PassQueen
a.k.a dumpsqueen.com

Enterprise Geodata Management Professional 2201

Esri EGMP2201

Version Demo

Total Demo Questions: 15

Total Premium Questions: 158

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Planning and Designing an Enterprise Geodatabase	40
Topic 2, Configuring an Enterprise Geodatabase	43
Topic 3, Managing an Enterprise Geodatabase	64
Topic 4, Troubleshooting Enterprise Geodatabases	10
Topic 5, Mix Questions	1
Total	158

QUESTION NO: 1

A GIS analyst needs to share a large repository of lidar data with the organization.

This lidar data will have surface constraints applied for breaklines. Which type of dataset should the GIS analyst use?

A. Mosaic dataset

B. Feature dataset

C. LAS dataset

ANSWER: C

Explanation:

LAS dataset is the appropriate dataset type because it is specifically designed to reference, organize, visualize, and analyze large collections of lidar point-cloud files, including LAS and ZLAS files. Rather than importing all points into a geodatabase, a LAS dataset stores references to the source lidar files and associated metadata, making it practical for a large repository while preserving efficient access to the underlying data.

A key capability of a LAS dataset is support for surface constraints. Breaklines can be added as hard- or soft-replace constraints to control how an elevation surface is generated from the lidar points. This enables hydrologically and topographically meaningful surfaces by enforcing linear features such as ridge lines, shorelines, and stream channels during surface rendering or analysis. The LAS dataset also provides point filtering by classification, return, flag, and other lidar attributes, allowing users throughout the organization to work consistently with the managed point-cloud collection and its constraints.

Esri documents LAS datasets as the mechanism for managing and analyzing large volumes of lidar data and documents how breakline features are incorporated as surface constraints. See What is a LAS dataset? and Manage breaklines.

QUESTION NO: 2

An editor reconciles a traditional version with its parent version.

The reconcile operation is configured to detect conflicts by attribute and to favor the edit version when conflicts occur.

Which statements describe this configuration?

A. A conflict is detected only when the same attribute is modified in both versions

B. Conflicting values from the edit version are retained during automatic conflict resolution

C. Any edit to the same row is treated as a conflict

D. Conflicting values from the target version are retained during automatic conflict resolution

ANSWER: A B

Explanation:

When conflicts are detected **by attribute**, a conflict is identified only when the same attribute of the same row was modified in both the edit version and the target version. Changes to different attributes of the same row can be merged without being treated as a conflict.

When the operation is configured to **favor the edit version**, the values from the edit version are retained for conflicts that are automatically resolved during reconcile. This allows the editor's conflicting changes to take precedence over the corresponding target-version values.

Detecting conflicts by row would identify a conflict whenever both versions modify the same row, regardless of the specific fields edited. Favoring the target version would retain the parent-version values for automatically resolved conflicts. See Esri documentation on conflict detection in versioning and reconcile and post changes.

QUESTION NO: 3

In a multi-user editing environment, what is the purpose of the "version" in versioning?

A. To create a snapshot of the data at a particular point in time

B. To allow multiple users to edit the same data simultaneously

C. To track changes made to the data over time

ANSWER: B

Explanation:

To allow multiple users to edit the same data simultaneously is correct. In an enterprise geodatabase, a version is an independent, named view of versioned data. It gives an editor a workspace in which edits can be made without immediately changing the shared default version or interfering with edits being performed by other users. Multiple editors can therefore work concurrently on the same underlying datasets, each in an appropriate version and according to the organization's access permissions and editing workflow.

When editing work is ready to be integrated, the changes in a child version can be reconciled with its parent version. Reconciliation identifies conflicts where applicable and incorporates changes from the parent into the editing version for review. The editor can then post the reconciled changes to the parent version, commonly the default version, so the approved edits become available in the shared representation of the data. This version-based workflow supports coordinated multiuser data maintenance while preserving controlled integration of edits and data integrity. Esri describes versions as views of an enterprise geodatabase that enable simultaneous editing and outlines reconciliation and posting in its versioning overview and reconcile and post documentation.

QUESTION NO: 4

A GIS analyst is creating a query layer from a database view that returns parcel geometry.

Which requirements must the SQL query meet for ArcGIS Pro to create the spatial query layer?

A. Return a unique integer object ID field

B. Return one spatial field

C. Register the source table as versioned

D. Add global IDs to the source table

ANSWER: A B

Explanation:

The query must return a **unique integer object ID field**. ArcGIS Pro uses this field to uniquely identify each row in the query layer. If the source view does not contain a suitable field, the SQL statement can provide one when the database supports an appropriate method.

