



Salesforce Tableau CRM Einstein Discovery Consultant

Salesforce Certified-Tableau-CRM-and-Einstein-Discovery-Consultant

Version Demo

Total Demo Questions: 28

Total Premium Questions: 281

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

PASSQUEEN.COM

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Data Layer	63
Topic 2, Security	50
Topic 3, Dashboard Design	85
Topic 4, Einstein Discovery	54
Topic 5, Implementation	29
Total	281

QUESTION NO: 1

Which three things can be done with the Einstein Analytics Dashboard Inspector? Choose 3 answers

- A. Automatically remove bottlenecks to make queries run faster.
- B. Get a list of recommendations on how to improve the performance of the dashboard.
- C. View all queries and the time it took to run each one.
- D. View the total time required to run all queries.
- E. See the final query for each step along with query results.

ANSWER: B C E

Explanation:

The Dashboard Inspector is a diagnostic tool for understanding how a CRM Analytics dashboard performs when it loads. **Get a list of recommendations on how to improve the performance of the dashboard.** is correct because Inspector evaluates dashboard design and query behavior, then surfaces actionable performance guidance to help an author optimize the dashboard.

View all queries and the time it took to run each one. is also correct. Its query information lets the author identify the individual queries associated with dashboard steps and inspect timing details, which is essential for locating the queries that contribute most to load delays.

See the final query for each step along with query results. is correct because Inspector provides step-level query debugging information. Seeing the final generated query and its returned results helps an author validate filters, bindings, aggregations, and the data used to render a dashboard widget. Together, recommendations, per-query timing, and final-query/result inspection support both performance tuning and troubleshooting. Salesforce documents dashboard performance analysis and optimization capabilities in its CRM Analytics resources, including the [dashboard optimization documentation](#) and the [CRM Analytics dashboard optimization Trailhead module](#).

QUESTION NO: 2

After getting approval of the dashboard layout design for a desktop, the Tableau CRM consultant is ready to start the design process for a mobile layout.

What are three considerations that the consultant should keep in mind when developing the layout? (Choose three.)

- A. If more than one layout is eligible, the one with the most device properties set is used. If there is a tie, the most recently defined layout is used.
- B. A layout for mobile is eligible for use when the device meets all the device properties set in the Layout panel.
- C. If no layouts are eligible for the mobile device, the first defined layout is used.
- D. If no layouts are eligible for the mobile device, an error message will be displayed.
- E. There are widgets that cannot be displayed on mobile layouts.

ANSWER: A B C

Explanation:

Tableau CRM dashboard layouts are selected at run time according to the properties of the viewing device, so the consultant should configure layouts with an understanding of the layout-selection rules. "A layout for mobile is eligible for use when the device meets all the device properties set in the Layout panel" is correct because every configured property for a layout must match the device before that layout can be considered eligible. When more than one eligible layout exists, "If more than one layout is eligible, the one with the most device properties set is used. If there is a tie, the most recently defined layout is used" correctly describes how Tableau CRM selects the more specific layout and resolves a tie. Finally, "If no layouts are eligible for the mobile device, the first defined layout is used" describes the fallback behavior. These rules make it important

to create a sensible baseline layout first, then add more targeted layouts for device form factors, orientation, and dimensions as needed. A consultant can therefore ensure that users receive the intended mobile experience while retaining a predictable fallback if a device does not match any specialized layout. Salesforce documents these dashboard layout eligibility, priority, and fallback rules in [Dashboard Layout Rules](#). For general dashboard design guidance, see Salesforce's [CRM Analytics Dashboards Quick Look](#).

QUESTION NO: 3

Which chart type is suitable for rendering five measures in a lens visualization?

- A. Treemap chart
- B. Scatter chart
- C. Metric Radar chart
- D. Stacked Bar chart

ANSWER: C

Explanation:

Metric Radar chart is suitable because CRM Analytics uses this visualization to compare several measures simultaneously in a single lens. Each selected measure is represented as a separate axis radiating from a common center, and the plotted shape makes the relative profile of those measures easy to assess at a glance. This is especially useful when the goal is to display a compact, multivariate view—for example, comparing performance measures such as revenue, margin, pipeline coverage, win rate, and customer retention together—rather than emphasizing a single value or a simple category comparison. A lens can include up to five measures in a metric radar visualization, making it the intended choice for the stated requirement. CRM Analytics chart selection should reflect both the data configuration supported by the chart and the analytical question the visualization needs to answer; metric radar charts are designed for this type of multi-measure comparison. See Salesforce's [CRM Analytics chart types documentation](#) and the [CRM Analytics lenses Trailhead module](#) for guidance on choosing and configuring lens visualizations.

