

DUMPS ARENA

Salesforce Certified Tableau Architect

Salesforce Analytics-Arch-201

Version Demo

Total Demo Questions: 22

Total Premium Questions: 220

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Data Strategy	1
Topic 2, Data Governance	13
Topic 3, Data Architecture	5
Topic 4, Tableau Server Architecture	55
Topic 5, Tableau Server Implementation	146
Total	220

QUESTION NO: 1

When creating a custom administrative view to monitor user activity in Tableau Server, which table in the Tableau repository schema should you focus on to interpret login and logout events?

- A. The 'workbooks' table to track user interactions with different workbooks
- B. The 'data_connections' table to monitor which data sources are being accessed
- C. The 'http_requests' table to analyze web requests made by user
- D. The 'historical_events' table to track specific user login and logout activities

ANSWER: D

Explanation:

The `historical_events` table to track specific user login and logout activities is correct because Tableau Server records auditable, time-based server activity as historical events. This makes it the appropriate repository area for an administrative view intended to understand authentication activity, including when users establish or end sessions. A useful custom view can use the event timestamp and associated user or actor information to identify access patterns, investigate activity during a particular period, and support operational or security reviews. In practice, repository queries commonly enrich these event records by joining related lookup tables, such as event-type metadata and user information, so the final visualization presents readable event names and user identities rather than only internal identifiers. Tableau documents the Tableau Server Repository as the source for creating custom administrative views and provides its schema reference, including historical event data. Because repository schemas and access methods can vary by Tableau Server release, query design should be validated against the documented schema for the deployed version. See Tableau's [Work with the Tableau Server Repository](#) guidance and its [Tableau Server Repository Data Dictionary](#).

QUESTION NO: 2

To automate the process of refreshing extracts on Tableau Server, which tool or API would be most effective for scheduling this task?

- A. Utilizing webhooks to trigger extract refreshes based on user interactions within the server
- B. Employing `tabcmd` to schedule and execute extract refreshes at specified intervals
- C. Using the Hyper API to manually initiate extract refreshes each time data changes
- D. Relying on manual refreshes through the Tableau Server web interface for precise control
- E. Using the Tableau Server REST API to create a schedule and associate extract-refresh tasks with it

ANSWER: E

Explanation:

Using the Tableau Server REST API to create a schedule and associate extract-refresh tasks with it is the most effective approach because Tableau Server schedules are the platform-native mechanism for recurring extract refresh automation. Through the REST API, an administrator or automation service can programmatically create and manage schedules, then add published data sources or workbooks to those schedules. Tableau Server runs the resulting refresh tasks according to the configured frequency, eliminating the need for a separate operating-system scheduler or manual intervention. This approach is especially useful when schedules must be provisioned consistently across sites, integrated into deployment pipelines, or managed from an external administrative application. The API also supports authenticated, repeatable administration, allowing a process to inspect and maintain the Server configuration as requirements change. Tableau documents schedule-management endpoints, including creating schedules and associating content with schedules, in its REST API reference. For general operational guidance on configuring refresh schedules, see Tableau's documentation on extract refresh tasks and schedules. [Tableau Server REST API Reference](#) and [Create or Edit a Schedule](#).

QUESTION NO: 3

An enterprise is merging its multiple Tableau sites into a single server for better management and efficiency. What should be the primary focus during this migration?

- A. Rapidly migrating all sites without a detailed review to accelerate the process
- B. Ensuring the compatibility and proper configuration of data connections across the merged sites
- C. Transferring only the most used dashboards and reports, disregarding less frequently used content
- D. Maintaining multiple backup servers in case the consolidation fails

ANSWER: B

Explanation:

