UiPath-IEPASv1 Dumps

Shared by Byrd on 20-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which statement correctly explains the process and constraints for adding nodes to an Automation Suite deployment?

Options:

A- Adding nodes is supported in both deployment profiles, but adding nodes to a single-node environment does not convert it into an HA-ready setup.
B- Adding nodes to a single-node evaluation profile will automatically convert it to a multi-node HA-ready production environment.
C- To add a new node, you must reconfigure the deployment profile and regenerate a new cluster configuration from scratch.
D- You can only add agent nodes post-installation; server nodes must be defined before the initial deployment.

Answer:

A

Explanation:

UiPath Automation Suite supports adding both server and agent nodes post-installation to scale compute and memory capacity across either single-node evaluation or multi-node production deployment profiles. However, adding nodes to a single-node evaluation cluster does not convert it into a multi-node High Availability (HA)-ready production deployment. Converting a single-node environment into a certified HA-ready setup requires specific architectural changes: migrating in-cluster databases to external SQL instances, configuring external object stores, provisioning a zone-redundant load balancer, deploying an odd number of server nodes (minimum three) to establish etcd quorum, and reconfiguring storage replication. Adding nodes to a single-node cluster merely provides additional compute capacity for specialized workloads (such as attaching a GPU agent node) while leaving the single control plane as a single point of failure. (Reference: UiPath Automation Suite Administration -- Cluster Scaling and Node Addition Constraints).

Question 2

Question Type: MultipleChoice

In a deployment across three availability zones, which tool can be used to route traffic to any of the load balancers according to a specified algorithm?

Options:

A- Traffic Manager
B- FluentD
C- ArgoCD
D- Prometheus

P2P exams

Answer:

A

Explanation:

When deploying UiPath Automation Suite across multi-zone or multi-region topologies utilizing independent regional load balancers, Azure Traffic Manager (or AWS Route 53) provides global DNS-based traffic routing. Azure Traffic Manager operates at the DNS transport layer, evaluating endpoint health and directing client connections to specific load balancer IP addresses based on configured routing methods. Supported algorithms include priority routing for active-passive failover, weighted round-robin for distributing loads proportionally, and geographic routing for latency optimization. The other listed tools serve distinct operational functions: FluentD aggregates cluster logs, ArgoCD automates declarative GitOps platform deployments, and Prometheus collects time-series monitoring metrics. Traffic Manager provides DNS-level endpoint health checking and routing algorithms for multi-load balancer architectures. (Reference: UiPath Automation Suite Cloud Architecture Guide -- Multi-Zone DNS and Traffic Management).

==========

P2P exams

Question 3

Question Type: MultipleChoice

When setting up the Load Balancer for Automation Suite, which types of nodes must be included in the Node Pool?

Options:

A- Only specialized agent nodes only.
B- Only server nodes.
C- Server nodes and non-specialized agent nodes.
D- Only agent nodes but not Specialized agent nodes.

Answer:

C

Explanation:

When provisioning the enterprise reverse-proxy load balancer for UiPath Automation Suite, the backend Node Pool must contain all server nodes and all non-specialized agent nodes. In an Automation Suite cluster, the ingress service mesh gateway (Istio) and application pods run distributed across the server nodes and standard worker (non-specialized agent) nodes. Including both node types in the load balancer pool ensures incoming traffic on ports 443 and 80 is balanced across all nodes hosting ingress controllers and application services. Conversely, specialized agent nodes---such as dedicated GPU nodes for Document Understanding and dedicated nodes for Automation Suite Robots---are tainted to prevent generic pod scheduling and must be excluded from the load balancer pool. Routing external web traffic to specialized nodes causes request drops because those nodes do not host ingress routing components. (Reference: UiPath Automation Suite Network and Load Balancing Configuration -- Backend Pool Composition).

==========

Question 4

Question Type: MultipleChoice

Which platforms are supported for deploying Automation Suite?