The query must also return **one spatial field** that contains the geometry to display. ArcGIS Pro uses that field, together with the geometry type and spatial reference identified during validation, to represent the results as a spatial layer. Registering the source data as versioned and adding global IDs are not requirements for creating a query layer. See Esri documentation on query layers and defining query layers.

QUESTION NO: 5

Which ArcGIS component allows storage of rules, relationships, and behavior for features?

A. Map document

B. Layer file

C. Geodatabase

ANSWER: C

Explanation:

Geodatabase is correct because it is ArcGIS's data management framework for storing geographic data along with the schema and integrity mechanisms that define how features are structured and interact. A geodatabase can store feature classes and tables, domains, subtypes, relationship classes, attribute rules, topology rules, and network datasets. These capabilities enable an organization to model real-world business rules and feature behavior directly with its authoritative data rather than merely storing geometry and attribute values. For example, relationship classes can formally associate related records, domains can constrain permitted attribute values, and topology rules can validate spatial relationships and data quality. In enterprise geodata management workflows, these rules and relationships are important because they help maintain consistent, reliable data for multiple editors, applications, and services. Esri describes the geodatabase as a collection of geographic datasets that can represent the properties and behavior of geographic information, including rules for maintaining data integrity. See Esri's What is a geodatabase? overview and its documentation on geodatabase design.

QUESTION NO: 6

An organization has an enterprise geodatabase used for editing and public use.

Editors are experiencing performance issues during peak hours.

The GIS data administrator needs to make sure that the editing and public usage do not affect each other. Which action should be taken?

A. Build new feature datasets

B. Separate permissions for public services

C. Create separate database instances

ANSWER: C

Explanation:

Create separate database instances is the appropriate action because it isolates the editing workload from the public-facing workload. Enterprise geodatabase editing can generate substantial transactional activity, including locks, version-management operations, index updates, and database resource consumption. At the same time, public map and feature services can create high volumes of concurrent read requests. Hosting both workloads against one database instance means that peak public demand can contend with editors for database CPU, memory, disk I/O, and connections.

A dedicated enterprise geodatabase for editing should remain the authoritative source of data. A separate enterprise geodatabase can host a replicated copy for public services, normally configured to support efficient read-only access. Geodatabase replication provides a supported mechanism to distribute data between geodatabases and synchronize changes from the editing geodatabase to the public-serving geodatabase on a schedule appropriate to the organization's currency requirements. This architecture permits each environment to be tuned, scaled, maintained, and monitored independently while keeping operational editing isolated from public access demand.

Esri describes replication as a way to distribute data across geodatabases and synchronize changes: ArcGIS Pro documentation: Replication and geodatabases. See also ArcGIS Enterprise documentation: Strategies for data access.

QUESTION NO: 7

In a distributed GIS environment, what is the primary consideration for designing a geodata synchronization strategy?

A. Data consistency and integrity

B. Data latency and synchronization

C. Data security and access control

ANSWER: A B

Explanation:

Data consistency and integrity are fundamental to a geodata synchronization strategy because distributed sites must share a reliable, coherent representation of the same business data. Synchronization must accurately transfer inserts, updates, and deletes while preserving relationships, schema compatibility, and rules that maintain the quality of the geodatabase. A strategy should also account for how conflicts are detected and reconciled when the same data is edited at more than one location. ArcGIS geodatabase replication supports this workflow by tracking changes and exchanging them between replica geodatabases, helping organizations maintain coordinated datasets across disconnected or geographically distributed environments.

Data latency and synchronization are also primary design considerations. Teams need to decide how often changes are exchanged, whether synchronization will occur on a schedule or on demand, and how intermittent connectivity or limited bandwidth affects the time required for updates to reach other sites. The acceptable delay depends on operational needs: field or emergency workflows may require frequent exchanges, while less time-sensitive editing can use periodic synchronization. Esri documents the synchronization process, including sending and receiving data changes, in its Synchronize changes between replicas guidance and describes distributed editing patterns in its Replication and geodatabases overview.

QUESTION NO: 8

A GIS data administrator must provide a group of users access to selected nonsensitive fields in an enterprise geodatabase feature class.

Which actions should the administrator take to enforce this requirement at the database level?

