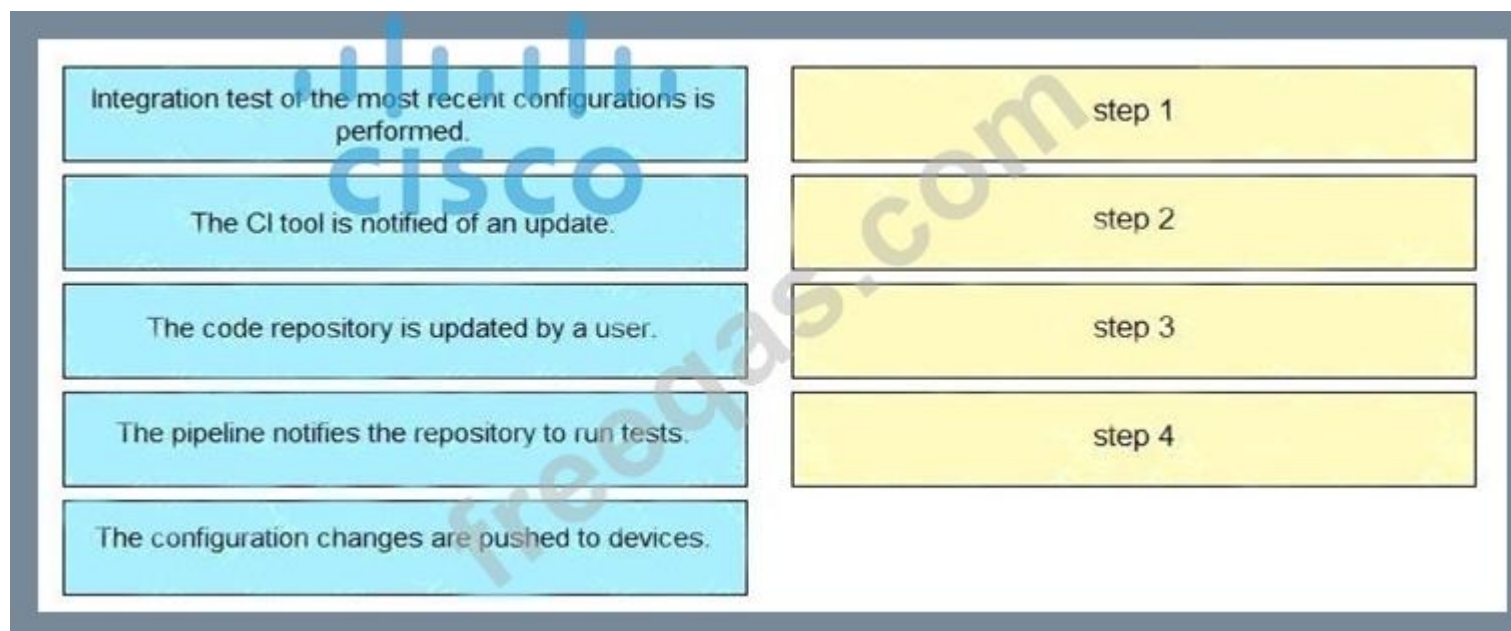
- A. message queue
- B. web server
- C. log parser
- D. time-series database

Answer: [\(SHOW ANSWER\)](#)

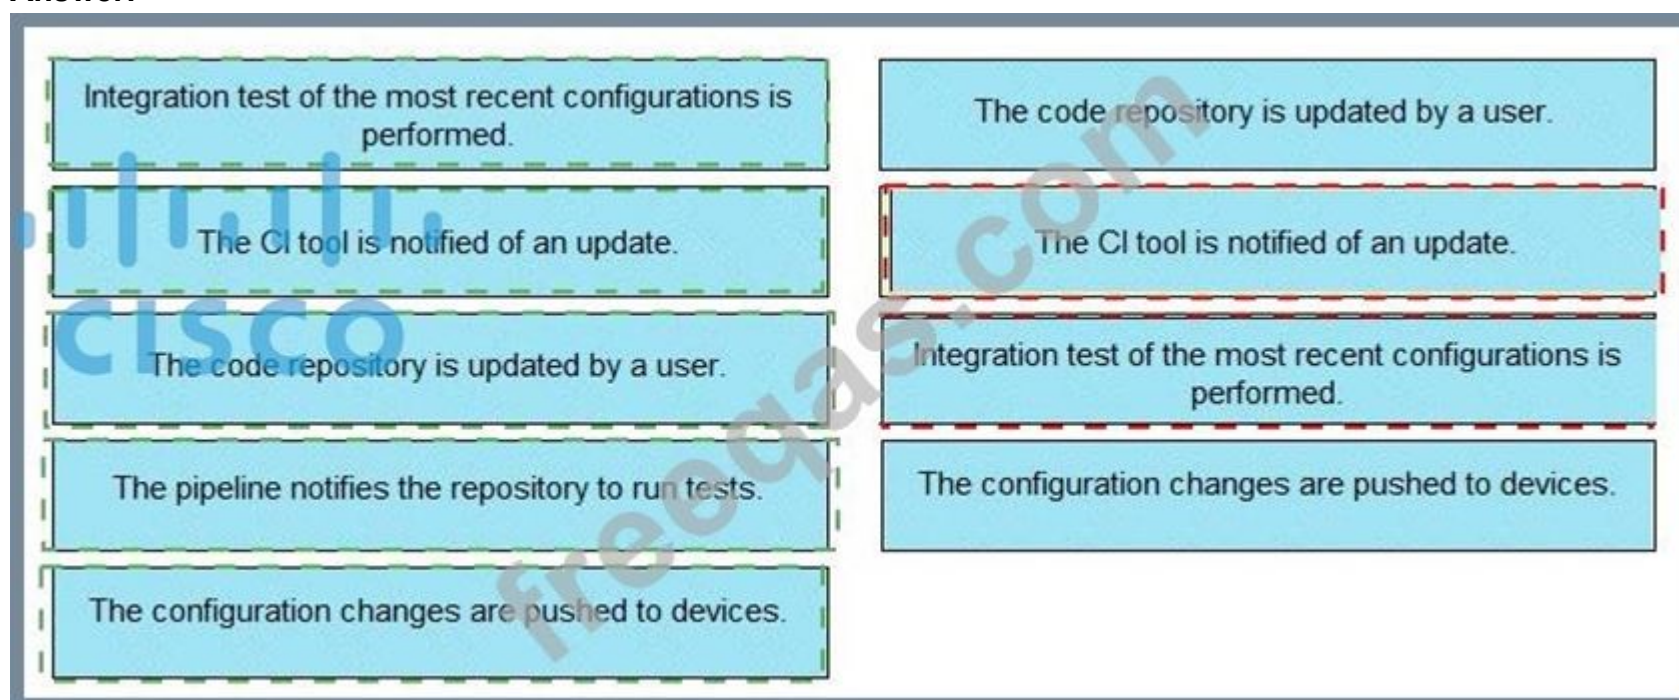
A log parser is a software component that is used to collect and parse log data from various sources, such as infrastructure devices, applications, and services. It is responsible for collecting log data from the various sources and then mapping it to a specific format that can be easily consumed and interpreted by the NOC team. The log parser can also be used to filter out irrelevant data and store the processed data in a time-series database. This data can then be used by the NOC team to generate visualizations and reports, which will help them identify and address any issues within the infrastructure. (Source: Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual Chapter 5, Understanding Logging and Log Analysis)

NEW QUESTION: 64

The IaC configuration for an application is being deployed using a CI/CD pipeline. Drag and drop the steps for this pipeline from the left into the correct order that they would be executed on the right. Not all options are used.