Options:

A- Windows Server
B- Any Kubernetes-based containerization platform
C- Only OpenShift
D- AKS, EKS, Openshift and RHEL Linux based clusters

Answer:

D

Explanation:

UiPath Automation Suite provides flexible enterprise deployment options across supported environments. The officially supported deployment platforms encompass: self-hosted Red Hat Enterprise Linux (RHEL) or compatible Linux-based clusters (leveraging the embedded RKE2 Kubernetes engine), Azure Kubernetes Service (AKS), Amazon Elastic Kubernetes Service (EKS), and Red Hat OpenShift Container Platform. Microsoft Windows Server is not supported for Automation Suite; Windows is used only for legacy standalone service installations. Furthermore, UiPath does not support arbitrary, uncertified Kubernetes distributions (such as generic K3s, Minikube, or unmanaged community Kubernetes). Each supported target platform undergoes extensive compatibility validation for storage integration, CNI networking, ingress controllers, and Security Context Constraints. Enterprise infrastructure teams must align with these four certified environments to receive official support. (Reference: UiPath Automation Suite Installation Guide -- Supported Target Platforms and Distribution Matrix).

==========

Question 5

Question Type: MultipleChoice

What is the minimum processor requirement for each cluster in a multi-node deployment?

Options:

A- 32 (v-)CPU/cores
B- 64 (v-)CPU/cores
C- 96 (v-)CPU/cores
D- 128 (v-)CPU/cores

Answer:

C

Explanation:

In an enterprise multi-node High Availability (HA)-ready deployment of UiPath Automation Suite, the minimum baseline aggregate processor requirement is 96 (v-)CPU cores per cluster. A baseline multi-node production deployment requires at least three server nodes to establish

control-plane quorum, along with accompanying agent worker nodes to host platform microservices, core shared dependencies (e.g., Kafka, Redis, RabbitMQ), and data processing pipelines. Sizing standards require each server node to be provisioned with a minimum of 16 to 32 vCPUs, resulting in an aggregate capacity requirement of at least 96 vCPU cores across the cluster nodes. This compute capacity ensures adequate headroom for control-plane stability, pod scheduling, and burst processing during high automation workloads. Clusters provisioned below this 96-core threshold risk CPU throttling and pod scheduling failures. (Reference: UiPath Automation Suite Sizing Guide -- Multi-Node Production Sizing Baselines).

==========

Question 6

Question Type: MultipleChoice

How many HAA licenses are needed to configure an Active/Passive deployment of Automation Suite?

Options:

A- Two licenses, one for each shard.
B- One license for each cluster.
C- Only one license is sufficient.
D- Two licenses, one for each cluster.

Answer:

D

Explanation:

Configuring an Active/Passive multi-site disaster recovery deployment of UiPath Automation Suite requires two separate High Availability Add-on (HAA / Redis Enterprise) licenses: one for each participating cluster. In an Active/Passive topology, both the primary production cluster and the standby secondary cluster run independent instances of the Redis High Availability Add-on to provide session caching, state replication, and message brokering. Because each cluster operates as an independent Kubernetes deployment with its own control plane and network boundary, an independent license must be activated on each site. UiPath technical specifications dictate that each cluster must be licensed for a minimum of two shards. Attempting to share a single HAA license across clusters or omitting licensing on the standby site prevents cluster synchronization and breaks failover readiness. (Reference: UiPath Automation Suite Disaster

Recovery Architecture -- HAA Redis Licensing Model).

==========

Question 7

Question Type: MultipleChoice

Which of the following statements is accurate when deploying the Orchestrator service in Automation Suite?

Options:

A- Only SQL server Enterprise editions are supported
B- Requires SQL Server version 2019 or higher for optimal performance
C- Requires a dedicated Microsoft SQL Server 2016, 2017, 2019, and 2022 Standard and Enterprise editions
D- Orchestrator does not require an SQL server because it is deployed as a service as part of Automation Suite

Answer:

B

Explanation:

When provisioning the Orchestrator service within UiPath Automation Suite, external Microsoft SQL Server 2019 or higher is required for optimal performance and full feature compatibility. Orchestrator acts as the central transaction hub for job scheduling, queue management, asset security, and robot execution logs, generating high transactional volumes. The modern database schemas utilized by containerized Orchestrator leverage SQL Server 2019 engine enhancements, including optimized row-level locking, intelligent query processing, and native JSON functions. UiPath supports both Microsoft SQL Server Standard and Enterprise editions, but requires version 2019 or newer to ensure performance and prevent schema compatibility issues. Orchestrator cannot operate without an external SQL Server, but it does not mandate a physically dedicated SQL instance, allowing it to share database instances with other compatible platform services. (Reference: UiPath Automation Suite Database Requirements -- Orchestrator Database Prerequisites).