Ensuring the compatibility and proper configuration of data connections across the merged sites is the primary focus because a Tableau consolidation succeeds only if published content can continue to reach its required data reliably after it is moved. Workbooks, published data sources, virtual connections, and flows can depend on embedded connection details, published data-source references, credentials, authentication methods, network access, drivers, and refresh schedules. These dependencies must be inventoried and validated in the target site so that users retain access to current data and scheduled operations continue as intended. The migration team should also confirm that the target environment has the necessary connectors and database drivers, that service-account or delegated credentials are configured appropriately, and that data-source permissions align with the consolidated site's governance model. Tableau's migration guidance specifically emphasizes planning, assessing content and dependencies, and testing migrated content before broad rollout. Tableau also documents that connections and credentials are central considerations when publishing and managing data sources. Addressing these details first protects the usability and trustworthiness of dashboards after consolidation. See Tableau's [migration documentation](#) and its guidance on [data-source connections](#).

QUESTION NO: 4

When planning to implement automated user provisioning for Tableau Cloud, how can the System for Cross-Domain Identity Management (SCIM) be effectively utilized?

- A. By manually updating user roles in Tableau Cloud whenever there are changes in the organization 's identity management system
- B. Integrating SCIM with the organization 's identity provider to automate the process of creating, updating, and deactivating user accounts in Tableau Cloud
- C. Using SCIM exclusively for periodic audits of user accounts rather than for ongoing user account management
- D. Configuring SCIM to allow users to self-provision their accounts directly in Tableau Cloud

ANSWER: B

Explanation:

Integrating SCIM with the organization's identity provider to automate the process of creating, updating, and deactivating user accounts in Tableau Cloud is the effective approach. Tableau Cloud supports SCIM as a standard-based mechanism for synchronizing user and group information from an identity provider, such as Okta, Microsoft Entra ID, or another SCIM-compatible provider. After an administrator configures SCIM and supplies the required Tableau Cloud credentials or token to the identity provider, identity lifecycle events can be propagated automatically. For example, a new employee assigned to the appropriate application or group in the identity provider can be created in Tableau Cloud, and changes to their identity information can be synchronized. When access is removed in the identity provider, SCIM can deactivate the corresponding Tableau Cloud user, supporting timely access removal and stronger governance. This centralizes identity administration in the organization's authoritative identity system instead of requiring repeated administration in Tableau Cloud. Tableau documents SCIM as a method for automating user provisioning and group synchronization in Tableau Cloud; see [Configure SCIM Provisioning for Tableau Cloud](#) and [SCIM Provisioning Overview](#).

QUESTION NO: 5

A company uses SQL Server row-level security based on the authenticated Windows user. Tableau Server users must see only the rows permitted for their own Active Directory identities when they open a published workbook. Which approach best preserves this database-level security model?

- A. Embed one shared SQL Server credential in the published data source
- B. Configure Kerberos delegation for viewer-specific Windows authentication to SQL Server
- C. Replace the live connection with an extract refreshed by a service account
- D. Grant all Tableau Server users database administrator access

ANSWER: B

Explanation:

Configuring Kerberos delegation so Tableau Server can use the viewer's Windows identity for the SQL Server connection is the correct approach. This allows SQL Server to evaluate row-level security against the individual authenticated user rather than against a shared database account. The deployment requires correctly configured service principal names and delegation for the Tableau Server service account and the SQL Server service.

An embedded SQL Server credential would cause database access to occur under the shared credential, preventing SQL Server from distinguishing individual viewers for this purpose. Extracting the data can be useful for performance, but it does not preserve real-time database-enforced user identity. Giving all users direct database administrator rights violates least-privilege security principles.

QUESTION NO: 6

After installing Tableau Server on a Windows system, you find that the server is not accessible from client machines in the network. What should be your first step in troubleshooting this network accessibility issue?

