



Salesforce Tableau CRM Einstein Discovery Consultant

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

Version Demo

Total Demo Questions: 34

Total Premium Questions: 341

Buy Premium PDF

<https://exam4future.com>

support@exam4future.com

exam4future.com

support@exam4future.com

Topic Break Down

Topic	No. of Questions
Topic 1, Analytics Discovery	34
Topic 2, Data	90
Topic 3, Dashboard	80
Topic 4, Security	60
Topic 5, Einstein Discovery	77
Total	341

QUESTION NO: 1

A company's Salesforce org has multi-currency enabled. This company's business intelligence team used Einstein Analytics to build a dataflow that creates a dataset, "OpportunityDataSet". This dataset is populated with data extracted from the standard object, Opportunity. One of the extracted fields is the standard field, Amount.

If a user explores the "OpportunityDataSet" in Einstein Analytics, in which currency will the Amount values be shown?

- A. In the connected user's currency
- B. In the integration user's currency
- C. In the currency that is set on the "currency" attribute in the dataset
- D. In the currency that is set on the "currency" attribute in the dataflow

ANSWER: B

Explanation:

In the integration user's currency is correct. When CRM Analytics (formerly Einstein Analytics) runs a dataflow that extracts currency values from Salesforce in an organization with multiple currencies enabled, the extraction runs under the dataflow's integration user. Salesforce uses that user's configured currency context when supplying currency values to the analytics dataflow. Consequently, the resulting dataset's extracted `Amount` values reflect the integration user's currency rather than being dynamically converted for each person who later opens an exploration.

This behavior matters when designing datasets that include Opportunity amounts, forecasts, or other monetary fields. The integration user should be deliberately configured with the currency appropriate for the reporting and analysis use case before the dataflow is run. If the organization needs values normalized to a common reporting currency, the analytics implementation should account for that requirement during data preparation and use the appropriate currency fields and conversion logic.

Salesforce documents the behavior of currency field values and currency conversion in its platform data model and analytics data-preparation guidance. See [Salesforce API Field Types](#) and [CRM Analytics Data Preparation](#).

QUESTION NO: 2

What are two core design principles to consider when building Analytics apps or dashboards? Choose 2 answers

- A. Emphasis: make sure they have space for important information, such as headlines and key charts.
- B. Clarity: make sure they are uncluttered and easy to interpret.
- C. Consistency: make sure they have a sense of familiarity to strengthen your users' ease of use.
- D. Balance: make sure they have a balance of different charts to get a more interesting design.

ANSWER: B C

Explanation:

Clarity: make sure they are uncluttered and easy to interpret. is a core design principle because an Analytics app or dashboard must enable people to understand information and act on it quickly. Clear layouts, meaningful labels, and an intentional visual hierarchy let users identify the key message without having to decipher unnecessary visual complexity. Clarity is especially important when users are reviewing business metrics under time constraints.

Consistency: make sure they have a sense of familiarity to strengthen your users' ease of use. is also a core principle. Applying consistent chart conventions, colors, placement, terminology, and interactions across dashboards helps users build familiarity. Once users learn how one page communicates information, they can more easily navigate and interpret the rest of the app. This lowers cognitive effort and supports a trustworthy, usable analytics experience. Salesforce's Analytics App Design guidance identifies clarity and consistency as foundational principles for designing effective analytics experiences. See the [Salesforce Trailhead: Principles of Good Design](#) and the [Salesforce dashboard design best practices](#).

QUESTION NO: 3

A small Business Intelligence team is overwhelmed with multiple requests to create new dashboards. An option is to use layout templates to help with dashboard development.

What are two advantages of using layout templates? Choose 2 answers

- A. Layout templates are equipped with design best practice, such as optimally positioned KPIs and filters
- B. Layout templates are fast to create because they are fixed and read-only
- C. Layout templates offer a consistent end user experience
- D. Layout templates offer a universal layout thus reducing the hours spent on design

ANSWER: C D

Explanation:

Layout templates offer a consistent end user experience because a team can reuse a standardized dashboard structure across dashboard builds. This promotes predictable placement and behavior of common dashboard elements, so users do not need to relearn navigation and interaction patterns for every new dashboard. Consistency is especially valuable when many dashboards are delivered to the same business audience, because it supports adoption and makes dashboard content easier to scan.