Answer:



Explanation

- The code repository is updated by a user.
- The CI tool is notified of an update.
- Integration test of the most recent configurations is performed
- The configuration changes are pushed to devices.

NEW QUESTION: 65

Drag and drop the tools from the left onto the tools function on the right Not all options are used.



Answer:

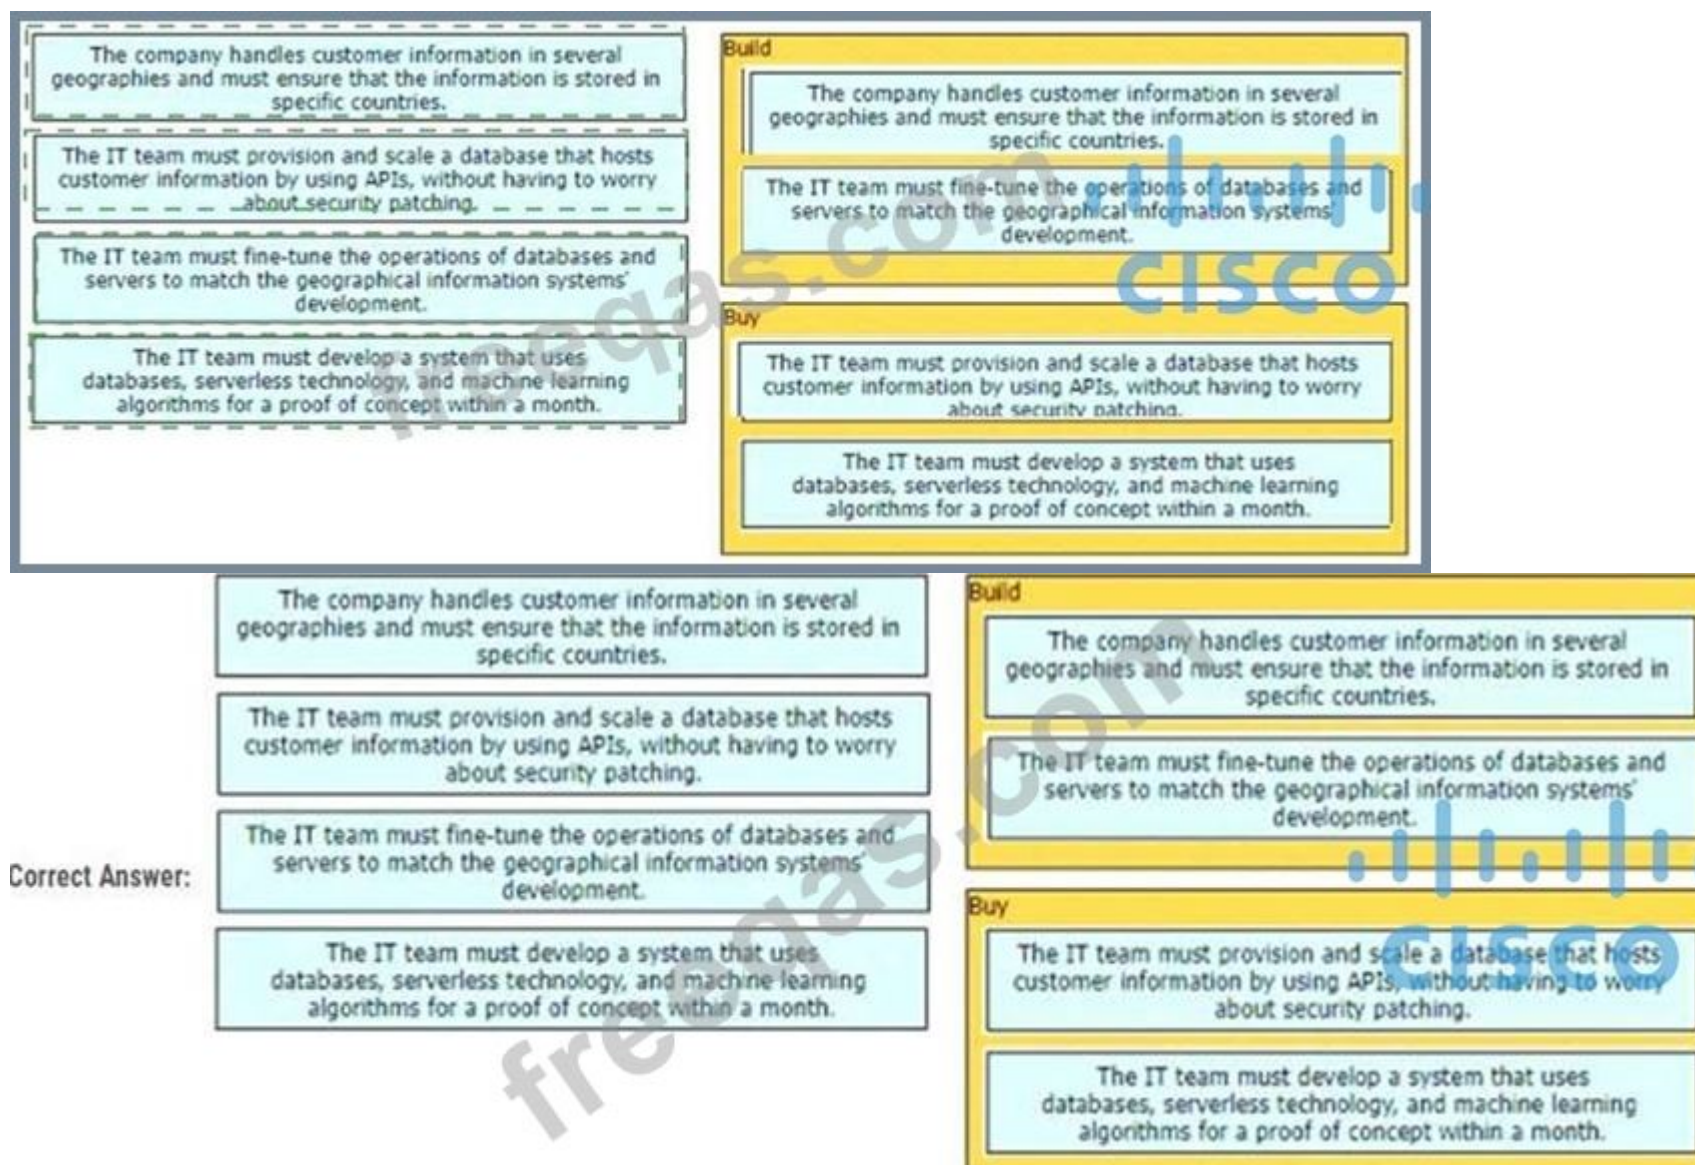


NEW QUESTION: 66

| | |
|---|--|
| <p>The company handles customer information in several geographies and must ensure that the information is stored in specific countries.</p> | <p>Build</p> <div style="border: 1px solid black; height: 40px; width: 100%;"></div> <div style="border: 1px solid black; height: 40px; width: 100%;"></div> |
| <p>The IT team must provision and scale a database that hosts customer information by using APIs, without having to worry about security patching.</p> | |
| <p>The IT team must fine-tune the operations of databases and servers to match the geographical information systems' development.</p> | |
| <p>The IT team must develop a system that uses databases, serverless technology, and machine learning algorithms for a proof of concept within a month.</p> | |
| | <p>Buy</p> <div style="border: 1px solid black; height: 40px; width: 100%;"></div> <div style="border: 1px solid black; height: 40px; width: 100%;"></div> |

Drag and drop the scenarios from the left onto the cloud strategy categories on the right.