QUESTION NO: 4

In a Compare Table formula, you can refer to other columns with:

- A. Their names
- B. Numbers (1..9)
- C. Letters (A..Z)
- D. All of the above

ANSWER: C

Explanation:

Letters (A..Z) is correct. Compare Table formulas use letter-based column references, assigned according to the columns' position in the table. This lets an author create a calculation based on values already displayed by the Compare Table, such as calculating a difference, percentage, or ratio between two displayed measures. For example, after the table assigns references to its columns, a formula can combine those references directly using arithmetic operators. The letters are therefore formula identifiers for the table columns, rather than the fields' dataset API names or numeric column positions. This convention makes formulas portable within the Compare Table configuration and helps distinguish table-calculation references from underlying CRM Analytics query fields. When configuring a Compare Table, confirm the displayed column order before writing or maintaining a formula, because the letter reference corresponds to that order. Salesforce's CRM Analytics documentation describes Compare Tables as a way to compare values and add formula-based calculations, while the dashboard designer learning material covers configuring table visualizations and their calculations. See the [Salesforce Compare Tables documentation](#) and [Salesforce CRM Analytics Dashboard Designer learning module](#).

QUESTION NO: 5

A consultant created an Einstein Analytics dashboard in a sandbox. Now, the dashboard needs to be migrated into production.

To complete the migration, what are the consultant's three options? Choose 3 answers

- A. Ant Migration Tool
- B. Analytics External Data API
- C. Analytics REST API
- D. Change sets
- E. Analytics dashboard connector

ANSWER: A C D

Explanation:

Ant Migration Tool, **Analytics REST API**, and **Change sets** are supported approaches for moving CRM Analytics (formerly Einstein Analytics) dashboard assets between a sandbox and production. The Ant Migration Tool uses the Salesforce Metadata API to retrieve metadata from the source org and deploy it to the target org. This makes it useful for repeatable, source-controlled, and automated release processes. Salesforce documents CRM Analytics metadata types, including dashboard-related assets, in its Metadata API reference: [AnalyticsDashboard Metadata Type](#).

Change sets provide an admin-friendly, declarative deployment path between connected Salesforce orgs. A consultant can upload the dashboard and its related CRM Analytics components from the sandbox and then validate and deploy the inbound change set in production. The **Analytics REST API** supports programmatic access to Analytics assets, including dashboard definitions and dashboard management operations. It can therefore be used by an integration or custom deployment process to read, create, and update dashboard assets across environments. Salesforce's API overview and resource documentation are available in the [CRM Analytics REST API Developer Guide](#).

QUESTION NO: 6

What's the best way to describe the template-info.json file?

- A. It's a read-only file with all the user documentation for a template.
- B. It's the boss of the template with all the information needed to create the app.
- C. It makes sure that the app includes all the right Salesforce data.
- D. It's automatically populated every time you create an app from the template.

ANSWER: B

Explanation:

It's the boss of the template with all the information needed to create the app. The `template-info.json` file is the primary configuration and metadata file for an Analytics (CRM Analytics/Tableau CRM) app template. Salesforce uses it to identify the template and to determine how an app should be created from that template. It contains essential template-level information, such as the template's name, label, description, version details, and references to the template assets and configuration that are packaged with it. In practical terms, it acts as the central manifest that ties the template's contents together and supplies the information required during app creation and installation. This is why the description of the file as the template's "boss" is an effective shorthand: it is not merely documentation or a file that is populated after each app is created; it provides the core instructions and metadata Salesforce needs to instantiate an app from the template.

Salesforce's template development documentation describes the template files and their roles, including the central template information file: [Salesforce Analytics Template Syntax documentation](#). For broader context on creating and distributing CRM Analytics apps with templates, see [Salesforce Trailhead: CRM Analytics Apps Basics](#).

QUESTION NO: 7

What are the two types of bindings? Choose 2:

- A. Results binding
- B. Data bindings
- C. Description bindings
- D. Selection bindings

ANSWER: A D

Explanation:

CRM Analytics dashboard bindings are categorized as result bindings and selection bindings. Result bindings use the output of a step, such as values returned by a query, to dynamically supply a property or input to another dashboard widget or step. This lets a dashboard respond to computed results without requiring a viewer to manually choose a value. Selection bindings use a viewer's selection in a widget, such as clicking a chart segment or choosing a value from a list, as the input for another widget or step. They support interactive, connected dashboard behavior by passing the selected value into the target's query, filter, or configuration. In both cases, bindings help dashboard authors create reusable and responsive experiences by connecting source and target components through JSON configuration. Salesforce documents these two binding categories and describes how they pass either query results or user selections between dashboard elements. See the Salesforce CRM Analytics bindings overview in the [CRM Analytics Developer Guide](#) and the related [Salesforce Help documentation on dashboard bindings](#).

QUESTION NO: 8

An Einstein Discovery consultant is preparing a dataset for a binary classification story that predicts whether cases will be escalated.