A. Create a database view that excludes the restricted fields

B. Grant privileges on the view to a database role

C. Share a layer file with the restricted fields hidden

D. Grant SELECT privileges on the base feature class to all users

ANSWER: A B

Explanation:

A database view can expose only the authorized columns from the underlying feature class. Users who access the view cannot retrieve fields omitted from its SQL definition, which makes a view appropriate for limiting access to sensitive attributes.

The administrator should grant the required privileges on the view to a database role and make the appropriate users members of that role. Using a role provides consistent, centralized privilege management for multiple users. The base feature class should not be granted directly to users who must not access its restricted fields. See Esri documentation on database views and enterprise geodatabase authentication and authorization.

QUESTION NO: 9

A GIS administrator learns that geodatabase users report decreasing performance when adding data from child versions to their map.

The organization uses a complex traditional version tree architecture Python script completes batch-reconcile/ post operations, compresses the geodatabase, and data owners rebuild indexes and update statistics Python script runs overnight with little to no geodatabase connections being made Which Analyze Datasets parameter should be checked?

A. Include System Tables

B. Analyze Archive Tables For Selected Datasets

C. Analyze Base Tables For Selected Datasets

D. Analyze Delta Tables For Selected Datasets

ANSWER: D

Explanation:

Analyze Delta Tables For Selected Datasets should be checked. In traditional versioning, edits made in child versions are stored in the adds and deletes delta tables rather than being immediately written to the base table. When a user adds a dataset from a child version to a map, ArcGIS and the underlying DBMS must query the versioned data representation, including the relevant delta tables, to construct the version's current view of the data. In a complex version tree, those tables can become substantial and their row distributions can change considerably after batch reconcile/post operations and a compress operation.

Updating statistics for the selected datasets' delta tables allows the database query optimizer to use current information about delta-table size and data distribution when choosing an execution plan. This is particularly appropriate in the stated overnight maintenance window, after reconcile/post and compression have completed and while connections are minimal. Esri recommends maintaining database statistics as part of traditional-versioning maintenance, and the Analyze Datasets tool provides a specific setting for analyzing delta tables associated with selected versioned datasets. Current statistics help maintain efficient map display and query performance for child versions. See Esri's Analyze Datasets tool reference and geodatabase maintenance guidance.

QUESTION NO: 10

You are designing a geodatabase to support a web GIS application that requires real-time data access. What is the primary consideration for optimizing database performance?

A. Data modeling and schema design

B. Indexing and query optimization

C. Connection pooling and caching

ANSWER: B C

Explanation:

Indexing and query optimization is essential for real-time web GIS workloads because database indexes allow the DBMS to locate the relatively small set of rows needed by a request without scanning an entire business table. Appropriate attribute indexes support common filter and join fields, while spatial indexes support efficient location and geometry queries. Query design should likewise target only required records and fields so that each map, feature, or query service request can be answered with minimal database work. Esri describes how indexes improve data retrieval and when they should be maintained in its ArcGIS Server indexing guidance.

Connection pooling and caching also supports responsive access under concurrent web traffic. ArcGIS Server service instances are reused to process requests, helping avoid the overhead of repeatedly creating server-side resources and database connections for every user request. Proper service capacity and instance management therefore help maintain throughput as demand rises. Where requests can safely use recently generated results, caching reduces repeated processing and database reads, preserving capacity for requests that must retrieve the newest operational data. Esri explains the role of service instances and the ArcSOC process in its ArcSOC process documentation.

QUESTION NO: 11

An organization uses branch versioning for a feature service.

Editors must be able to work in their own named versions, while reviewers must be able to view those versions but not edit them. Which version access level should be assigned to the named versions?

A. Protected

B. Public

C. Private

ANSWER: A

Explanation:

Protected is the appropriate version access level. A protected version can be viewed by users who have access to the service, but only the version owner can edit the version. This allows reviewers to inspect an editor's work without being able to modify it.

A private version can only be accessed by its owner, preventing reviewers from viewing it. A public version permits other authorized users to edit the version, which does not meet the requirement to restrict editing to the owner. See Esri documentation on version access permissions.

QUESTION NO: 12

A GIS data administrator is configuring attribute rules for an asset feature class.

The asset_status field must automatically receive a value when a new asset is created. Editors must also be prevented from saving an asset with an installation date later than the current date.

Which attribute rule types should be used?

A. Calculation rule

B. Constraint rule

C. Validation rule

D. Relationship rule

ANSWER: A B

Explanation:

A calculation rule is used to automatically populate an attribute value. The rule can calculate the initial value of asset_status when an asset is inserted, ensuring that the field is consistently populated without requiring manual entry.