Answer:



NEW QUESTION: 67

Refer to the exhibit.

```

$ kubectl create deployment hello-app --
image=k8s.gcr.io/echoserver:1.4
deployment.apps/hello-app created

$ kubectl get pods
NAME                                READY   STATUS    RESTARTS   AGE
hello-app-857b7d747f-xg8kj          1/1     Running   0           35s

$ kubectl get services
NAME          TYPE        CLUSTER-IP   EXTERNAL-IP   PORT(S)    AGE
kubernetes   ClusterIP   10.96.0.1    <none>        443/TCP    44h

```

Which action allows the development team to reach the deployed application?

- A. Create a service to expose the logic running in the pod
- B. Delete the deployment and redeploy by using the latest tag for the container image
- C. Delete the deployment and redeploy by using a ReplicaSet.
- D. Create an init container to initialize routes for the containers in the pod

Answer: (SHOW ANSWER)

NEW QUESTION: 68

Fill in the blanks to complete the statement.

When creating multiple containers, in order to allow communication with each other, you must create a _____ of type _____.

A. podvolume

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 69

Refer to the exhibit.

```
1 {
2   "_index": "linux-servers",
3   "_type": "_doc",
4   "_id": "mr250XQ8MiQCMA9yzEQ6",
5   "_version": 1,
6   "_score": null,
7   "_source": {
8     "message": "<38>Aug 29 12:08:37 centos systemd-logind: Removed session
9     19802.",
10    "@timestamp": "2020-03-29T09:08:37.280Z",
11    "type": "syslog",
12    "@version": "1",
13    "hostname": "vcs.local.lan",
14    "host": "172.16.40.168"
15  },
16  "fields": {
17    "@timestamp": [
18      "2020-08-29T12:03:37.280Z"
19    ]
20  },
21  "sort": [
22    1598692117280
23  ]
24 }
```



The JSON object represents a single entry on a centralized log server, but log data cannot be processed because of the format.

What causes the issue?

- A. A hostgroup must be defined
- B. The process name in the message must be parsed into a field
- C. The priority of the message must be to the server
- D. The "_type" must represent the process type

Answer: **B** ([LEAVE A REPLY](#))

NEW QUESTION: 70

A company has experienced serious growth and the hosted applications are getting more and more hits. Based on the increased work load, applications have started to show signs of reduced performance, which is affecting the user experience.

Which two server metrics should be tracked to ensure that the customer experience remains within acceptable limits? (Choose two.)

- A. application feature sprawl
- B. CPU peak usage
- C. CPU average usage

D. microservices count

E. CPU frequency

Answer: B,C (LEAVE A REPLY)

Section: Packaging and Delivery of Applications

NEW QUESTION: 71

Which two actions help limit the attack surface of your Docker container? (Choose two.)

A. Run only a single service in each container.

B. Run all services in a single image.

C. Use version tags for base images and dependencies.

D. Use Kali Linux as a base image.

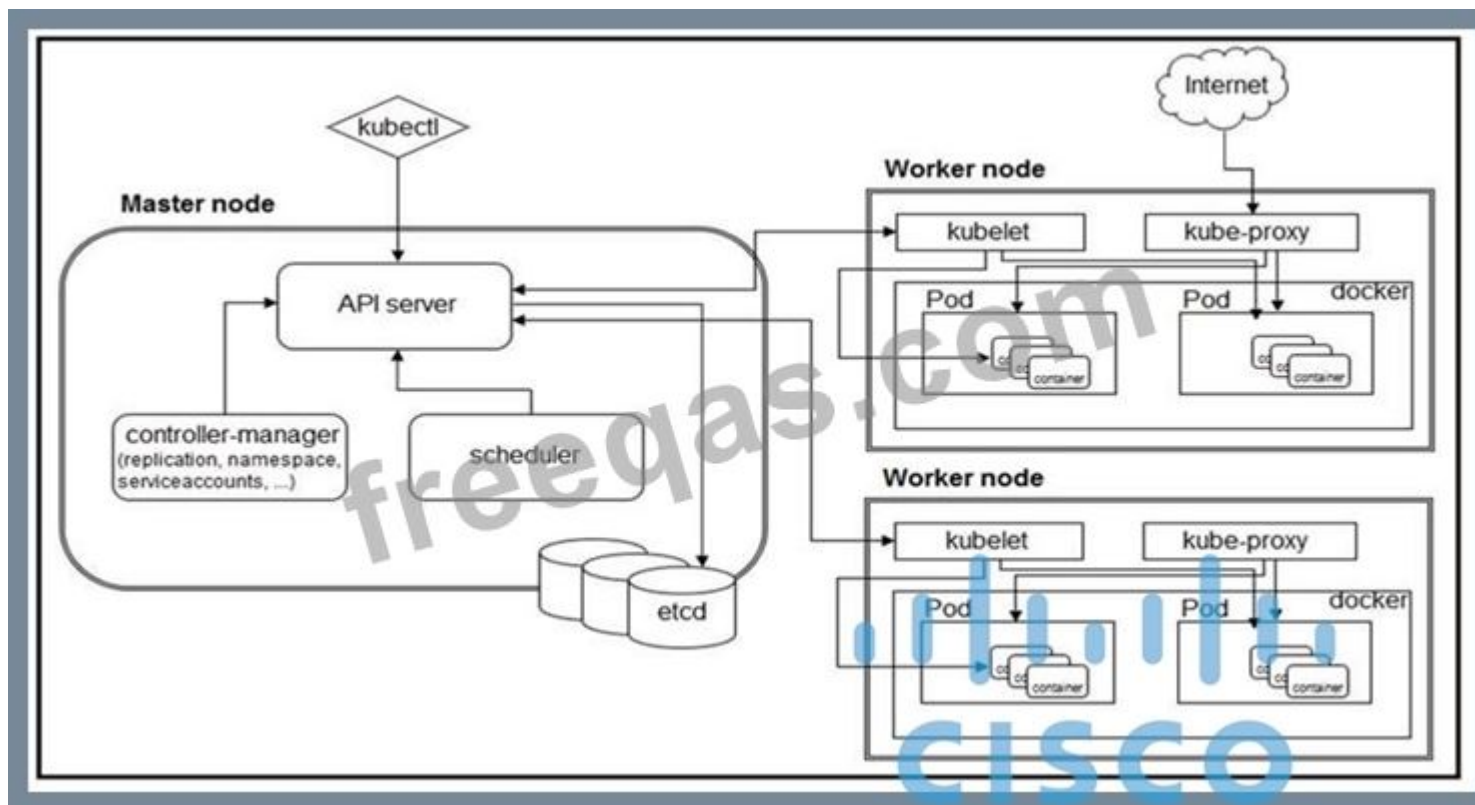
E. Download images over HTTPS supporting sites.

Answer: C,E (LEAVE A REPLY)

Section: Packaging and Delivery of Applications

NEW QUESTION: 72

Refer to the exhibit.



A developer needs to scale the existing pods within the worked nodes. Which object should be edited to achieve this goal?