Layout templates offer a universal layout thus reducing the hours spent on design because authors can begin with a reusable arrangement rather than designing the dashboard canvas from scratch for each request. The template establishes the overall structure while allowing the dashboard creator to focus effort on the relevant data, widgets, and business-specific content. This accelerates delivery for a small BI team and provides a repeatable foundation for dashboard development. Salesforce describes dashboard design and layout capabilities in its CRM Analytics dashboard learning materials, including approaches for arranging and reusing dashboard content. See [Salesforce Trailhead: CRM Analytics Dashboards](#) and [Salesforce Help: CRM Analytics Dashboards](#).

QUESTION NO: 4

The model quality metrics of an Einstein Discovery story indicate that the GINI coefficient in the four folds are 0.82, 0.83, 0.84, and 0.75, respectively.

Which two actions should a consultant take? Choose 2 answers

- A. Deploy the story, because the variation in the metrics is within the normal range.
- B. Confirm that the overall GINI coefficient is good prior to deploying the story.
- C. Research and check the dataset for outliers in the target field and the main predictors that are shown on top of the story.
- D. Do not deploy the story immediately, and research why one fold is performing worse than the others.

ANSWER: B D

Explanation:

Confirm that the overall GINI coefficient is good prior to deploying the story. The overall GINI value is the primary indicator of the story's ability to distinguish between outcomes across the validation process. A consultant should verify that this aggregate measure meets the organization's requirements and supports the intended decision-making use case before enabling predictions in production. Salesforce presents model-quality metrics so practitioners can assess predictive performance and confidence before deployment.

Do not deploy the story immediately, and research why one fold is performing worse than the others. Cross-validation folds are separate samples of the training data, so their quality metrics should be reasonably consistent. Although the first three folds have closely grouped GINI values, the lower 0.75 result warrants investigation to ensure that the model is stable and will generalize reliably. The consultant should examine whether characteristics of the records in that validation fold, such as

data distribution or influential feature values, are driving the reduced performance. Addressing the reason for the inconsistency before deployment supports a more dependable Einstein Discovery model.

See Salesforce guidance on [Einstein Discovery model quality](#) and the [Einstein Discovery basics](#).

QUESTION NO: 5

A company created a dataset called myExternalDataset. It contains a dimension, called UserRegion, that is populated by the standard User object's picklist: Region.

The company has security requirements for this dataset:

1. Only a subset of users should be able to access the dataset.
2. Users who can access the dataset should see only data that belongs to their region.

In which two ways can an Einstein Analytics consultant help them implement these requirements? Choose 2 answers

- A.** Save the dataset under "My Private App" so only a subset of selected users can access the dataset.
- B.** Use a security predicate in the dataset so users can see only rows belonging to their respective regions.
- C.** Create a sharing rule in Salesforce based on the Region field in the user object and set the dataset to inherit sharing from Salesforce.
- D.** Save the dataset under a new application and give "viewer" access to a public group containing the list of users allowed to see the dataset.

ANSWER: B D

Explanation:

Use a security predicate in the dataset so users can see only rows belonging to their respective regions. CRM Analytics security predicates provide row-level security by evaluating attributes of the current user and matching them to values in dataset rows. In this case, the predicate can compare the running user's User.Region value with the UserRegion dimension in myExternalDataset. Consequently, each authorized user queries the same dataset but receives only records for the region associated with that user's Salesforce User record. This approach is designed for datasets whose access needs differ by user or group while retaining a centrally managed dataset.

Save the dataset under a new application and give "viewer" access to a public group containing the list of users allowed to see the dataset. Application sharing controls who can access the app and its contained assets, including the dataset. Granting Viewer access to a public group establishes the required audience without granting access individually. Combining application-level sharing for access entitlement with a dataset security predicate for region-based row filtering satisfies both requirements independently and consistently. Salesforce documents these capabilities in its [CRM Analytics app sharing guidance](#) and [security predicates guidance](#).