Which two actions should the consultant take to help create a reliable model? Choose 2 answers

- A. Remove or exclude fields that reveal the outcome after the case has been escalated.
- B. Review records with missing values in the outcome field.
- C. Use the Case ID as the primary predictor.
- D. Replace the binary outcome with a long text field.

ANSWER: A B

Explanation:

Remove or exclude fields that reveal the outcome after the case has been escalated. is correct because these fields can create data leakage. A predictor that becomes available only after the target event occurs can make model-quality metrics appear artificially strong while producing unreliable predictions when the model is used on active cases.

Review records with missing values in the outcome field. is also correct because the story needs a valid outcome for the records used to train and validate the model. The consultant should determine whether records with missing escalation outcomes should be excluded or whether the source process needs to be corrected.

Record IDs are generally not meaningful business predictors and should also be reviewed, but changing the target field to a text description does not address data quality or model reliability. See [Einstein Discovery Basics](#).

QUESTION NO: 9

The Tableau CRM team at a company has created three dataflows.

myDataflowOne: this dataflow tasks 2 hours to run.

myDataflowOne: this dataflow tasks 1 hours and 30 minutes to run

myDataflowOne: this dataflow tasks 1 minute and 30 seconds to run.

If all three dataflows run, how many count towards the limit?

- A. 3
- B. 0
- C. 1
- D. 2

ANSWER: D

Explanation:

2 is correct because Tableau CRM dataflow-run limits count runs based on their duration. A dataflow run that takes at least two minutes counts toward the organization's daily dataflow-run allowance. The run lasting two hours and the run lasting one hour and 30 minutes both exceed that threshold, so each consumes one run from the limit. The run lasting one minute and 30 seconds is below two minutes and therefore does not count toward the daily limit. This duration-based treatment is useful because it prevents short, lightweight processing jobs from consuming the same run capacity as longer data-integration workloads. Administrators should still monitor all dataflow executions for operational purposes, including short runs, but only the qualifying longer runs are included when calculating this particular usage limit. Salesforce documents Tableau CRM dataflow behavior and limits in its CRM Analytics implementation guidance; see [Salesforce Help: Dataflow Limits](#) and [Trailhead: CRM Analytics Data Preparation](#).

QUESTION NO: 10

Universal Containers has a dashboard for Sales Managers. They need the ability to visualize the number of Closed Won opportunities by month, quarter, or year, and then display the result in a single chart. An Einstein Consultant created a static step to display three values: ClosedDate_month, ClosedDate_quarter, and ClosedDate_year.

What should the consultant do next?

- A. Use nested binding to update the grouping in the chart.
- B. Use selection binding to update the measure in the chart.
- C. Use result binding to update the measure in the chart.
- D. Use selection binding to update the grouping in the chart.

ANSWER: D

Explanation:

Use selection binding to update the grouping in the chart. The static step functions as a user-facing selector whose chosen value represents the field by which opportunity records should be grouped: ClosedDate_month, ClosedDate_quarter, or ClosedDate_year. A selection binding passes the currently selected value from that step into the chart query. The chart's group clause can therefore be dynamically set to the selected date-granularity field, allowing one chart to aggregate Closed Won opportunities at monthly, quarterly, or yearly levels without creating separate charts.

This design keeps the measure—such as the count of opportunities—constant while changing the query dimension used for aggregation. In CRM Analytics dashboard bindings, selection bindings are intended to use a value selected by a dashboard viewer to dynamically affect another step. The binding should reference the static step's selected value and insert it into the target chart's grouping definition. Salesforce documents the binding concepts and selection-driven interactions in the [CRM Analytics Dashboard Bindings Developer Guide](#). For additional dashboard-building context, see Salesforce's [CRM Analytics Dashboards documentation](#).

QUESTION NO: 11

In a dataset, there are multiple boolean fields. When displayed on any dashboard, the boolean fields should all be displayed in the same way: a value of true should result in the display of the word "Yes" in green; a value of false should result in the display of the word "No" in red.

In which two ways can this be accomplished?

- A. Create an XMD node in the dataflow to change the label and color of the values in the fields.
- B. In the explorer, select the boolean fields and use the "edit values" option on the fields to change values and colors.
- C. Download the XMD for the dataset and change the values and color for the fields and then upload to the dataset.
- D. In the explorer, select the boolean fields, then create a derived dimension that references the boolean fields, and then modify the values and colors of the derived dimension.