A. ReplicaSet

B. PriorityClass

C. Deployment

D. Pod

Answer: A (LEAVE A REPLY)

NEW QUESTION: 73

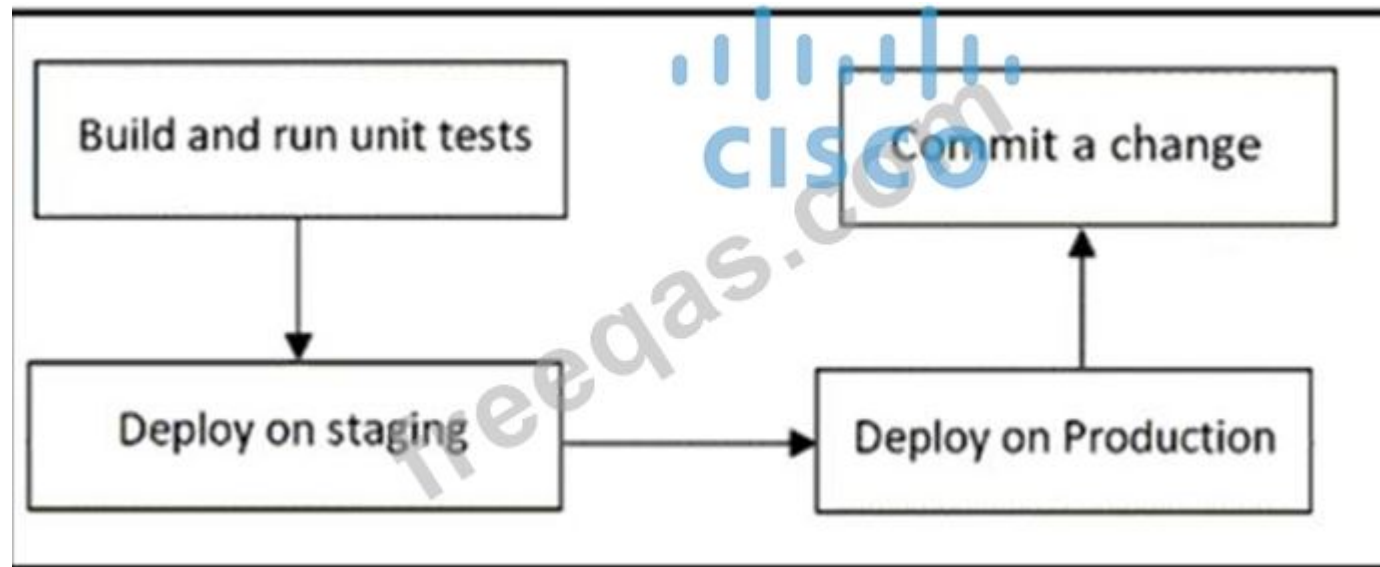
A DevOps engineer has built a new container and must open port 8080 for intercontainer communication. Which command must be added in a Dockerfile to accomplish this goal?

- A. FIREWALL ADD-PORT 8080
- B. OPEN PORT 8080
- C. PORT 8080
- D. EXPOSE 8080

Answer: D (LEAVE A REPLY)

NEW QUESTION: 74

Refer to the exhibit.



An organization has issues with code-based failures after implementing a CI/CD pipeline to automate the builds and deployment phases of an application.

Which action must be added to the pipeline, after the application is deployed in the staging environment to minimize failures and to ensure a successful continuous deployment?

- A. Restructuring and monitoring tests must be run after it is promoted to production
- B. Restructuring and monitoring tests must be run before it is promoted to production
- C. Functional and nonfunctional tests must be run after it is promoted to production
- D. Functional and nonfunctional tests must be run before it is promoted to production

Answer: D (LEAVE A REPLY)

According to Cisco's Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, it is important to run functional and non-functional tests before promoting the application to production. This is essential in order to identify and address any issues that could cause code-based failures. Additionally, a pipeline should include automated tests that can be used to verify that processes are running correctly. By running these tests before promoting to production, you can ensure that any potential issues are identified and addressed before the application is released.

NEW QUESTION: 75

A developer has created a deployment that should launch a pod to run their database service. The pod should launch with a metadata name of "Cisco-DB," and the developer has added it to the "Cisco" namespace in their deployment.

Which Kubernetes command confirms that the service is running and usable?

- A. `kubectl -n Cisco get services | grep "Cisco-DB"`
- B. `kubectl -n Cisco get pods | grep "Cisco-DB"`
- C. `kubectl get pods | grep "Cisco-DB"`
- D. `kubectl -n Cisco get service | grep "Cisco-DB"`

Answer: B (LEAVE A REPLY)

The correct answer is B. `kubectl -n Cisco get pods | grep "Cisco-DB"`. This command allows you to view the pods in the Cisco namespace and filter the list for the pod named "Cisco-DB". This will confirm that the pod is running and usable.

NEW QUESTION: 76

Which two statements about Infrastructure as Code are true? (Choose two.)

- A. Infrastructure as Code must use the same programming language as the application.
- B. Test-driven development practices make use of Infrastructure as Code.
- C. Infrastructure as Code refers to automated testing libraries.
- D. Infrastructure as Code is based on practices from software development.
- E. DevOps builds upon Infrastructure as Code.

Answer: (SHOW ANSWER)

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the **newest 300-910 exam dumps**, the PrepPdf.com 300-910 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 77

A user wants to deploy a new service to a Kubernetes cluster. Which two commands accomplish this goal?

(Choose two.)

- A. Apply
- B. Install
- C. Deploy
- D. Create
- E. Expose

Answer: A,E (LEAVE A REPLY)

Section: Cloud and Multicloud

NEW QUESTION: 78

When static routes are added to a router in a network using a CI/CD pipeline, an Ansible playbook is used to make these changes. Which steps must be added to the pipeline to validate that the changes have the intended effect on the traffic flow?

- A. Add a step to ping a host on each of the static routes before the change, and a step after to repeat the same check. Calculate the difference between the two checks for review.

- B.** Add a step to run the debug ip routing command before the change, and add a step after the change to issue the no debug ip routing command.
- C.** Add a step to run the same playbook again with the debug option enabled and use grep on the log output to ensure that the commands are not applied again.
- D.** Add a step to capture the routing table before the change, and add a step after the change to capture it again. Calculate the difference between the two for review.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 79

Which Docker command is used to start an interactive Bash shell in a running container named "test"?

- A.** docker attach -it test /bin/bash
- B.** docker run -it test /bin/bash
- C.** docker exec -it test /bin/bash
- D.** docker run test /bin/bash

Answer: C (LEAVE A REPLY)

NEW QUESTION: 80

An interface on a router as a Layer 3 link has been configured to another device by updating an Ansible playbook that is executed via a CI/CD pipeline. The Ansible playbook was written to push the configuration change using the ios_config module. Which automated test validates that the interface is functioning as it should?

- A.** Automated testing is already built in because the playbook execution is successful only after Ansible validates that it applied the change.
- B.** Add a pipeline step that runs the playbook a second type. If it does not modify the device on the second run, then the interface is correctly configured.
- C.** Add a pipeline step that captures and parses the show ip interface output so that the operational state can be referenced.
- D.** Add a pipeline step that captures and parses the packet flow for the interface.

Answer: C (LEAVE A REPLY)

Explanation

Adding a pipeline step that captures and parses the show ip interface output so that the operational state can be referenced is the best way to validate that the interface is functioning as it should. This can be done by using the ioscommand module to capture the output of the show ip interface command, and then using the iosparse module to parse the output and extract the relevant interface information. The ios_parse module allows you to validate the interface state, to ensure that it is up and running as expected. Reference: Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, Chapter 6, Section 6.3.6.

NEW QUESTION: 81

A company has a CI/CD pipeline implemented to automate application deployment. An urgent change is required in the production environment. A developer produces and tests the required change in the company's environment, but the pipeline fails when pushed. Which action conforms to the rules of a CI/CD pipeline and enables the pipeline to pass?