QUESTION NO: 6

In Einstein Discovery:

- A.** 'What Is The Difference' insights are comparative insights that help you better understand the relationships between explanatory variables and the goal (target outcome variable) in your story. These insights, based on a statistical analysis of your dataset, help you figure out which factors contribute to the biggest changes in the outcome variable. Einstein Discovery uses waterfall charts to help you visualize comparisons in What Is The Difference insights.
- B.** 'Why It Happened' insights help you take a deeper look into the exact factors that led to an outcome. Why It Happened s/Q insights drill deeper into the various factors that contributed to your story's goal. These insights are based on a statistical analysis of your dataset. Einstein Discovery uses waterfall charts to help you visualize Why It Happened insights.
- C.** 'Predictions and Improvements' insights help you explore what might happen in the future. For example, you can interactively perform "what if analyses in your story. Einstein Discovery provides you with predictions and suggested improvements based on a statistical analysis of your dataset and predictive analytics. To help you visualize these insights,

Einstein Discovery uses:

- waterfall charts for predictions
- bar charts for suggested improvements

D. 'What Happened' insights are the primary insights in your story. They are descriptive insights that help you explore, at an y/ Q overview level, what factors contributed to the outcome, based on a statistical analysis of your dataset. Einstein Discovery uses bar charts to help you visualize What Happened insights.

ANSWER: A C D

Explanation:

Einstein Discovery organizes story insights around distinct analytical questions, and the correct statements accurately describe three of those insight types. *What Happened* provides descriptive, high-level insight into the outcome and the factors associated with it. This initial exploration is presented with bar charts, making it easy to see how explanatory variables relate to the story goal. *What Is The Difference* is comparative: it evaluates differences between groups or values and identifies factors associated with the largest changes in the target outcome. Waterfall charts are useful here because they show the relative contribution of factors to the comparison. *Predictions and Improvements* applies the predictive model to estimate future outcomes and lets users perform interactive what-if analysis. Prediction results are visualized in waterfall charts, while suggested improvements are shown in bar charts so users can identify the changes most likely to improve the predicted result. These insight categories combine descriptive, comparative, and predictive analysis within an Einstein Discovery story. See Salesforce's [Einstein Discovery Basics](#) Trailhead module and the [Einstein Discovery documentation](#) for more detail.

QUESTION NO: 7

A Tableau CRM consultant has just completed deployment of an analytic app containing a recipe plus several datasets and dashboards. While conducting post deployment a smoke test, the new datasets don't seem to have migrated.

What post migration step has likely been forgotten?

What post migration step has likely been forgotten?

- A. Apply security predicates on datasets
- B. Run the recipe
- C. Provide read access on the datasets
- D. Go to Analytics Settings in Setup and approve the deployment

ANSWER: B

Explanation:

Run the recipe is the required post-migration action. In Tableau CRM, a recipe is a data-preparation process that reads source data, applies transformations, and writes its output to one or more datasets. Deploying an analytic app migrates the recipe definition, its associated assets, and dashboard metadata, but the deployment itself does not necessarily execute the recipe in the target org. Therefore, the output datasets expected from that recipe are not created or refreshed until the recipe runs successfully in the destination environment.

As part of smoke testing, the consultant should open Analytics Studio in the target org, locate the deployed recipe, run it, and confirm that its job completes successfully. After the run, the consultant can verify that the recipe output datasets exist, contain the expected records and fields, and can supply data to the deployed dashboards. If the datasets must remain current after go-live, the recipe should also be given an appropriate refresh schedule. Salesforce describes recipes as tools for combining, transforming, and outputting data into datasets in its [Tableau CRM Recipes Trailhead module](#). See also Salesforce's [recipe documentation](#) for recipe processing concepts.

QUESTION NO: 8

In the Insights list, insights are ranked in what order?

- A. Time
- B. Geospatial
- C. Statistical significance
- D. Alphabetical
- E. There is no order

ANSWER: C

Explanation:

Statistical significance is correct. Einstein Discovery evaluates patterns found in the data and prioritizes the insights that have the strongest statistical evidence. This ranking helps users focus first on findings that are least likely to be attributable to random variation in the data and are therefore more reliable for decision-making. Statistical significance is assessed through the platform's underlying statistical analysis of the relationship or difference identified by an insight. In practice, the Insights list is designed as a prioritized review experience: rather than merely presenting findings in a neutral display order, it surfaces the most meaningful statistically supported discoveries first. This is particularly useful when an analysis generates many potential drivers, segments, comparisons, or correlations, because it directs attention toward results with the greatest analytical validity. Salesforce describes Einstein Discovery as using statistical analysis and machine learning to uncover meaningful insights and explain outcomes. For background on how Einstein Discovery identifies and presents data-driven findings, see the [Einstein Discovery Basics Trailhead module](#) and Salesforce's [Einstein Discovery documentation](#).

