# Free Salesforce Analytics-Con-301 Practice Questions

**Shared by Talley on 13-08-2026**

**For More Free Questions and Preparation Resources**

Check the Links on Last Page

## Question 1

Question Type: MultipleChoice

A client has a dashboard that uses a bar chart to visualize sales by Sub-Category and a detail table that has all the orders for the products within Sub-

Category. The table has more than 10,000 rows of data and is slow to load.

A consultant plans to add an action so when the client interacts with the bar chart, only the relevant data appears in the table.

What will provide the fastest rendering of the dashboard?

### Options:

A- Add a filter action, set 'Run action on' to Select, and set 'Clearing the selection will' to Exclude all values.
B- Add a highlight action and set Target Highlighting to Sub-Category.
C- Add a highlight action and set Target Highlighting to All Fields.
D- Add a filter action, set 'Run action on' to Menu, and set 'Clearing the selection will' to Show all values.

### Answer:

A

### Explanation:

To optimize the dashboard rendering, particularly when dealing with a large dataset, a filter action is the most effective tool. Here's why the specified choice is optimal:

Add a filter action: This action creates a direct filter on the detail table based on the selection in the bar chart. It ensures that only data related to the selected sub-category is loaded into the table, significantly reducing load time and improving performance.

Set 'Run action on' to Select: This setting means the filter action will be triggered as soon as the user selects a bar in the bar chart. Immediate activation of the filter ensures that the dashboard is interactive and responsive.

Set 'Clearing the selection will' to Exclude all values: When the selection is cleared, this setting ensures that no data is shown, which avoids loading the entire dataset unnecessarily. This maintains performance when no sub-category is actively selected.

Reference

This strategy follows Tableau's performance best practices by using actions to limit the amount

of data processed and rendered, as detailed in the Tableau User Guide and training materials on Dashboard Actions for optimizing large datasets.

## Question 2

Question Type: MultipleChoice

A client wants guidance for Creators to build efficient extracts from large data sources.

What are three Tableau best practices that the Creators should use? Select three.

### Options:

A- Keep only the data required for analysis by using extract filters.
B- Use aggregate data for visible dimensions, whenever possible.
C- Use only live connections as they are always faster than extracts.
D- Include all the data from the original data source in the extract.
E- Hide all unused fields.

### Answer:

A, B, E

### Explanation:

To build efficient extracts from large data sources, it is crucial to minimize the load and optimize the performance of the extracts:

A . Keep only the data required for analysis by using extract filters: This best practice involves using filters to reduce the volume of data extracted, thus focusing only on the data necessary for analysis.

B . Use aggregate data for visible dimensions, whenever possible: Aggregating data at the time of extraction reduces the granularity of the data, which can significantly improve performance and reduce the size of the extract.

E . Hide all unused fields: Removing fields that are not needed for analysis from the extract reduces the complexity and size of the data model, which in turn enhances performance and speeds up load times.

These practices are endorsed in Tableau's official documentation and training sessions as effective ways to enhance the performance of Tableau extracts and optimize dashboard responsiveness.

## Question 3

Question Type: MultipleChoice

A client has a database that stores widget inventory by day and it is updated on a nonstandard schedule as shown below.

**Widget Inventory Table**

| Date | Change | Widget Count |
|---|---|---|
| 12/31/2023 | -20 | 1080 |
| 01/02/2024 | -30 | 1050 |
| 01/03/2024 | +60 | 1110 |

They want a data visualization that shows widget inventory daily, however their business unit does not have the ability to modify the data warehouse

structure.

What should the client do to achieve the desired result?

### Options:

A- Create a temporary table in the database.
B- Use Tableau Desktop to visualize null values.
C- Update the Widget Inventory Table to be a daily snapshot.
D- Use Tableau Prep to add new rows.

### Answer:

D

### Explanation:

For a client who needs a daily visualization of widget inventory but cannot modify the data warehouse structure, the best approach is to use Tableau Prep to add new rows. Tableau Prep can be used to manipulate the existing dataset by adding missing date entries and appropriately adjusting inventory counts based on available data. This allows the creation of a complete daily snapshot for visualization without needing changes to the underlying database structure.

## Question 4

Question Type: MultipleChoice

A client is considering migrating from Tableau Server to Tableau Cloud.

Which two elements are determining factors of whether the client should use Tableau Server or Tableau Cloud? Select two.

### Options:

A- Whether or not the client plans to leverage single sign-on (SSO)
B- Whether or not there are large numbers of concurrent extract refreshes
C- Whether or not the client needs the ability to connect to public, cloud-based data sources
D- Amount of data storage used on the client's existing server

### Answer:

A, B

### Explanation:

When considering a migration from Tableau Server to Tableau Cloud, two critical factors to consider are the client's need for single sign-on (SSO) and the volume of concurrent extract refreshes.