- A.** Manually deploy the change into the production environment
- B.** Use a backup pipeline
- C.** Replicate the production environment in the pipeline before applying

D. Remove the failing test.

Answer: [\(SHOW ANSWER\)](#)

Explanation

This ensures that the change is applied in the same environment that it was tested in, and that the pipeline will pass. This is in accordance with the rule of a CI/CD pipeline that states that the same environment should be used for testing and deploying code into production. As stated in the Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, "the same environment should be used for testing and deploying code into production, or the pipeline will fail."

NEW QUESTION: 82

Refer to the exhibit.

```
<34>1 2020-10-11T22:14:15.003Z 161.20.30.44 su - ID47 - BOM'su root' failed for DevNetUser on /dev/pts/8
```

The text represents a syslog message sent from a Linux server to a centralized log system.

Based on the format of the log message, how must the functionality of the log parser be extended to improve search capabilities?

- A. Configure the Linux machine to add a UID field to messages
- B. Filter out the text of the message to speed up searches
- C. Convert the date to the time zone of the system
- D. Reverse lookup the IP address to add a hostname field

Answer: [B \(LEAVE A REPLY\)](#)

NEW QUESTION: 83

Construct an Ansible script to gather information about target routers and then use it to apply no ip redirects to every interface on each device. Drag and drop the code from the bottom onto the correct location in the exhibit to complete the tasks section of the Ansible playbook so that it accomplishes your goal.

```

tasks:
- name: Get info from devices
  
  provider: "{{ credentials }}"
   dev_info
- name: Add 'no ip redirects' to all interfaces
  
  provider: "{{ credentials }}"

  lines:
  - "no ip redirects"
  parents: "interface {{ item.key }}"
  with_items: "{{ dev_info['ansible_facts']['ansible_net_interfaces'] | dict2items }}"

```


Answer:

```

tasks:
- name: Get info from devices
  
  provider: "{{ credentials }}"
   dev_info
- name: Add 'no ip redirects' to all interfaces
  
  provider: "{{ credentials }}"

  lines:
  - "no ip redirects"
  parents: "interface {{ item.key }}"
  with_items: "{{ dev_info['ansible_facts']['ansible_net_interfaces'] | dict2items }}"

```


NEW QUESTION: 84

Fill in the blank to complete the statement.


A user wants a Kubernetes deployment to run three separate pods of a web application at one time. In the deployment YAML, the user must configure the _____ field in the _____ subsection.

Answer:

selector, container

NEW QUESTION: 85

Refer to the exhibit.



The IT team is creating a new design for a logging system. The system must be able to collect logs from different components of the infrastructure using the SNMP protocol. When the data is collected it will need to be presented in a graphical UI to the NOC team.

What is the architectural component that needs to be placed in the unlabeled box to complete this design?

- A. time-series database
- B. log parser
- C. web server
- D. message queue

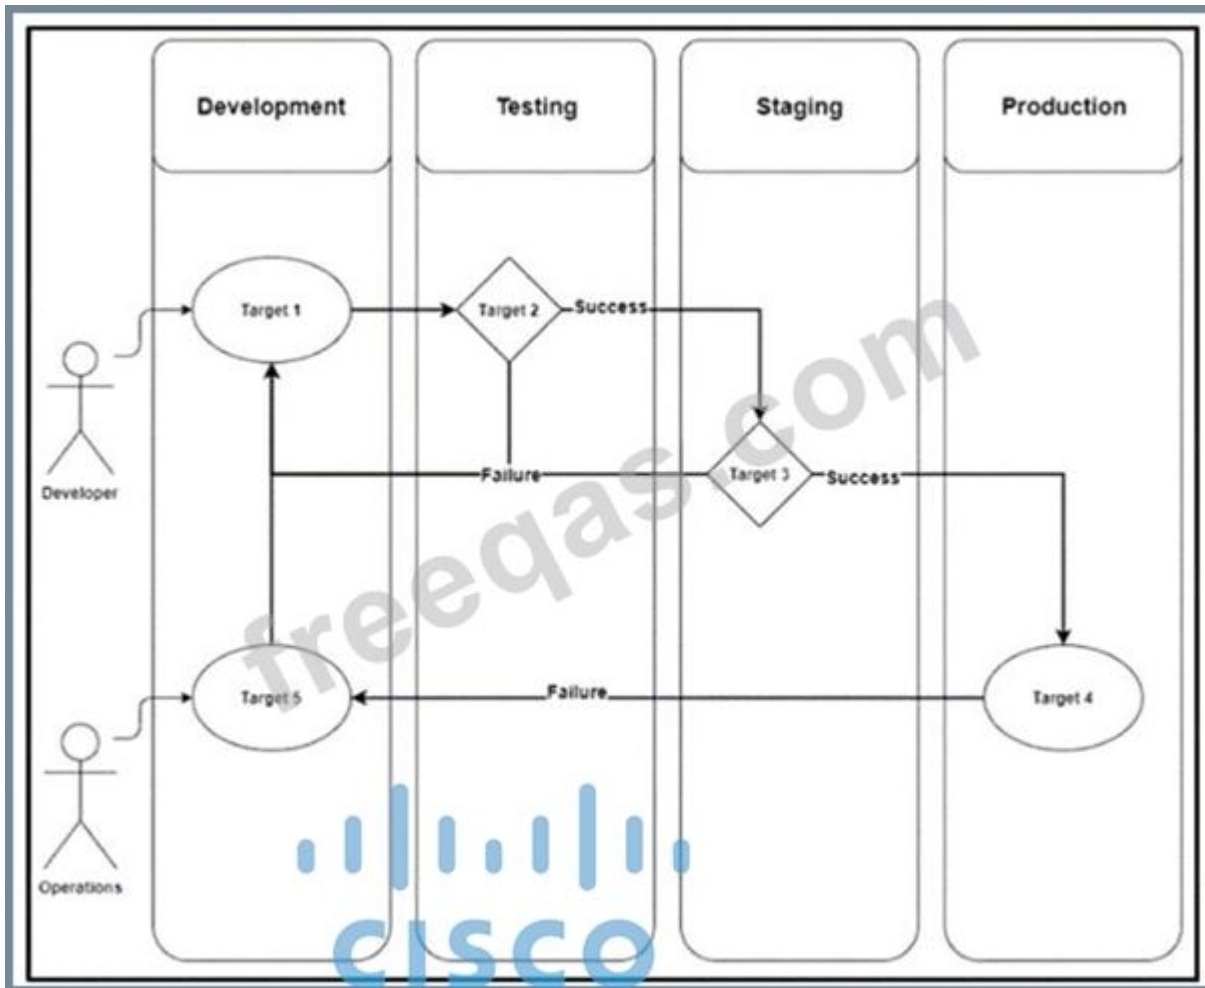
Answer: B (LEAVE A REPLY)

Explanation

A log parser is a software component that is used to collect and parse log data from various sources, such as infrastructure devices, applications, and services. It is responsible for collecting log data from the various sources and then mapping it to a specific format that can be easily consumed and interpreted by the NOC team. The log parser can also be used to filter out irrelevant data and store the processed data in a time-series database. This data can then be used by the NOC team to generate visualizations and reports, which will help them identify and address any issues within the infrastructure. (Source: Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual Chapter 5, Understanding Logging and Log Analysis)

NEW QUESTION: 86

Refer to the exhibit.



A development team is designing an application that will include multiple components and services. To streamline the process, CO/CD must be implemented.

Drag and drop the CI/CD pipeline stages from the left onto the targets on the right.