A constraint rule is used to prevent invalid edits. The rule evaluates the installation date during the edit operation and returns an error message when the date is later than the current date, preventing the editor from saving the invalid feature. See Esri documentation on attribute rule types and constraint attribute rules.

QUESTION NO: 13

An organization needs to edit GIS data using web services.

The data must be stored locally in the organization's servers.

Specific business fields must be indexed in the database to help with performance. Which storage should be used for the data?

A. Enterprise geodatabase

B. File geodatabase

C. Hosted relational database

ANSWER: A

Explanation:

An enterprise geodatabase is the appropriate storage choice because it stores spatial and tabular GIS data in an organization-managed relational database management system, such as Microsoft SQL Server, Oracle, or PostgreSQL, on infrastructure controlled by the organization. This satisfies the requirement that the data remain local to the organization's servers while providing the database capabilities needed for enterprise operations.

Enterprise geodatabases are designed for use with ArcGIS Enterprise and ArcGIS Server. Registered enterprise geodatabase data can be published as feature services that support authorized web-based editing. This architecture lets the organization centrally manage data while exposing editing workflows through web services. The relational database also permits database administrators or data owners to create indexes for attribute fields used frequently in queries, searches, joins, or editing workflows. Proper attribute indexing reduces the amount of data the database must scan and can substantially improve service responsiveness; geodatabases also maintain spatial indexes for efficient spatial access.

This combination of organization-controlled server storage, editable feature services, and configurable business-field indexes makes enterprise geodatabase the suitable choice. Esri documents enterprise geodatabases at Enterprise geodatabases and explains field indexing at Attribute indexes.

QUESTION NO: 14

In a distributed GIS environment, what is the primary consideration for designing a geodata replication strategy?

A. Data consistency and integrity

B. Data latency and synchronization

C. Data security and access control

ANSWER: A B

Explanation:

Data consistency and integrity are central to a geodata replication strategy because replicated geodatabases are copies of related datasets that may be edited at different locations. A sound design defines which datasets are replicated, where edits are permitted, how changes are exchanged, and how conflicts are detected and reconciled. These controls help ensure that the participating geodatabases remain logically consistent and that authoritative edits are retained accurately throughout the replication lifecycle.

Data latency and synchronization are equally essential design considerations. Replication is commonly asynchronous: changes are packaged and sent through synchronization rather than being immediately available at every site. Consequently, the organization must establish an appropriate synchronization frequency, account for network availability and transfer volumes, and decide how much delay is acceptable for each operational workflow. The replica type and data direction—such as one-way or two-way exchange—should support those timeliness requirements while preserving reliable change tracking. Together, consistency controls and a realistic synchronization schedule enable distributed teams to work with dependable data while accommodating disconnected or bandwidth-constrained environments. Esri documents the replica creation, change tracking, synchronization, and conflict-management workflow in its Introduction to geodatabase replication and Synchronize changes guidance.

QUESTION NO: 15

Multiple editors in a web application need to collaboratively edit the same dataset using the following requirements:

• Each editor works in isolation until they come to a good stopping point

• The editor shares their edits with the other editors at this point

• The editor again goes into isolation to continue editing until the next stopping point is reached

• At any point, an editor can choose to see the edits that other editors have shared without sharing their own edits

How should the dataset be registered?

A. Branch versioned

B. Traditional versioned with editor tracking enabled

C. Traditional versioned with the option to move edits to base

ANSWER: A

Explanation:

Branch versioned is correct because branch versioning supports collaborative, service-based editing workflows in web applications. The dataset is published through a feature service, and editors can work in separate named versions, which keeps each editor's in-progress work isolated from the shared version. When an editor reaches an appropriate stopping point, they can reconcile and post their completed changes to the shared target version, making those changes available to the wider editing group.

This workflow also supports the requirement to obtain changes shared by other editors without making one's own unshared edits visible. An editor can reconcile their named version with its ancestor version to incorporate and inspect shared changes, then continue editing in isolation. Posting remains a separate, deliberate operation, so reconciling does not itself publish the editor's current work. This separation of isolated named-version edits, reconcile, and post operations provides the controlled collaboration cycle described in the scenario.

Branch versioning is designed for web-service workflows and is managed through ArcGIS Enterprise feature services and version management capabilities. Esri documents the branch-versioning workflow and its use with feature services at Branch version scenarios and Reconcile and post changes.