ANSWER: C

Explanation:

Downloading the XMD for the dataset and changing the values and color for the fields and then uploading to the dataset is the correct approach because extended metadata (XMD) controls dataset presentation metadata that is applied consistently wherever the dataset is used, including dashboards. In the XMD JSON, the relevant boolean fields can be configured with member-level metadata for their stored values. The `true` member can be given the display label `Yes` and a green color, while the `false` member can be given the display label `No` and a red color. Because this metadata belongs to the dataset rather than an individual dashboard widget, the configuration supplies the required reusable, standardized presentation across dashboard visualizations that use those fields. After editing, uploading the XMD makes the metadata available to CRM Analytics for the dataset. This is particularly appropriate when several fields need identical member labels and color treatment, since the configuration can be maintained centrally and deployed with the dataset metadata. Salesforce documents XMD as the mechanism for customizing dataset metadata and describes how it affects labels, formatting, and other display behavior in CRM Analytics. See the [Salesforce XMD developer documentation](#) and the [Salesforce Help documentation for extended metadata](#).

QUESTION NO: 12

Which three parameters are used in a compare table formula field to calculate and display a field as percentage value? (Choose three.)

- A. Length
- B. Column Name
- C. Column Color
- D. Format
- E. Calculation

ANSWER: A D E

Explanation:

In a CRM Analytics compare table, a formula field needs a **Calculation** to define the mathematical expression that produces the percentage. This calculation can divide one aggregated value by another or use the appropriate percentage-oriented calculation logic, depending on the business question being answered. The formula field's **Format** controls how the resulting numeric value is rendered to viewers; applying a percentage format displays a ratio such as 0.25 as 25%, rather than as a decimal. The **Length** parameter controls the display width allocated to the formula column in the compare table, ensuring that the percentage output is presented cleanly and consistently alongside the other columns. Together, Calculation determines the value, Format determines its percentage representation, and Length determines its table presentation. Salesforce's compare-table calculation guidance describes configuring formula calculations and their display properties

when exploring data in CRM Analytics. See the Trailhead unit on [calculating with the compare table](#) and the [Salesforce CRM Analytics table widget documentation](#).

QUESTION NO: 13

The Tableau CRM team at a company created a dataset based on the Sales_Target__c custom object. The VP of Sales reports seeing the message "No results found" when opening the dataset to explore it. Other users below the VP in the role hierarchy can see rows on the same dataset.

Which problem might be causing this issue?

- A. The Security Predicates set up at the dataset level are preventing the VP from seeing data.
- B. The Salesforce profile for the VP does not have read permission on some fields of the Sales_Target__c custom object.
- C. The dataset is inheriting sharing from Salesforce and the VP can see more than 3000 rows.
- D. The Salesforce profile for the VP does not have read permission on the Sales_Target__c custom object.

ANSWER: C

Explanation:

The dataset is inheriting sharing from Salesforce and the VP can see more than 3000 rows. Salesforce sharing inheritance applies the Salesforce record-sharing model when Tableau CRM queries the dataset. A senior user such as a VP commonly receives access to records owned by, or shared with, users below them in the role hierarchy. This can result in a much larger accessible record set than that of an individual contributor. For inherited sharing, Tableau CRM has a documented limit of 3,000 shared-record IDs that can be evaluated for a user. When the user's Salesforce sharing access exceeds that limit, Tableau CRM can be unable to return dataset rows, resulting in the "No results found" behavior even though users with narrower access can view data. The fact that lower-level users can see rows while the VP cannot is consistent with the VP exceeding the sharing-inheritance record limit because of broader hierarchical access. Review whether the dataset uses inherited sharing and, where appropriate, use a security predicate designed for the required access model instead. See Salesforce's [CRM Analytics dataset security documentation](#) and [inherited sharing documentation](#).

QUESTION NO: 14

A consultant is configuring dashboard bindings so that a user can select a region in a list widget and use that selection to filter other dashboard steps.

Which two statements about dashboard binding functions are true? Choose 2 answers

- A. The selection binding function returns values selected by a user in a widget.
- B. The result binding function accesses values returned by a step query.
- C. The selection binding function can be used only with static text widgets.
- D. The result binding function can access only Salesforce object records.

ANSWER: A B

Explanation:

The selection binding function returns values selected by a user in a widget. is correct because selection bindings use the current selection from an interactive widget, such as a list, chart, or table. A consultant can pass those selected values into another step's query to create responsive dashboard filtering.

The result binding function accesses values returned by a step query. is also correct. Result bindings reference the data produced by a query, rather than a user interaction. They are useful when a dashboard component needs to use a query result, such as a calculated value or a returned dimension member, in another query or widget property.

Bindings enable dashboard components to interact without requiring a separate dashboard for every view. See the [CRM Analytics Dashboard Bindings documentation](#).

QUESTION NO: 15

A company wants to allow users who belong to an account team to see all the Opportunities associated with that Account in Tableau CRM.

Which two actions accomplish this requirement? (Choose two.)

- A. Create a master-detail relationship between the Salesforce Account and Opportunity objects.
- B. Apply sharing inheritance.
- C. In the dataflow, extract the AccountTeamMember object and augment it with the Opportunity object using 'AccountId' as the join field and apply following security predicate: 'AccountTeamMember.UserId' == "\$User.Id".
- D. In the dataflow, extract the OpportunityTeamMember object and augment it with the Opportunity object using 'OpportunityId' as the join field and apply the following security predicate: 'OpportunityTeamMember.UserId' == "sUser.Id".