Single Sign-On (SSO): Tableau Cloud supports SSO, which can streamline user authentication and enhance security. If the client plans to leverage SSO, Tableau Cloud may be a suitable choice1.

Concurrent Extract Refreshes: The number of concurrent extract refreshes is a significant factor because it impacts performance and resource allocation. Tableau Server might be more appropriate if the client has a high volume of concurrent extract refreshes, as it allows for more control over the infrastructure to manage these workloads2.

## Question 5

Question Type: MultipleChoice

A Tableau Cloud client has requested a custom dashboard to help track which data sources are used most frequently in dashboards across their site.

Which two actions should the client use to access the necessary metadata? Choose two.

## Options:

A- Connect directly to the Site Content data source within the Admin Insights project.
B- Query metadata through the GraphiQL engine.
C- Access metadata through the Metadata API.
D- Download metadata through Tableau Catalog.

## Answer:

B, C

## Explanation:

To track which data sources are used most frequently across a site in Tableau Cloud, the client should use the GraphiQL engine and the Metadata API. The GraphiQL engine allows for interactive exploration of the metadata, making it easier to construct and test queries1. The Metadata API provides access to metadata and lineage of external assets used by the content published to Tableau Cloud, which is essential for tracking data source usage2.

# Question 6

Question Type: MultipleChoice

A client wants to see the average number of orders per customer per month, broken down by region. The client has created the following calculated field:

Orders per Customer: {FIXED [Customer ID]: COUNTD([Order ID])}

The client then creates a line chart that plots AVG(Orders per Customer) over MONTH(Order Date) by Region. The numbers shown by this chart are far higher

than the customer expects.

The client asks a consultant to rewrite the calculation so the result meets their expectation.

Which calculation should the consultant use?

## Options:

A- {INCLUDE [Customer ID]: COUNTD([Order ID])}
B- {FIXED [Customer ID], [Region]: COUNTD([Order ID])}
C- {EXCLUDE [Customer ID]: COUNTD([Order ID])}
D- {FIXED [Customer ID], [Region], [Order Date]: COUNTD([Order ID])}

## Answer:

B

## Explanation:

The calculation {FIXED [Customer ID], [Region]: COUNTD([Order ID])} is the correct one to use for this scenario. This Level of Detail (LOD) expression will calculate the distinct count of orders for each customer within each region, which is then averaged per month. This approach ensures that the average number of orders per customer is accurately calculated for each region and then broken down by month, aligning with the client's expectations.

The initial calculation provided by the client likely overestimates the average number of orders per customer per month by region due to improper granularity control. The revised calculation must take into account both the customer and the region to correctly aggregate the data:

FIXED Level of Detail Expression: This calculation uses a FIXED expression to count distinct order IDs for each customer within each region. This ensures that the count of orders is correctly grouped by both customer ID and region, addressing potential duplication or misaggregation issues.

Accurate Aggregation: By specifying both [Customer ID] and [Region] in the FIXED expression, the calculation prevents the overcounting of orders that may appear if only customer ID was considered, especially when a customer could be ordering from multiple regions.

Level of Detail Expressions in Tableau: These expressions allow you to specify the level of granularity you need for your calculations, independent of the visualization's level of detail, thus offering precise control over data aggregation.

# Question 7

Question Type: MultipleChoice

A customer plans to do an in-place upgrade of their single node Tableau Server from 2023.1 to the most recent version.

What is the correct sequence to prepare for an in-place upgrade?

## Options:

A- * In the production environment:* Disable scheduled tasks.* Uninstall Tableau Server 2023.1.* Run the upgrade script for the most recent version of Tableau Server.* Confirm everything works as expected and test new features.
B- * In the production environment:* Disable scheduled tasks.* Run the upgrade script for the most recent version of Tableau Server.* Confirm everything works as expected and test new features.
C- * In a non-production environment:* Install the most recent version of Tableau Server.* Back up the existing production environment.* Restore settings and backup into the non-production environment.* Confirm everything works as expected and test new features.* Redirect user traffic from the production environment to the non-production environment.
D- * In a non-production environment:* Clone a copy of existing production environment to create a VM snapshot.* Restore the VM snapshot into the non-production environment.* Run the upgrade script for the most recent version of Tableau Server.* Confirm everything works as expected and test new features.* Redirect user traffic from the production environment to the non-production environment.

## Answer:

B

## Explanation:

Tableau's official in-place upgrade steps for single-node Tableau Server require:

Disable scheduled tasks (to prevent extract failures during upgrade).

Run the upgrade script for the new version.

Validate and test the server after upgrade.

The critical detail:

An in-place upgrade means upgrading the existing production environment directly, not installing in a separate environment or cloning VMs.