Answer:



NEW QUESTION: 87

A DevOps engineer is designing a precheck validation of the network state in a CI/CD pipeline and must implement this workflow:

- * Take a source Docker image named alpine
- * Define two stages named build and push
- * Check network connectivity before the stages run
- * Fetch the latest Docker image
- * Create a new Docker image and push it to the registry
- * Tag the new Docker image as the latest version

Drag and drop the code snippets from the bottom onto the boxes in the GitLab CI configuration to achieve the requirements for the design.

Not all options are used.

```

image: alpine
stages:
  - build
  - push

[REDACTED]:

- echo -n SCI_JOB_TOKEN | docker login -u gitlab-ci-token
--password-stdin SCI_REGISTRY

- [REDACTED]

Build:
  stage: build
  script:
    [REDACTED] --pull --build-arg VCS_REF=SCI_COMMIT_SHA
    --build-arg VCS_URL=SCI_PROJECT_URL --tag SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA
    - docker push SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA
Push:
  stage: push
  script:
    - [REDACTED] SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA

SCI_REGISTRY_IMAGE:latest
  - docker push SCI_REGISTRY_IMAGE:latest

```

docker tag

ping \$DEFAULT_GATEWAY -c 1

before_script

telnet

CONNECT 192.168.1.5000

Answer:

```
image: alpine
stages:
  - build
  - push

before_script:
  - echo -n SCI_JOB_TOKEN | docker login -u gitlab-ci-token
  --password-stdin SCI_REGISTRY
  - ping $DEFAULT_GATEWAY -c 1

Build:
  stage: build
  script:
    - docker build --pull --build-arg VCS_REF=SCI_COMMIT_SHA
    --build-arg VCS_URL=SCI_PROJECT_URL --tag SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA
    - docker push SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA
Push:
  stage: push
  script:
    - docker tag SCI_REGISTRY_IMAGE:SCI_COMMIT_SHA
    SCI_REGISTRY_IMAGE:latest
    - docker push SCI_REGISTRY_IMAGE:latest
```

```
docker tag
ping $DEFAULT_GATEWAY -c 1
before_script
telnet
```

```
before_script
ping $DEFAULT_GATEWAY -c 1
docker build
docker tag
```

When static routes are added to a router in a network using a CI/CD pipeline, an Ansible playbook is used to make these changes. Which steps must be added to the pipeline to validate that the changes have the intended effect on the traffic flow?

- A.** Add a step to run the debug ip routing command before the change, and add a step after the change to issue the no debug ip routing command.
- B.** Add a step to capture the routing table before the change, and add a step after the change to capture it again. Calculate the difference between the two for review.
- C.** Add a step to ping a host on each of the static routes before the change, and a step after to repeat the same check. Calculate the difference between the two checks for review.
- D.** Add a step to run the same playbook again with the debug option enabled and use grep on the log output to ensure that the commands are not applied again.

Answer: B (LEAVE A REPLY)

Section: Automating Infrastructure

NEW QUESTION: 89

An interface on a router as a Layer 3 link has been configured to another device by updating an Ansible playbook that is executed via a CI/CD pipeline. The Ansible playbook was written to push the configuration change using the ios_config module. Which automated test validates that the interface is functioning as it should?

- A.** Add a pipeline step that runs the playbook a second time. If it does not modify the device on the second run, then the interface is correctly configured.
- B.** Automated testing is already built in because the playbook execution is successful only after Ansible validates that it applied the change.
- C.** Add a pipeline step that captures and parses the packet flow for the interface.
- D.** Add a pipeline step that captures and parses the show ip interface output so that the operational state can be referenced.

Answer: B (LEAVE A REPLY)

Adding a pipeline step that captures and parses the show ip interface output so that the operational state can be referenced is the best way to validate that the interface is functioning as it should. This can be done by using the ioscommand module to capture the output of the show ip interface command, and then using the iosparse module to parse the output and extract the relevant interface information. The ios_parse module allows you to validate the interface state, to ensure that it is up and running as expected. Reference: Cisco Implementing DevOps Solutions and Practices using Cisco Platforms (DEVOPS) Study Manual, Chapter 6, Section 6.3.6.

NEW QUESTION: 90

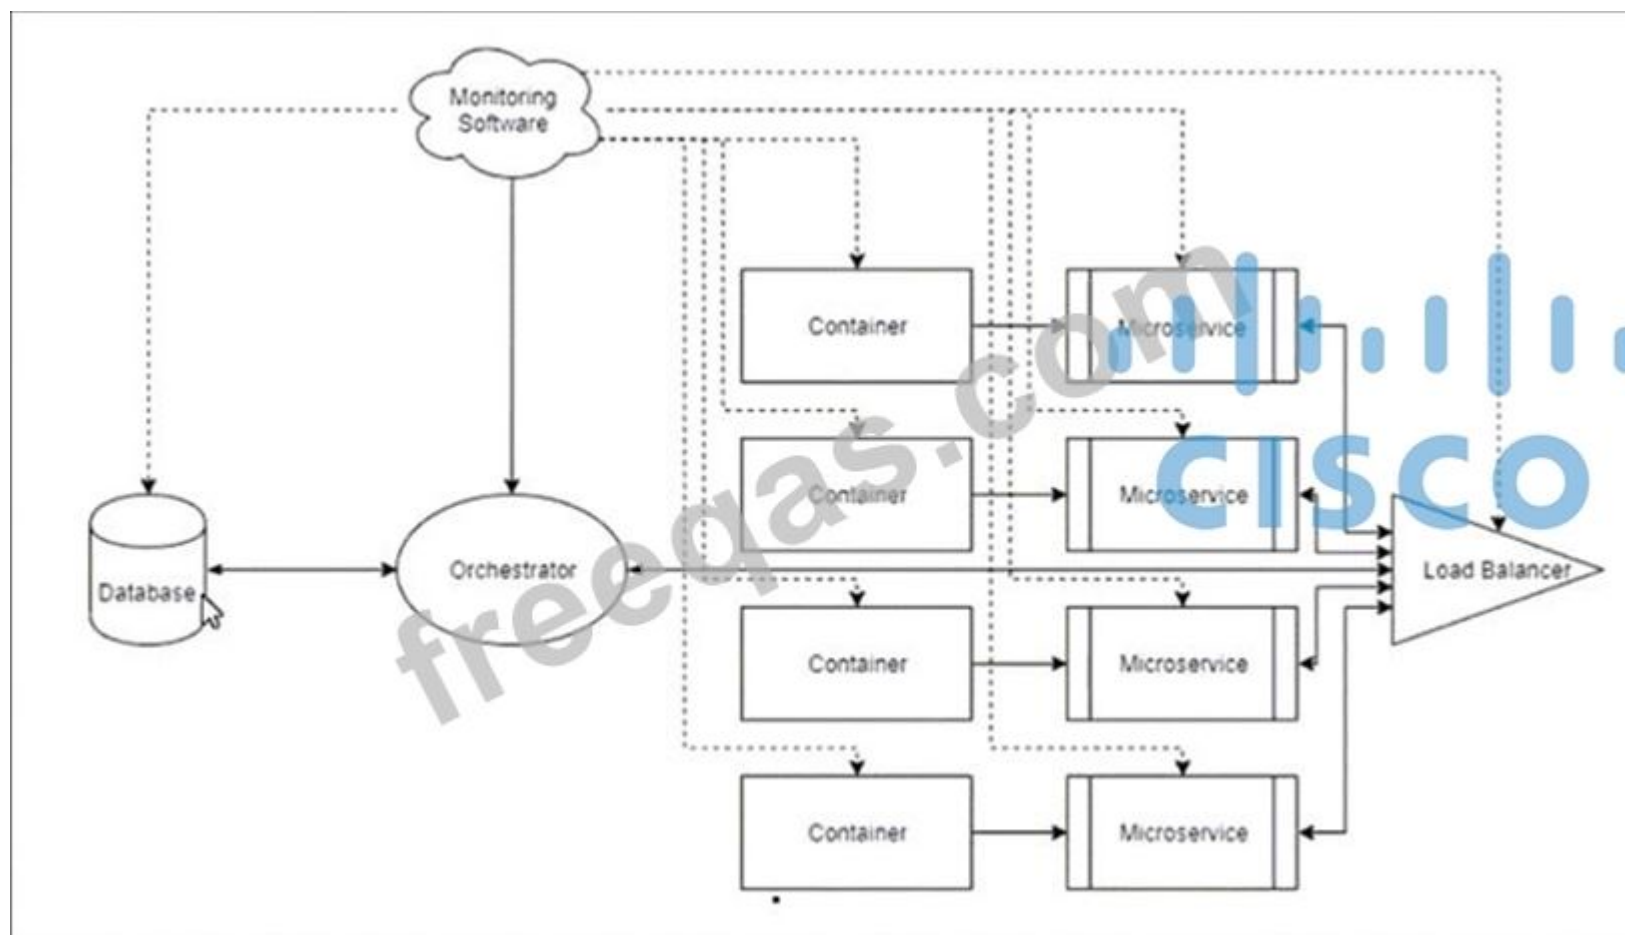
What are two testing scenarios of the chaos engineering principle? (Choose two.)