ANSWER: B C

Explanation:

Applying sharing inheritance and creating a dataset-level security predicate based on account-team membership together provide the required Tableau CRM access model. Sharing inheritance enables opportunity records to inherit the applicable access context from their parent account relationship, so the security model follows the account-to-opportunity hierarchy rather than requiring a separately maintained opportunity-team assignment for every opportunity.

The dataflow must also extract `AccountTeamMember` and augment it onto Opportunity using `AccountId`. This makes the account-team user identifier available on each relevant opportunity record in the resulting dataset. The predicate `'AccountTeamMember.UserId' == "$User.Id"` is evaluated for the currently logged-in user and returns only opportunity rows associated with accounts on whose teams that user appears. This is an appropriate use of row-level security because Tableau CRM predicates operate against fields materialized in the dataset after dataflow processing. The approach therefore gives account-team members access to all related opportunities while retaining record-level filtering for other users.

Salesforce documents CRM Analytics security predicates and inherited security behavior in the [CRM Analytics Security Guide](#). See also Salesforce Help on [security predicates](#).

QUESTION NO: 16

Before using bindings, you can try using facets to specify interactions between widgets.

- A. False
- B. True

ANSWER: B

Explanation:

True is correct. In CRM Analytics (formerly Tableau CRM/Einstein Analytics), faceting is the straightforward, declarative way to create common dashboard interactions. A facet lets a selection in one widget, such as a chart or list, filter the data displayed by other widgets. This approach is appropriate when the intended interaction is simply filtering related dashboard content based on a user's selection, and it can often be configured without editing dashboard JSON.

Bindings are intended for cases where the interaction requires more control than a standard facet provides. They can pass values or selections between dashboard components and dynamically alter queries, filters, or other widget properties. Because bindings involve JSON syntax and query structures, they are generally more complex to configure and maintain.

Salesforce guidance therefore supports starting with facets for standard filtering interactions and moving to bindings when the dashboard needs dynamic behavior that faceting cannot express. This keeps dashboard design simpler while still allowing advanced interactivity when required. See Salesforce's [CRM Analytics bindings introduction](#) and [dashboard faceting documentation](#).

QUESTION NO: 17

The Universal Containers company plans to upload target data from an external tool to Einstein Analytics so they can calculate the Sales team target attainments.

The target data changes every month, so the datasets need to be updated on a monthly basis. The target data is a CSV file that contains the Salesforce ID of the sales representative, the target amount, and the month of the target. For each sales representative, the file contains a target for every month of the current year as well as all previous years.

Based on this information, which operation should a consultant use with the Analytics External Data API to upload the file?

- A. Update
- B. Append
- C. Overwrite
- D. Upsert

ANSWER: C

Explanation:

Overwrite is the appropriate operation because each monthly CSV is a complete replacement snapshot of the target dataset. It includes target records for every sales representative for the current year and all prior years, rather than containing only newly created or changed rows. Using an overwrite operation replaces the existing dataset contents with the full contents of the newly uploaded file, ensuring that corrections to historical targets and updates to current-month targets are reflected consistently. This approach also prevents duplicate target records from accumulating across monthly loads, which is essential when calculating attainment measures by representative and month. The Analytics External Data API supports an overwrite action for this full-refresh loading pattern: the upload creates a new version of the dataset from the supplied external data. A complete, authoritative monthly file is therefore well suited to a scheduled overwrite process. Salesforce documents the external-data upload workflow and its dataset-processing operations in the [Analytics REST API External Data documentation](#). Additional implementation guidance for bringing external data into CRM Analytics is available in [Salesforce Help: Integrate External Data](#).

QUESTION NO: 18

You can get data into Einstein Discovery:

- A. By connecting a USB drive to your laptop
- B. By using an Einstein analytics dataset
- C. Only if your database administrator does it for you
- D. Only if Salesforce does it for you

ANSWER: B

Explanation:

By using an Einstein analytics dataset is correct because Einstein Discovery uses datasets in Salesforce CRM Analytics (formerly Tableau CRM/Einstein Analytics) as a primary source for creating stories and running predictions. A dataset is a prepared, tabular collection of records and fields that CRM Analytics can build from Salesforce objects, external data, or other supported data connections. Once the dataset is available, a consultant or business user can select it when creating an Einstein Discovery story, identify the outcome to analyze, and let Einstein Discovery evaluate patterns and generate

insights. This dataset-based approach is important because it provides a governed and reusable analytics layer: data can be refreshed, transformed, and secured before it is used for discovery. The dataset should contain sufficient historical records and an appropriate outcome field so Einstein Discovery can train a meaningful model. Salesforce's learning materials describe using CRM Analytics datasets as the foundation for Einstein Discovery analysis and explain the preparation required before creating a story. See [Salesforce Trailhead: Einstein Discovery Basics](#) and [Salesforce Help: Einstein Discovery Overview](#).