- A. Reinstalling the network drivers on the Windows server hosting Tableau Server
- B. Checking the Windows server's firewall settings to ensure the necessary ports for Tableau Server are open
- C. Upgrading the client machines to a compatible operating system version
- D. Configuring a static IP address for the Windows server hosting Tableau Server

ANSWER: B

Explanation:

Checking the Windows server's firewall settings to ensure the necessary ports for Tableau Server are open is the appropriate first troubleshooting step because a host-based firewall can block inbound client traffic even when Tableau Server has installed and started successfully. Client browsers normally connect to Tableau Server over HTTP or HTTPS, using the ports configured for the Tableau Server gateway—commonly port 80 for HTTP and port 443 for HTTPS. If those inbound connections are not permitted on the Windows host, users on the network cannot reach the Tableau Server sign-in page or published content. Confirm that the applicable inbound firewall rules allow traffic from the intended client networks and that the ports match any custom gateway configuration. This check is quick, low-risk, and directly tests the most common local network control affecting a newly installed server. Tableau's port documentation identifies the gateway ports and other ports that may be required depending on the deployment architecture, while its installation guidance notes that network and firewall configuration must permit required communication. See [Tableau Server Ports](#) and [Install Tableau Server](#).

QUESTION NO: 7

When implementing extract encryption in Tableau Server, what is a crucial step to secure the data extracts stored on the server?

- A. Configuring a VPN tunnel for all data extract transfers to and from Tableau Server
- B. Enabling at-rest encryption for data extracts within Tableau Server's configuration settings
- C. Implementing a network intrusion detection system to monitor extract file accesses
- D. Increasing the storage capacity of the server to accommodate the additional space required by encrypted extracts

ANSWER: B

Explanation:

Enabling at-rest encryption for data extracts within Tableau Server's configuration settings is the crucial implementation step because it protects the contents of Tableau extract files while they reside on Tableau Server disks. Tableau Server supports extract encryption at rest for .hyper extracts, using encryption keys managed by Tableau Server through its key-management capabilities. After an administrator enables this feature through Tableau Services Manager configuration and applies the pending changes, newly created extracts are encrypted; existing extracts must be refreshed or republished to become encrypted. This is a data-at-rest control: it reduces the risk that someone who obtains access to extract files or the underlying storage can read their contents without the applicable encryption keys. Architects should also plan the rollout, validate that the server environment meets Tableau's requirements, and understand the operational effect on publishing and refresh workflows. Tableau documents both the configuration process and the behavior for existing versus newly created extracts in its Tableau Server security guidance. See [Tableau Server: Encrypt Data Extracts](#) and [Tableau Server Security](#).

QUESTION NO: 8

A company is transitioning to Tableau Cloud but still has critical data in on-premises databases that need to be accessed in real-time. What is the best solution for integrating these data sources with Tableau Cloud?

- A. Utilize Tableau Builder for real-time data integration
- B. Implement Tableau Bridge to establish a live connection to on-premises databases
- C. Migrate all on-premises data to the cloud before using Tableau Cloud
- D. Rely solely on Tableau Cloud's native capabilities for on-premises data integration

ANSWER: B

Explanation:

Implement Tableau Bridge to establish a live connection to on-premises databases is the appropriate solution because Tableau Bridge is Tableau Cloud's supported connectivity client for data that remains inside a private network. An organization installs Bridge on a Windows machine that has network access to the on-premises database, then associates the relevant data source with Tableau Cloud. This allows Tableau Cloud to securely reach data that would not otherwise be directly accessible from the cloud service, without requiring inbound firewall access to the organization's network.

For a real-time requirement, Bridge can support live queries: when a viewer interacts with a workbook using a supported live connection, the query is routed through the Bridge client to the private-network database and the current results are returned to Tableau Cloud. Bridge can also maintain extracts through scheduled or managed refreshes when an extract-based design is appropriate, but live connectivity is the key capability that meets the stated need for current on-premises data. The architecture preserves the existing database location while enabling governed Tableau Cloud analytics and central administration. Tableau documents Bridge as the mechanism for keeping private-network data sources connected to Tableau Cloud and describes its support for both live connections and extract refreshes. See [Tableau Bridge Overview](#) and [Use Tableau Bridge for Live Connections](#).

QUESTION NO: 9

When optimizing caching for Tableau Server to improve dashboard performance, which setting is most effective to adjust?