==========

Question 8

Question Type: MultipleChoice

What minimum hardware specifications are recommended for a Microsoft SQL Server that hosts Automation Suite databases?

Options:

A- 4 vCPU, 16 GB RAM, 128 GB SSD
B- 6 vCPU, 24 GB RAM, 200 GB SSD
C- 8 vCPU, 32 GB RAM, 256 GB SSD
D- 10 vCPU, 40 GB RAM, 512 GB SSD

Answer:

C

Explanation:

External relational databases play a critical role in state persistence for UiPath Automation Suite. The official infrastructure specifications recommend a hardware baseline of 8 (v-)CPUs, 32 GB of RAM, and a 256 GB enterprise SSD for a dedicated Microsoft SQL Server hosting platform databases. UiPath Automation Suite consolidates relational data for multiple services---including Orchestrator, Identity Server, Action Center, Data Service, and Test Manager---onto this SQL database engine. This hardware sizing ensures the database engine can maintain adequate buffer cache memory, execute complex transactional queries, manage parallel index maintenance, and provide sustained I/O throughput with low latency. Sizing below these specifications can lead to high database CPU utilization, memory pressure, connection pool exhaustion, and performance bottlenecks under peak automation loads. (Reference: UiPath Automation Suite Database Requirements -- External Microsoft SQL Server Sizing).

==========

Question 9

Question Type: MultipleChoice

What needs to be done to enable Handwriting recognition for an online installation?

Options:

A- The parameter should be configured in the cluster_config.json.
B- Handwriting is always enabled for online installation.
C- It can be enabled from the AI Center interface.
D- It can be enabled by calling the Document Understanding API endpoint.

Answer:

B

Explanation:

In UiPath Automation Suite deployments where Document Understanding is provisioned, Handwriting Recognition capabilities for OCR and digitization are automatically enabled out-of-the-box for all online installations. Because online clusters possess outbound internet connectivity to UiPath public cloud infrastructure and secure model endpoints, the platform automatically downloads, registers, and exposes the handwriting recognition ML pipelines without requiring manual administrative configuration. Additional administrative intervention---such as manually downloading offline ML package bundles, staging model tarballs, and declaring configuration parameters in `cluster_config.json`---is required only for air-gapped or offline deployment modes. In an online environment, the handwriting recognition capability is initialized automatically as part of the core Document Understanding service rollout. (Reference: UiPath Automation Suite Document Understanding Guide -- Feature Enablement and OCR Configuration).

==========

Question 10

Question Type: MultipleChoice

Which parameter adds machines as server nodes to an existing Automation Suite cluster?

Options:

A- -k
B- -j server
C- -j
D- -k server

Answer:

B

Explanation:

During the automated scaling or manual node join workflow of a UiPath Automation Suite deployment on Linux, the installation shell script (`configureUiPathAS.sh` or `install-uipath.sh`) utilizes the `-j` parameter to define the node role during cluster join operations. To join an additional machine as a Kubernetes control-plane server node, the flag `-j server` must be explicitly specified. The `-j` flag indicates a 'join' action, and the `server` argument instructs the installer to register the machine as an RKE2 control-plane member, joining it to the existing etcd consensus ring and deploying the core Kubernetes management services. Omitting the `server` argument or passing `-j agent` joins the host as a standard worker node instead. The `-k` flag is reserved strictly for executing prerequisite validation checks and cannot be combined with node-joining parameters. (Reference: UiPath Automation Suite Node Operations -- Adding Server Nodes via CLI).

==========

To Get Premium Files for UiPath-IEPASv1 Visit
https://www.p2pexams.com/products/uipath-iepasv1

For More Free Questions Visit
https://www.p2pexams.com/uipath/pdf/uipath-iepasv1