QUESTION NO: 19

Benito has Viewer access to a new Einstein Analytics app. He wants to build and edit a dashboard using the Seed Bank dataset in this new app. What must an admin do for Benito to do this?

- A. Grant Benito access to all datasets in the org.
- B. Assign Benito a permission set that contains the permission: Create and Edit Analytics Dashboards.
- C. Grant Benito Editor access to the app that contains the Seed Bank dataset and assign him a permission set that contains the Create and Edit Analytics Dashboards permission.
- D. Assign Benito a permission set that contains the permission: Create and Manage Analytics Apps.

ANSWER: C

Explanation:

Granting Benito Editor access to the app that contains the Seed Bank dataset, together with the *Create and Edit Analytics Dashboards* permission, gives him the required combination of app-level access and functional permission. Editor access allows a user to create and modify assets within the shared Analytics app, while the dashboard permission authorizes the user to create and edit Analytics dashboards. Access to the relevant app also makes the app's available data resources, including the Seed Bank dataset, usable in that app context, subject to any applicable dataset security predicates. This follows the Analytics sharing model: app roles control what a person can do with content in an app, and user permissions control whether they can author dashboards. An administrator should assign the dashboard-authoring permission through a permission set and share the app with Benito at the Editor level. Salesforce's Analytics app-sharing guidance describes the distinction between viewer and editor capabilities, and the Trailhead module on sharing Analytics apps covers the permissions needed for users to build content. See [Share CRM Analytics Apps](#) and [Salesforce Help: Share Analytics Apps](#).

QUESTION NO: 20

A consultant has deployed an Einstein Discovery prediction to Salesforce. Business users want prediction results to be useful and trustworthy after deployment.

Which two actions should the consultant recommend? Choose 2 answers

- A. Monitor prediction quality and business outcomes over time.
- B. Provide users with the prediction explanation and top factors.
- C. Prevent users from viewing the prediction outcome after deployment.
- D. Assume that a deployed prediction remains accurate without further review.

ANSWER: A B

Explanation:

Monitor prediction quality and business outcomes over time. is correct because data patterns, business processes, and outcome distributions can change after deployment. Reviewing prediction behavior and actual outcomes helps the organization identify when the story should be reassessed or retrained.

Provide users with the prediction explanation and top factors. is also correct because explanations help users understand the major factors influencing an individual prediction. This supports informed business decisions rather than treating a prediction as an unexplained score.

Predictions should complement human judgment and established business processes. Salesforce describes prediction explanations, deployment, and ongoing story management in the [Einstein Discovery Basics](#) Trailhead module.

QUESTION NO: 21

In an org, some of the Activity records do not have related Opportunity records. Additionally, there are Opportunity records that do not have related Activity records.

Which SAQL statement will fetch all Opportunities, even if they don't have associated Activities, and fetch related Activity data, if it exists?

- A. `q = cogroup Opportunity by 'Id', Activities by 'OpportunityId';`
- B. `q = cogroup Opportunity by 'Id', Activities by 'OpportunityId' left;`
- C. `q = cogroup Opportunity by 'Id' right, Activities by 'OpportunityId';`
- D. `q = cogroup Opportunity by 'Id' left, Activities by 'OpportunityId';`

ANSWER: D

Explanation:

`q = cogroup Opportunity by 'Id' left, Activities by 'OpportunityId';` is correct because the `left` modifier makes the Opportunity stream the preserved side of the cogroup operation. SAQL therefore retains every grouped Opportunity ID, including IDs for Opportunities that have no matching Activity record. Where an Activity has an `OpportunityId` that matches an Opportunity `Id`, the corresponding Activity group is available in the cogroup result and can be used in subsequent `foreach`, aggregation, or projection logic.

This approach is appropriate when Opportunities are the required population for a dashboard or analysis and Activity information is supplemental. The cogroup first groups each stream using its specified key and then combines the groups; the left-preserving behavior ensures that an empty matching Activity group does not eliminate an Opportunity from the result. Activities without a related Opportunity do not affect the required result population, because the business requirement is to return all Opportunities and associated Activity details when present. Salesforce documents cogroup as the SAQL mechanism for joining grouped streams and supports outer-join behavior through join modifiers such as `left`. See the [Salesforce SAQL cogroup documentation](#) and the [Trailhead SAQL joins module](#).

QUESTION NO: 22

A consultant is configuring a dashboard that contains a list widget for product family. Users can select one or more product families and use the selection to filter a chart step.