- A. Setting the cache to refresh every time a view is loaded to ensure the most up-to-date data is always used
- B. Configuring the cache to be cleared at a regular, scheduled interval that aligns with the data refresh schedule
- C. Disabling caching entirely to force real-time queries for all dashboard views
- D. Increasing the server's RAM to enhance its overall caching capability

ANSWER: B

Explanation:

Configuring the cache to be cleared at a regular, scheduled interval that aligns with the data refresh schedule is the most effective approach because it balances dashboard responsiveness with appropriate data freshness. Tableau Server can reuse cached query results and rendered view information, reducing repeated database work and lowering load times for frequently accessed dashboards. Keeping that cache available between data updates allows users to benefit from those reused results rather than requiring a new query for every interaction or page load.

Aligning cache invalidation with the extract-refresh or source-data-refresh cadence ensures that cached content is discarded when a newly refreshed data set should be served. This gives administrators a predictable freshness policy: users receive performant cached results during a refresh interval and updated results after the relevant data refresh completes. It also helps avoid unnecessary pressure on the database and Tableau Server processes while still meeting the organization's data-latency requirements. Tableau documents caching as an important factor in view performance and recommends managing refresh behavior according to the required data freshness. See Tableau's [Server performance tuning guidance](#) and [data cache documentation](#).

QUESTION NO: 10

During the troubleshooting of SAML authentication issues in Tableau Server, what is a common area to investigate?

- A. The network bandwidth and latency between the Tableau Server and the SAML provider
- B. The time synchronization between Tableau Server and the SAML identity provider
- C. The storage capacity of the Tableau Server to handle SAML requests
- D. The version compatibility of the web browser used to access Tableau Server

ANSWER: B

Explanation:

The time synchronization between Tableau Server and the SAML identity provider is a key area to investigate because SAML authentication relies on assertions with time-bound validity conditions. An identity provider typically includes timestamps such as when an assertion becomes valid and when it expires. Tableau Server validates those timestamps against its own system clock. If the server and identity provider clocks differ by more than the permitted tolerance, Tableau Server can treat an otherwise valid assertion as expired or not yet valid, resulting in sign-in failures. Confirm that the Tableau Server hosts and the identity provider use a reliable Network Time Protocol source, have the correct time zone and daylight-saving settings, and maintain closely aligned system clocks. This check is especially important when failures are intermittent, occur after infrastructure changes, or produce SAML errors associated with assertion validity windows. Tableau's SAML troubleshooting guidance identifies clock synchronization as an important consideration when validating SAML responses and their conditions. For configuration and troubleshooting context, see [Tableau Server SAML Authentication](#) and [Troubleshoot SAML Authentication](#).

QUESTION NO: 11

During the installation of Tableau Server on Linux, which action is crucial to ensure proper system group and file system permissions are set?

- A. Assigning the Tableau Server user to the root group to ensure full system access
- B. Creating a dedicated Tableau user and group, and setting appropriate ownership and per-missions on the Tableau directories
- C. Configuring all users on the Linux system to have administrative privileges for the duration of the Tableau Server installation
- D. Disabling the Linux system's firewall to prevent it from interfering with file permissions

ANSWER: B

Explanation:

Creating a dedicated Tableau user and group, and setting appropriate ownership and per-missions on the Tableau directories is the required approach because Tableau Server needs a controlled operating-system identity and reliable access to its installation and data locations. On Linux, Tableau Server processes run under the Tableau service account rather than as an unrestricted root user. The account and its group must be able to read, write, and execute where needed within the Tableau Server directories, including locations used for configuration, logs, temporary files, repository data, and service operations. Assigning ownership and permissions deliberately lets Tableau Server start and manage its components successfully while applying the principle of least privilege. It also prevents permission failures after installation, upgrades, restarts, or maintenance operations, and avoids exposing Tableau data or server files to unrelated local users. The installer and administrative setup procedures require elevated privileges for installation tasks, but the resulting service account and filesystem access should remain appropriately scoped. Salesforce's Linux installation guidance describes the Tableau account and directory requirements, while its deployment documentation provides the operational context for administering Tableau Server securely. See [Tableau Server on Linux Installation](#) and [Tableau Server Security](#).