- A.** removing all users from a version control system
- B.** blocking developers' building access
- C.** unplugging a core switch device
- D.** maxing out CPU cores on an Elasticsearch cluster
- E.** executing routine in driver code to emulate I/O errors

Answer: C,D (LEAVE A REPLY)

NEW QUESTION: 91

Refer to the exhibit.



How does the architecture respond to a usage spike?

- A. The Orchestrator increases the resources assigned to containers to increase performance
- B. The load balancer interacts with the Orchestrator to optimize the balancing algorithm to match the incoming flow of requests
- C. The monitoring software instructs the Orchestrator to increase the count of containers and to configure the load balancer accordingly
- D. The database prioritizes quick and immediate queries over complicated ones to reduce the request queue

Answer: B (LEAVE A REPLY)

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the **newest 300-910 exam dumps**, the PrepPdf.com 300-910 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 92

Fill in the blanks to complete the line of Python code that sends a message to a Webex Teams room or person.

```
response = requests. [ ] ("https://api.ciscopark.com/v1/ [ ] ", headers=message_header, data=message_data)
```

Answer:

post, messages

NEW QUESTION: 93

Microservices architecture pattern has been applied and the system has been architected as a set of services.

Each service is deployed as a set of instances for throughput and availability. In which two ways are these services packaged and deployed? (Choose two.)

- A. Service instances must be isolated from one another.
- B. Service must be independently deployable and scalable.
- C. Service are written using the same languages, frameworks, and framework versions.
- D. Service must be dependent, deployable, and scalable.
- E. Service instances do not need to be isolated from one another.

Answer: ([SHOW ANSWER](#))

Section: Packaging and Delivery of Applications

NEW QUESTION: 94

Which Dockerfile produces an efficient image rebuild when the exposed port Python dependency, or MyApp source code is modified?

A)

```
FROM python
EXPOSE 8082
WORKDIR /home/app
COPY ./my-app/requirements.txt /home/app/requirements.txt
RUN pip install -r requirements.txt
COPY ./my-app/ /home/app
```

B)

```
FROM python:3.7.4-buster
EXPOSE 8082
WORKDIR /home/app
COPY ./my-app/ /home/app
RUN pip install -r requirements.txt
```

C)

```
FROM python:latest
WORKDIR /home/app
RUN pip install -r requirements.txt
COPY ./my-app/ /home/app
EXPOSE 8082
```

D)

```
FROM python:3.7.4-buster
WORKDIR /home/app
COPY ./my-app/requirements.txt /home/app/requirements.txt
RUN pip install -r requirements.txt
COPY ./my-app/ /home/app/
EXPOSE 8082
```

- A. Option B
- B. Option A
- C. Option C
- D. Option D

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 95

FILL BLANK

Fill in the blanks to complete the statement.

When creating multiple containers, in order to allow communication with each other, you must create a _____ of type _____ .

Answer:

pod, volume

Section: Packaging and Delivery of Applications

NEW QUESTION: 96

Refer to the exhibit.

```
json_data = """
{
  "matchers": [
    {
      "name": "alertname",
      "value": "DevNetExampleAlert",
      "isRegex": false
    },
    {
      "name": "job",
      "value": "core",
      "isRegex": false
    }
  ],
  "startsAt": "2019-07-10T23:42:37.565Z",
  "endsAt": "2019-07-11T01:42:37.565Z",
  "createdBy": "Cisco DevNet",
  "comment": "Silencing this alert while working on it.",
  "id": null
}
"""
```

The snippet contains the JSON string that will be sent to the Prometheus AlertManager to silence an alert. AlertManager accepts only the content type of application 'json' Which Python code correctly sends an API call to perform action?

- A. response= request.post(url. data=json_data)
- B. response= request.post(url. json=json.loads(json_data))
- C. response= request.post(url. json=json_data)
- D. response= request.post(url, data=json=(json_data))

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 97

Refer to the Exhibit.



Answer:



NEW QUESTION: 98

Refer to the exhibit.

Build Command

```
#!/bin/bash
# apt-get -y install python3-pip
# pip install --upgrade pip
rm -rf automationSandboxTest || true
git clone https://github.com/oborys/automationSanboxTest.git
export SEARCHPATH="$PWD/automationSanboxTest/"
for k in $(cat $SEARCHPATH/requiriemnts.txt | cut -d '>' -f 1 | cut -d '<' -f 1
| cut -d '=' -f 1
do
    python -m pip install $k
done
echo
for k in $(find $SEARCHPATH -name *.py)
do
    echo > msg.txt
    python $k > /dev/null 2> msg.txt || true
    export CODE=$(grep -c ^ msg.txt)
    if [ $CODE != 0 ]
    then
        echo "File: $k" >> error.message.txt
        cat msg.txt >> error_message.txt
        echo >> error.message.txt
    fi
done
rm -rf msg.txt || true
cat error_message.txt
if [ $(cat error_message.txt| wc -l) != 0 ]
then
    exit 1
fi
```



```
++find /data/bms/webapps/jenkins/workspace/team_team_devnet-learning-labs-
automation/Always_On_Sandbox_testing/automationSandboxTest/ -name '*.py'
+ for k in `$(find $SEARCHPATH -name *.py)`
+ echo
+ python /data/bms/webapps/jenkins/workspace/team_team_devnet-learning-labs-
automation/Always_On_Sandbox_testing/automationSandboxTest/alwaysOnSandboxCh
eck.py
++ grep -c '^' msg.txt
+ export CODE=0
+ CODE=0
+ '[' 0 '!=' 0 ']'
+ rm -rf msg.txt
+ cat error_message.txt
cat: error_message.txt: No such file or directory
Build step 'Virtualenv Builder' marked build as failure
Notifying upstream projects of job completion
Finished: FAILURE
```

Part of Console Output

How should the Jenkins job be troubleshooted based on the error provided?

- A. Place the code in a container and run the job again.
- B. Install dependencies.
- C. Verify what the responding file created.
- D. Update pip.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 99

Which two practices help make the security of an application a more integral part of the software development lifecycle? (Choose two.)