Which two binding functions can the consultant use to pass values from the selection widget? Choose 2 answers

- A. `cell()`
- B. `column()`
- C. `aggregate()`
- D. `register()`

ANSWER: A B

Explanation:

column() is correct because it returns the values from a specified column in the selection step result. It is useful when a user can select multiple product-family rows and the target query needs the collection of selected values for a filter.

cell() is also correct because it returns an individual value from a specified row and column in a step result. It is appropriate when the binding needs a single selected value, such as the value from a single-select widget or a selected row used to drive a KPI query.

Dashboard bindings dynamically pass widget values into queries, filters, and widget properties. Salesforce documents binding functions and dashboard JSON configuration in the [CRM Analytics Dashboard Bindings documentation](#).

QUESTION NO: 23

The Universal Containers Sales team is looking to build a dashboard that shows the total revenue from their top 10 accounts (based on revenue) and make it dynamic with the filters of the dashboard. What two steps are required to accomplish this? (Choose two.)

- A. Create a query with a selection based interaction as a filter using the limited query to pass in the 10 Account Id's.
- B. Create a query and sort descending by revenue and limit it to 10 results.
- C. Create a query that groups by Account Id and sorts descending by revenue and limit it to 10 results.
- D. Create a query with a results based interaction as a filter using the limited query to pass in the 10 Account Id's.

ANSWER: C D

Explanation:

To identify the top ten accounts correctly, the query must aggregate revenue at the account level. "Create a query that groups by Account Id and sorts descending by revenue and limit it to 10 results." produces one result per account, orders those aggregate revenue totals from highest to lowest, and restricts the result set to the ten accounts that qualify. This limited query is the source of the dynamic top-account list.

"Create a query with a results based interaction as a filter using the limited query to pass in the 10 Account Id's." then uses the Account IDs returned by that top-ten query to filter the query that displays total revenue. Because dashboard filters are applied to the relevant query steps, the top-ten query recalculates when a dashboard filter changes. Its returned Account IDs are consequently refreshed and passed to the revenue query, keeping the displayed accounts and revenue aligned with the current filter context. This design separates ranking logic from the dependent display query while preserving interactivity. Salesforce documents dashboard bindings, including result bindings that use query output in another step, in the [CRM Analytics Dashboard Bindings Developer Guide](#). See also Salesforce's [Dashboard Bindings documentation](#) for binding behavior and supported use cases.

QUESTION NO: 24

A company wants to allow users who belong to an account team to see all the Opportunities associated with that Account in Einstein Analytics.

Which two actions accomplish this requirement? Choose 2 answers

- A. Apply sharing inheritance.
- B. Create a master-detail relationship between the Salesforce Account and Opportunity objects.
- C. In the dataflow, extract the AccountTeamMember object and augment it with the Opportunity object using 'AccountId' as the join field and apply following security predicate: 'AccountTeamMember.UserId' == '\$User.Id'.
- D. In the dataflow, extract the OpportunityTeamMember object and augment it with the Opportunity object using 'OpportunityId' as the join field and apply the following security predicate: 'OpportunityTeamMember.UserId' == '\$User.Id'.

ANSWER: A C

Explanation:

Apply sharing inheritance. is appropriate because sharing inheritance enables CRM Analytics to respect the parent Account's record-access model when presenting related Opportunity data. This keeps the analytics access model aligned with Salesforce sharing, so access granted through an account team can be inherited by data associated with that Account rather than requiring a separately maintained user-by-opportunity access list.

In the dataflow, extract the AccountTeamMember object and augment it with the Opportunity object using 'AccountId' as the join

field and apply following security predicate: 'AccountTeamMember.UserId' == "\$User.Id". provides the required record-level dataset security. AccountTeamMember supplies the relationship between a user and an Account. Augmenting that data onto Opportunity records by Account Id makes the account-team user identifier available on each related opportunity row. The security predicate then evaluates the running user's Salesforce user Id and returns only opportunities whose Account includes that user on its team. Together, inherited sharing and the predicate ensure account-team membership drives visibility for every Opportunity belonging to the corresponding Account. Salesforce describes CRM Analytics dataset security and predicates in its [CRM Analytics Security Trailhead module](#) and provides related implementation guidance in [CRM Analytics security documentation](#).

QUESTION NO: 25

A company uses a CRM Analytics dataset to display account information to regional managers. Each dataset row contains a Region field, and each manager's Salesforce User record contains a matching Region field.

Which security predicate enforces row-level access so that managers see only records for their own region?

- A. 'Region' == "\$User.Region"
- B. 'Region' == "\$User.Id"
- C. 'Region' == "\$Organization.Id"
- D. 'Region' in ["\$User.ProfileId"]

ANSWER: A

Explanation:

'Region' == "\$User.Region" is correct because a security predicate compares a field in each dataset row with an attribute of the running user. When the dataset Region value matches the logged-in user's Region value, CRM Analytics returns the row; otherwise, it excludes the row from the user's query results.