QUESTION NO: 12

During the installation of Tableau Server on a Linux system, you encounter a failure with the error message indicating a permissions issue. What is the first step you should take to resolve this issue?

- A. Reinstalling the Linux operating system to ensure a clean environment for Tableau Server
- B. Checking and modifying the file and directory permissions where Tableau Server is being installed
- C. Increasing the RAM and CPU resources allocated to the Linux server
- D. Configuring the Linux server to use a different file system

ANSWER: B

Explanation:

Checking and modifying the file and directory permissions where Tableau Server is being installed is the appropriate first action because a permissions-related installer failure normally means that the account performing installation, or the Tableau service account, cannot access, create, modify, or execute required files in a relevant path. Start by identifying the exact path and account cited in the installer output or logs, then verify ownership and read, write, and execute permissions on the Tableau Server installation location and any required temporary or data directories. This approach addresses the reported condition directly while preserving the server's existing configuration and allows the installation to be retried after the required access has been granted.

Tableau's Linux installation guidance requires installing under an appropriate administrative account and specifies filesystem and directory requirements that must be met before setup can succeed. Tableau also documents that the unprivileged Tableau account is used to run Tableau Server services, so it must be able to access the locations needed for server operation. Review the documented installation prerequisites and ensure that ownership and permissions conform to the selected installation and data-directory configuration before continuing. See [Install Tableau Server on Linux](#) and [Tableau Server for Linux Knowledge Base](#).

QUESTION NO: 13

In configuring a strategy to monitor Tableau Server process metrics, what is an essential aspect to include for ensuring proactive maintenance?

- A. Setting up alerts for when server processes reach their maximum capacity thresholds
- B. Analyzing historical data once a month to identify long-term trends in server performance
- C. Only monitoring the most frequently used dashboards and data sources on the server
- D. Conducting a manual review of server logs whenever a new Tableau version is deployed

ANSWER: A

Explanation:

Setting up alerts for when server processes reach their maximum capacity thresholds is essential because proactive maintenance depends on detecting resource pressure early enough to act before it becomes a user-facing outage or performance incident. Tableau Server depends on multiple coordinated processes, and sustained high utilization or capacity conditions can indicate that a process is becoming a bottleneck. Alerts turn continuously collected operational metrics into an actionable signal for administrators, enabling them to investigate workload changes, assess process placement and capacity, and make scaling or configuration decisions before requests begin failing or degrading significantly.

Tableau's Resource Monitoring Tool is designed to help administrators observe the health and performance of Tableau Server deployments and configure alerts for conditions requiring attention. Effective alerting should be paired with clear thresholds, an owner or response procedure, and periodic review to ensure that thresholds reflect the organization's normal workload and service expectations. This creates a continuous monitoring practice rather than relying on retrospective checks. See Tableau's documentation on [Resource Monitoring Tool alerts](#) and [collecting Tableau Server performance data](#).

QUESTION NO: 14

A large retail company with a high volume of daily Tableau users requires a configuration that optimizes query performance and user experience. Which configuration setting should be prioritized?

- A. Decrease the " vizqlserver.session.expiry.timeout " value to reduce session timeouts
- B. Increase the " backgrounder.querylimit " value to allow more concurrent queries
- C. Reduce the " cache.server.timeout " value to lower the caching time
- D. Increase the " vizqlserver.querylimit " value to allow more concurrent queries

ANSWER: D

Explanation:

Increase the " vizqlserver.querylimit " value to allow more concurrent queries is the appropriate configuration to prioritize because VizQL Server is the Tableau Server process that renders views and handles the interactive query activity generated when users open, filter, drill into, and otherwise interact with dashboards. In an environment with many concurrent daily users, the number of queries that a VizQL Server process can issue simultaneously can become a constraint on interactive responsiveness. Raising `vizqlserver.querylimit`, when supported by sufficient database, CPU, memory, and VizQL Server capacity, enables more interactive requests to proceed concurrently rather than waiting behind an artificially low per-process query limit. This can reduce queueing pressure and improve the perceived responsiveness of commonly used dashboards. The setting should be adjusted through controlled load testing and monitored alongside VizQL and data-source performance, since a higher concurrency limit must be matched to available downstream database capacity. Tableau documents server configuration keys and their use with TSM in its [TSM configuration reference](#). For architecture and capacity planning context, see Tableau's [server performance monitoring guidance](#).