Option B matches exactly the documented operational steps for an in-place upgrade.

Why the other options are incorrect:

A: Uninstalling Tableau Server is not part of an in-place upgrade workflow.

C: This describes a migration or test upgrade in a separate environment, not an in-place upgrade.

D: VM cloning and restoring is not required for in-place upgrades and is not part of Tableau's official procedure.

Only B represents the correct in-place upgrade sequence.

Tableau Server in-place upgrade instructions: disable tasks run upgrade script test.

Single-node upgrade guidelines noting that uninstalling is not required.

Tableau Server administration materials clarifying the difference between in-place upgrade and sandbox testing.

## Question 8

Question Type: MultipleChoice

A client calculates the percent of total sales for a particular region compared to all regions.

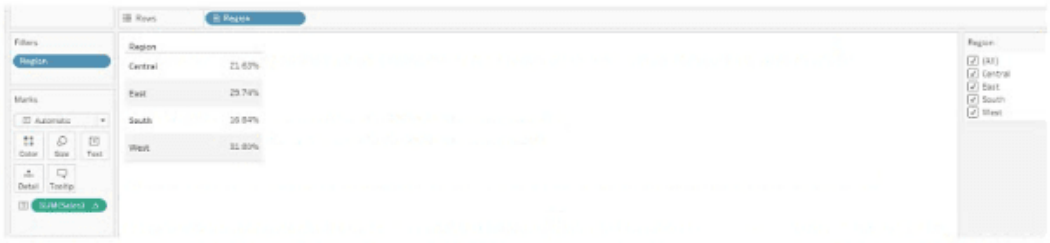

The Sales percentage is inadvertently recalculated each time the filter is applied to the Region.

Which calculation will fix the automatic recalculation on the % of total field?

### Options:

A- {FIXED [Region]:[Sales]}/{FIXED: SUM([Sales])}
B- {FIXED [Region]:sum([Sales])}
C- {FIXED [Region]:sum([Sales])}/{FIXED :SUM([Sales])
D- {FIXED [Region]:sum([Sales])}/SUM([Sales]}

### Answer:

D

### Explanation:

To correctly calculate the percent of total sales for a particular region compared to all regions, and to ensure that the calculation does not get inadvertently recalculated with each region filter application, the recommended calculation is:

{FIXED [Region]: sum([Sales])}: This part of the formula computes the sum of sales for each region, regardless of any filters applied to the view. It uses a Level of Detail expression to fix the sum of sales to each region, ensuring that filtering by regions won't affect the calculated value.

SUM([Sales]): This part computes the total sum of sales across all regions and is recalculated dynamically based on the filters applied to other parts of the dashboard or worksheet.

Combining the two parts: By dividing the fixed regional sales by the total sales, we get the proportion of sales for each region as compared to the total. This calculation ensures that while the denominator adjusts according to filters, the numerator remains fixed for each region, accurately reflecting the sales percentage without being affected by the region filter directly.

Reference

This calculation follows Tableau's best practices for using Level of Detail expressions to manage computation granularity in the presence of dashboard filters, as outlined in the Tableau User Guide and official Tableau training materials.

## Question 9

Question Type: MultipleChoice

A consultant wants to improve the performance of reports by moving calculations to the data layer and materializing them in the extract.

Which calculation should the consultant use?

### Options:

A- ZN([Sales])*(1 - ZN([Discount]))
B- CASE [Sector Parameter]WHEN 1 THEN 'green'WHEN 2 THEN 'yellow'
C- SUM([Profit])/SUM([Sales])
D- POWER(ZN(SUM([Sales]))/LOOKUP(ZN(SUM([Sales])), FIRST()),ZN(1/(INDEX()-1)))- 1END

### Answer:

C

### Explanation:

To improve performance by moving calculations to the data layer and materializing them in the extract, the consultant should choose calculations that benefit from pre-computation and significantly reduce the load during query time:

Aggregation-Level Calculation: The formula SUM([Profit])/SUM([Sales]) calculates a ratio at an aggregate level, which is ideal for pre-computation. Materializing this calculation in the extract means that the complex division operation is done once and stored, rather than being recalculated every time the report is accessed.

Performance Improvement: By pre-computing this aggregate ratio, Tableau can utilize the pre-calculated fields directly in visualizations, which speeds up report loading and interaction times as the heavy lifting of data processing is done during the data preparation stage.

Materialization in Extracts: This concept involves pre-calculating and storing complex aggregations or calculations within the Tableau data extract itself, improving performance by reducing the computational load during visualization rendering.

To Get Premium Files for Analytics-Con-301 Visit

https://www.p2pexams.com/products/analytics-con-301

For More Free Questions Visit

https://www.p2pexams.com/salesforce/pdf/analytics-con-301