This design allows one shared dataset and dashboard to provide region-specific views without creating separate datasets for each manager. The consultant must ensure that the user field is populated consistently and that the dataset field uses values compatible with the user attribute. Salesforce documents security predicates and user-based row-level access in the [Secure Data in a CRM Analytics Dataset](#) Trailhead unit.

QUESTION NO: 26

Einstein Recommends Improvements in 2 Areas

Strongest Predictors

Outliers

Uncommonly large or small numbers, potentially from data entry errors, or rare situations, can produce misleading charts or predictions.

DailyQuantity

Removing values below 0 and above 2,489

- ☐ Remove outliers (Recommended)
- ☒ Do Nothing

Cancel

Customize Story

Create New Story

A consultant has created a story to maximize the daily sales quantity of consumer products in stores. After creating a story, the consultant is presented with this recommendation by Einstein Discovery (see graphic).

What are two appropriate actions to take? (Choose two.)

- A. Remove the outliers as suggested by Einstein and deploy the model.
- B. Remove the outliers as suggested by Einstein, and verify using model metrics and story insights if the quality improved.
- C. Discuss with the client if values below 0 and above 2,489 are so uncommon that they should perhaps be left out of the story.
- D. Manually remove the sales below 0 (negative sales must be a data issue), but keep the large values (the more data, the better the model will be).

ANSWER: B C

Explanation:

Removing the outliers as suggested by Einstein, and verifying using model metrics and story insights if the quality improved, is appropriate because an outlier recommendation is an analytical prompt rather than an instruction to deploy immediately. The consultant should assess the impact of the data change by rebuilding or comparing the story and reviewing model quality indicators and insights. This validates that excluding unusual records improves the predictive usefulness of the model without unintentionally removing meaningful business variation.

Discussing with the client whether values below 0 and above 2,489 are sufficiently uncommon to exclude is also essential. Einstein Discovery can identify statistically unusual values, but it cannot determine whether those values represent valid business events, such as returns, corrections, exceptional promotions, or unusually large orders. The client supplies the business context needed to decide whether the records are errors, exceptional but valid observations, or cases outside the scope of the prediction objective. Combining Einstein's statistical recommendation with stakeholder validation and subsequent model-quality review is the sound data-preparation practice. See Salesforce's [Einstein Discovery Basics](#) and [Einstein Discovery data preparation documentation](#).

QUESTION NO: 27

A consultant is working with the client to define goals and metrics for a dashboard that will be built. The consultant is in the design phase prior to starting the actual build. Which two design considerations should be determined at this time? (Choose two.)

- A. How to get as much information as possible onto one page.
- B. How to quickly convey key information.
- C. Which colors to use and how to use them.
- D. Which charts to use to best support the defined goal.

ANSWER: B D

Explanation:

How to quickly convey key information is a core dashboard-design decision because a dashboard should let its intended audience understand performance, trends, and exceptions with minimal effort. During discovery and design, the consultant should identify the audience's primary questions, the decisions the dashboard must support, and the most important metrics to emphasize. This creates a focused information hierarchy before any build work begins.

Which charts to use to best support the defined goal should also be determined during design. Chart selection must follow the analytical purpose of each metric: for example, comparisons, changes over time, distributions, relationships, or progress toward a target. Choosing visualizations based on the agreed business goal helps ensure that the dashboard communicates the intended insight clearly and supports action. Tableau's guidance similarly emphasizes starting with the question being answered and selecting visual encodings that make the relevant pattern easy to see. These decisions form the dashboard's structure and communication strategy, allowing implementation details to follow a deliberate design rather than being driven by available screen space or incidental styling choices.

Reference: [Tableau Dashboard Best Practices](#) and [Tableau Visual Analysis Best Practices](#).

QUESTION NO: 28

How many external files can be uploaded to Einstein Analytics on rolling 24 hour basis?

- A. 100 files per dataset
- B. 300 files per dataset
- C. 500 files per dataset
- D. 50 files Per dataset

ANSWER: D

Explanation:

50 files Per dataset is correct. Einstein Analytics, now referred to as CRM Analytics/Tableau CRM in much of Salesforce's current documentation, applies an upload limit of 50 external-data files for a dataset during a rolling 24-hour period. A rolling period is measured continuously from each upload rather than resetting at a fixed daily boundary. This limit is relevant when data is loaded from external sources through the Analytics external-data upload process, including API-driven upload jobs. Teams that need to refresh data more frequently should design their integration and dataset strategy around this per-dataset limit, such as by consolidating data into fewer uploads where appropriate and scheduling loads with the rolling window in mind. Salesforce documents the external-data resource and its upload workflow in the [CRM Analytics REST API Developer Guide](#). For broader platform limits and guidance on monitoring usage, consult Salesforce's [CRM Analytics Limits documentation](#).