QUESTION NO: 15

In the context of Tableau Server, what is an important consideration when configuring access to the Metadata API for external applications?

- A. Allowing unrestricted access to the Metadata API from any external application
- B. Configuring the Metadata API to provide real-time updates to external applications
- C. Implementing OAuth for secure, token-based authentication for external applications accessing the Metadata API
- D. Ensuring external applications have direct database access for synchronized metadata retrieval

ANSWER: C

Explanation:

Implementing OAuth for secure, token-based authentication for external applications accessing the Metadata API is an important consideration because integrations should authenticate through a controlled, revocable mechanism rather than expose Tableau credentials or metadata services directly. Tableau's Metadata API is a GraphQL API that provides access to metadata governed by Tableau's existing authorization model. An external application must therefore establish an authenticated Tableau session or use an appropriately configured token-based integration, and the identity represented by that token must have permission to access the requested Tableau content and metadata.

OAuth-based connected-app patterns support this need by allowing an administrator to register and manage the external application, define its trust relationship with Tableau, and limit access through token lifetimes and scopes. This makes it practical to apply least privilege, rotate or revoke access when needed, and maintain an auditable boundary between Tableau Server and third-party services. The application should be granted only the permissions necessary for its metadata use case. Tableau documents Metadata API access and authentication requirements in its [Metadata API documentation](#), and describes managed external-integration authentication in its [Tableau Server connected apps documentation](#).

QUESTION NO: 16

When configuring a backgrounder process on a specific node in a Tableau Server deployment, what should be considered to ensure optimal performance of the backgrounder node?

- A. The backgrounder node should have a faster network connection than other nodes
- B. The node should have more processing power and memory compared to other nodes in the deployment
- C. The backgrounder node should be placed in a geographically different location than the primary server
- D. The node should run on a different operating system than the other nodes for compatibility

ANSWER: B

Explanation:

The node should have more processing power and memory compared to other nodes in the deployment. Tableau Server Backgrounder executes resource-intensive asynchronous work, including extract refreshes, flow runs, subscriptions, alerts, and some data-driven processing. Its throughput is therefore primarily constrained by the CPU capacity and available RAM on the node where Backgrounder runs. Providing a dedicated or comparatively higher-capacity node helps background jobs run concurrently with less contention and prevents scheduled work from competing directly with interactive services for system resources. The required capacity should be sized according to the volume, frequency, duration, and concurrency of extract refreshes and other scheduled tasks, with monitoring used to identify queueing or sustained resource pressure. Tableau's deployment guidance also supports separating processes across nodes as workloads scale, so that resource-heavy background processing can be independently sized and managed. See Tableau's guidance on [Backgrounder processes](#) and the [Tableau Server hardware recommendations](#).

QUESTION NO: 17

For a multinational corporation implementing Tableau, what is the most important consideration for licensing and ATR compliance?