- A. Ensure that the code repository server has enabled drive encryption and stores the keys on a Trusted Platform Module or Hardware Security Module.
- B. Add a step to the CI/CD pipeline to modify the release plan so that updated versions of the software are made available more often.
- C. Use only software modules that are written by the internal team.
- D. Add a step to the CI/CD pipeline that runs a static code analysis tool during the pipeline execution.
- E. Add a step to the CI/CD pipeline that runs a dynamic code analysis tool during the pipeline execution.

Answer: (SHOW ANSWER)

Explanation

Adding a step to the CI/CD pipeline that runs a dynamic code analysis tool and a static code analysis tool during the pipeline execution helps make the security of an application a more integral part of the software development lifecycle. Dynamic code analysis tools search for coding errors and vulnerabilities while the application is running, while static code analysis tools scan the source code for potential errors and vulnerabilities. This ensures that any security issues are identified and addressed before the application is deployed. Additionally, using automated tools helps to reduce the amount of manual effort required for security testing and can reduce the risk of security flaws being introduced.

NEW QUESTION: 100

Get Site Health

Operation Id: `getSiteHealth`

Description: Returns Overall Health information for all sites

GET /dna/intent/api/v1/site-health

Request Parameters

Query

timestamp | String

epoch time(in milliseconds) when the Site hierarchy is required

Get Overall Network Health

Operation Id: `getOverallNetworkHealth`

Description: Returns Overall Network Health information by Device category (Access, Distribution, Core, Router, Wireless) for any given point of time

GET /dna/intent/api/v1/network-health

Request Parameters

Query

timestamp | String

epoch time(in milliseconds) when the Network health data is

Refer to the exhibit. A developer is creating a health check monitoring script that queries information from the Cisco DNA Center platform. The script must trigger an alert if a site health statistic named `accessGoodCount` drops below 80 and if a network statistic named `latestHealthScore` is 95 or less.

Drag and drop the code snippets from the bottom onto the blanks in the code to monitor the site and network health on a Cisco DNA Center platform instance. Options may be used more than once. Not all options are used.

```

BASE_URL = 'https://sandboxdnac.cisco.com'
NETWORK_HEALTH_URL = '/dna/intent/api/v1/network-health'
SITE_HEALTH = '/dna/intent/api/v1/site-health'
timestamp = datetime.timestamp()
data = {
    'X-Auth-Token': "asfds"
}
info = {
    [redacted] : timestamp
}
while True:
    [redacted]
    response = requests.request('GET', url,
    headers=data, [redacted] =info)
    if response.json()[0]['accessGoodCount'] < 80:
        trigger_site_alert()
        [redacted]
    response = requests.request('GET', url,
    headers=data, [redacted] =info)

```

url = BASE_URL + SITE_HEALTH

params

url = BASE_URL + NETWORK_HEALTH_URL

'query'

"info"

'timestamp'

Answer:

```

BASE_URL = 'https://sandboxdnac.cisco.com'
NETWORK_HEALTH_URL = '/dna/intent/api/v1/network-health'
SITE_HEALTH = '/dna/intent/api/v1/site-health'
timestamp = datetime.timestamp()
data = {
    'X-Auth-Token': "asfds"
}
info = {
    'query' : timestamp
}
while True:
    url = BASE_URL + SITE_HEALTH
    response = requests.request('GET', url,
    headers=data, params = info)
    if response.json()[0]['accessGoodCount'] < 80:
        trigger_site_alert()
    url = BASE_URL + NETWORK_HEALTH_URL
    response = requests.request('GET', url,
    headers=data, params = info)

```




NEW QUESTION: 101

Refer to the exhibit.

```
kind: pipeline
name: test1
platform:
  os: linux
  arch: amd64
steps:
- name: test
  image: postgres:9-alpine
  commands:
  - sleep 10
  - psql -U postgres -d test -h database -c "SELECT version();"
services:
- name: database
  image: postgres
  environment:
    POSTGRES_DB: test
    POSTGRES_USER: postgres
---
```

kind: pipeline
name: test2
platform:
 os: linux
 arch: amd64
steps:
- name: test
 image: postgres:9-alpine
 commands:
 - sleep 10
 - psql -U postgres -d test -h database -c "SELECT version();"
services:
- name: database
 image: postgres
 environment:
 POSTGRES_DB: test
 POSTGRES_USER: postgres



What is the user doing with Drone in this automated test?

- A. testing only the amd64 architecture
- B. testing a PostgreSQL DB against multiple architectures
- C. testing Alpine Linux versus Ubuntu Linux
- D. testing PostgreSQL deployment in an Alpine Linux VM

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 102

Refer to the exhibit.

```
1 {
2   "_index": "linux-servers",
3   "_type": "_doc",
4   "_id": "mr25OXQBMiQCMA9yzEQ6",
5   "_version": 1,
6   "_score": null,
7   "_source": {
8     "message": "<38>Aug 29 12:08:37 centos systemd-logind: Removed session
9     19802.",
10    "@timestamp": "2020-03-29T09:08:37.280Z",
11    "type": "syslog",
12    "@version": "1",
13    "hostname": "vcs.local.lan",
14    "host": "172.16.40.168"
15  },
16  "fields": {
17    "@timestamp": [
18      "2020-08-29T12:03:37.280Z"
19    ]
20  },
21  "sort": [
22    1598692117280
23  ]
24 }
```

The JSON object represents a single entry on a centralized log server, but log data cannot be processed because of the format.

What causes the issue?

- A. The process name in the message must be parsed into a field
- B. A hostgroup must be defined
- C. The "_type" must represent the process type
- D. The priority of the message must be to the server

Answer: C (LEAVE A REPLY)

NEW QUESTION: 103

What are two advantages of using Configuration Management Tools? (Choose two.)

- A. reduction in policy violations caused by human errors
- B. reduction in administration costs
- C. reduction in network changes already performed automatically
- D. reduction of on-premises networking equipment
- E. reduction in networking team skills

Answer: B,C (LEAVE A REPLY)

Section: Logging, Monitoring, and Metrics

NEW QUESTION: 104

A development team uses Kubernetes for application development. Any changes on ConfigMap are performed manually for each development, test, and production environment. The edits are performed to deploy applications. This approach causes inconsistent deployments across all environments.

Which practice improves the consistency of the deployments?

- A.** In the master branch where the ConfigMaps are stored, create a branch for each environment that contains an environment-specific ConfigMap.
- B.** Generate the ConfigMaps specific to the environment by using a templating language such as Jinja2 and store the ConfigMaps in unique branches of a repository
- C.** Create a unique repository for each environment that contains ConfigMaps for that environment to ensure that each environment can be deployed independently
- D.** Implement environment variables within the ConfigMaps and store the variable definitions separately from the master branch where the ConfigMaps are stored

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 105

Microservices architecture pattern has been applied and the system has been architected as a set of services.

Each service is deployed as a set of instances for throughput and availability. In which two ways are these services packaged and deployed? (Choose two.)

- A.** Service are written using the same languages, frameworks, and framework versions.
- B.** Service instances do not need to be isolated from one another.
- C.** Service must be dependent, deployable, and scalable.
- D.** Service must be independently deployable and scalable.
- E.** Service instances must be isolated from one another.

Answer: ([SHOW ANSWER](#))

Valid 300-910 Dumps shared by PrepPdf.com for Helping Passing 300-910 Exam! PrepPdf.com now offer the **newest 300-910 exam dumps**, the PrepPdf.com 300-910 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com 300-910 dumps with Test Engine here: <https://www.preppdf.com/Cisco/300-910-prepaway-exam-dumps.html> (130 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)