- A. Opting for the cheapest available licensing option to minimize costs
- B. Ignoring ATR compliance as it is not crucial for multinational operations
- C. Choosing a licensing model that aligns with the global distribution of users and adheres to ATR requirements
- D. Selecting a licensing model based solely on the preferences of the IT department

ANSWER: C

Explanation:

Choosing a licensing model that aligns with the global distribution of users and adheres to ATR requirements is essential because a multinational Tableau deployment must be designed around who will use Tableau, where they are located, and the rights and restrictions governing that use. Licensing is not merely a procurement decision: the organization needs sufficient named-user or capacity-based entitlements for its expected creator, explorer, viewer, and embedded-analytics usage, as applicable. It must also maintain controls and records that support compliance with the contractual terms applicable to its deployment across regions. Considering geographic user distribution early helps the organization select the appropriate Tableau offering, plan identity and site administration, forecast growth, and avoid deployment designs that create entitlement or contractual-compliance risk. Tableau's licensing guidance describes the available licensing approaches and how licenses govern access to Tableau capabilities, while Tableau's legal terms establish the conditions under which products and services may be used. These considerations should be evaluated together with regional operational needs, rather than treated as separate technical and purchasing exercises. See Tableau's [licensing models overview](#) and [end-user license agreement information](#).

QUESTION NO: 18

A rapidly expanding retail company is planning to deploy Tableau for its nationwide operations. What is the most important factor to consider for ensuring the scalability of the Tableau deployment?

- A. Limiting the number of users to control system load
- B. Focusing only on current data requirements without considering future growth
- C. Choosing a deployment model that can scale with increasing data volume and user count
- D. Using a single server regardless of increasing data and user requirements

ANSWER: C

Explanation:

Choosing a deployment model that can scale with increasing data volume and user count is the central scalability consideration because Tableau deployment capacity must grow in step with both workload demand and organizational adoption. A nationwide retailer will typically add stores, transactions, data sources, dashboards, and concurrent users over time. The architecture therefore needs to support expansion through appropriately sized and, where needed, distributed Tableau Server processes, additional nodes, and ongoing capacity planning. Scalability is not simply a hardware choice made at initial installation; it requires forecasting usage patterns, monitoring resource utilization, and designing for peak concurrency, extract refresh activity, and increasingly complex analytics workloads. Tableau's deployment guidance emphasizes using performance and scalability testing to understand expected demand and then scaling the environment based on the results. A flexible deployment model also permits the organization to increase capacity without redesigning the entire platform as business requirements evolve. This approach protects dashboard responsiveness and refresh reliability while allowing broader self-service analytics adoption across retail operations. Tableau provides guidance for planning enterprise deployments and scaling Server topology in its [Performance and Scalability](#) documentation, as well as process and node architecture details in the [Tableau Server Processes](#) documentation.

QUESTION NO: 19

When managing a Tableau Server environment on a Linux system, which method is recommended for deploying automated backup scripts?

- A. Configuring the scripts to run automatically via the Tableau Server web interface
- B. Using cron jobs to schedule and execute backup scripts at regular intervals
- C. Relying on a third-party cloud service to handle all backup processes
- D. Manually initiating backup scripts through the Linux terminal as needed

ANSWER: B

Explanation:

Using cron jobs to schedule and execute backup scripts at regular intervals is the appropriate method for automating Tableau Server backups on Linux. Tableau Server backups are created with the `tsm maintenance backup` command, which can be placed in a shell script along with operational details such as a backup destination, logging, retention handling, and error notifications. Linux cron is designed to run such noninteractive commands according to a defined schedule, allowing administrators to create predictable, repeatable backups without relying on someone to start the process manually. The scheduled job should run under an account with the required Tableau Server permissions and access to the selected backup location. Administrators should also ensure the backup target has adequate capacity and that backups are copied or otherwise protected separately from the Tableau Server host, so they remain available for recovery after a server failure. Tableau documents the backup command and its requirements in its [Tableau Server Backup and Restore documentation](#). For the Linux scheduling mechanism, cron's scheduled-task format and operation are described in the [crontab manual page](#). This combination is a standard, supportable operational pattern for regular Tableau Server backup automation.

QUESTION NO: 20

A large organization plans to consolidate several Tableau Server instances into a single server. What is the most important consideration to ensure a successful consolidation?

- A. Consolidating all servers simultaneously to minimize the transition period
- B. Thoroughly planning the integration of data sources, user permissions, and content from each server
- C. Focusing solely on the technical aspects and not on the user impact of consolidation
- D. Immediately decommissioning all other servers before starting the consolidation process

ANSWER: B

Explanation:

Thoroughly planning the integration of data sources, user permissions, and content from each server is the most important consideration because consolidation is both a platform migration and an organizational change. The destination Tableau Server must be sized and configured to support the combined workload, while every workbook, data source, extract, schedule, subscription, project, and user or group entitlement needs an intentional migration and validation plan. Mapping authentication and identity-store details, preserving project governance, assessing duplicate content, and testing refreshes and permissions before cutover helps ensure that users retain appropriate access and that trusted reporting continues to operate. A phased approach with inventories, pilot migrations, stakeholder communication, validation, and rollback planning reduces business disruption while allowing administrators to resolve compatibility or ownership issues before moving critical content. Tableau's migration guidance emphasizes assessing the source environment, planning users and content, and validating migrated assets; its Server administration documentation also describes the central role of users, groups, projects, and permissions in governing access. See [Tableau Server migration documentation](#) and [Tableau Server permissions documentation](#).

QUESTION NO: 21

After performing load testing on Tableau Server, you observe a significant increase in response times during peak user activity. What is the most appropriate action based on this result?

- A. Immediately add more hardware resources, such as RAM and CPU, to the server
- B. Analyze server configurations and optimize performance settings before considering hardware upgrades
- C. Reduce the number of concurrent users allowed on the server to decrease load
- D. Ignore the results as temporary spikes in response times are normal during peak periods

ANSWER: B

Explanation:

Analyze server configurations and optimize performance settings before considering hardware upgrades is the appropriate response because load-test results should first be used to identify the source of the bottleneck and validate whether the existing deployment is configured efficiently. Tableau Server performance depends on how its processes are allocated, the number and placement of Background processes, VizQL capacity, caching behavior, repository performance, and the characteristics of the published workbooks and data sources. Reviewing these areas can reveal an imbalanced process topology, resource contention, inefficient queries, or workbook design issues that can be addressed without immediately increasing infrastructure cost.

Tableau recommends establishing a performance baseline and using monitoring and diagnostic information to understand system behavior under load. The results can then guide targeted tuning and, if utilization and bottleneck evidence still demonstrate insufficient capacity after optimization, provide a defensible basis for scaling CPU, memory, storage, or server nodes. This measured approach ensures that hardware changes address a verified limiting resource rather than masking a configuration or workload issue. See Tableau's [Performance and scalability guidance for Tableau Server](#) and [Tableau Server performance tuning guidance](#).

QUESTION NO: 22

When troubleshooting Kerberos authentication issues related to SPNs in Tableau Server, what common problem should be investigated first?

- A. Checking if the Kerberos tickets are expiring too quickly
- B. Verifying that the SPNs are correctly set for the Tableau Server service account
- C. Ensuring that the network firewall allows Kerberos traffic to pass through
- D. Confirming that all users have Kerberos enabled on their client machines

ANSWER: B

Explanation:

Verifying that the SPNs are correctly set for the Tableau Server service account is the appropriate first investigation because Kerberos depends on an Active Directory service principal name that uniquely maps the Tableau Server service's HTTP endpoint to the domain account running that service. When a client requests a Kerberos service ticket, Active Directory uses the requested HTTP SPN to identify the target account. A missing SPN, an SPN registered to the wrong account, or a duplicate SPN can prevent the ticket from being issued or accepted, causing authentication failures such as repeated credential prompts or fallback behavior. Administrators should confirm the Tableau Server run-as/service account, identify the fully qualified DNS names users employ to reach Tableau Server, and use the Microsoft `setspn` utility to query and validate the corresponding HTTP SPNs in Active Directory. This verification establishes that the core Kerberos identity mapping is valid before investigating secondary environmental factors. Tableau's Kerberos configuration guidance describes registering SPNs for the Tableau Server service account, while Microsoft documents querying and managing SPNs with `setspn`. See [Tableau Server Kerberos authentication documentation](#) and [Microsoft setspn documentation](#).