QUESTION NO: 9

The Universal Containers company uses a Microsoft Azure SQL Data Warehouse to gather information about sales representatives' objectives. They want to use Einstein Analytics to gain insights from this data; therefore, they need to automatically load the warehouse's data into an Einstein Analytics dataset on a daily basis. The data also needs to be transformed and merged with data from the company's org.

Which three Einstein Analytics user interface features can be used to complete these requirements? Choose 3 answers

- A. Dataflow
- B. Microsoft Azure SQL Data Warehouse Connector
- C. Einstein Analytics Connector for Excel
- D. Recipes
- E. Analytics REST API

ANSWER: A B D

Explanation:

Dataflow, **Microsoft Azure SQL Data Warehouse Connector**, and **Recipes** together provide the required user-interface-based integration capabilities. The Microsoft Azure SQL Data Warehouse Connector establishes the external connection and makes the warehouse tables available as a source in CRM Analytics. This lets the consultant retrieve the sales-objective data without building a custom API integration. A dataflow can be configured to run on a schedule, including daily, and is designed for repeatable ingestion, transformation, and augmentation of analytics data. It can combine the external source with Salesforce org data and write the resulting records to a CRM Analytics dataset.

Recipes provide the visual data-preparation experience for selecting fields, applying transformations, joining or augmenting data, and producing an output dataset. A recipe can also be scheduled for recurring runs, making it suitable when the transformation and merge logic is more conveniently managed through the point-and-click preparation interface. In practice, the connector supplies the Azure source, while either a dataflow or recipe can orchestrate and schedule the preparation work; both are valid UI features that address the stated loading and transformation requirements. Salesforce documents

QUESTION NO: 10

A consultant is asked to create a dashboard for sales to see the forecasted sales per product family. After gathering the requirements, the consultant decides to leverage the timeseries SAQL statement to predict sales.

Which two actions must be performed when writing the SAQL query?

Choose 2 answers

- A. Determine if the timeseries required parameter, prediction interval, should be 80 or 95.
- B. Investigate which prediction model is best for the data available.
- C. Validate that there is enough data in the dataset to make a good prediction.
- D. Include the fill SAQL statement to fill any missing dates.

ANSWER: C D

Explanation:

Validate that there is enough data in the dataset to make a good prediction. is required because a time-series forecast depends on a sufficient history of consistently recorded observations. The `timeseries` statement analyzes historical measure values across time to detect trend and seasonal patterns, then extends those patterns into forecast periods. Before building the query, the consultant should confirm that each product family has enough usable date-based data to support a meaningful forecast rather than relying on sparse or very short histories.

Include the fill SAQL statement to fill any missing dates. is also required when the source data does not contain every expected time period. Forecasting operates on an evenly spaced time series. Using `fill` to generate missing date buckets establishes a continuous timeline and supplies an appropriate value for periods that have no underlying rows, allowing the time-series operation to process the sequence correctly. This is particularly important for sales data, where some product families might have no recorded sales in a given week or month even though that period must remain represented in the forecasting series.

Salesforce documents the SAQL time-series forecasting syntax and its use with date-based measures in the [SAQL timeseries statement reference](#). For additional SAQL concepts, see the [Salesforce Trailhead SAQL module](#).

QUESTION NO: 11

Permissions of which 2 internal users are used by Analytics to gain access to Salesforce data?

- A. Each user has a corresponding Profile assigned (Analytics Cloud Integration User and Analytics Cloud Security User profiles).
- B. When you query a dataset that has row-level security based on the User object, Analytics uses the permissions of the Security User to access the User object and its fields.
- C. Analytics uses the permissions of the Integration User to extract data from Salesforce objects and fields when a dataflow job runs.

ANSWER: B C

Explanation:

Analytics uses two internal users for distinct data-access operations: the Analytics Cloud Integration User and the Analytics Cloud Security User. The Integration User is the identity used when a dataflow runs to extract and process Salesforce object and field data. Consequently, object-level and field-level access needed by a dataflow must be available through the

Integration User's permissions. The Security User is used when Analytics evaluates dataset row-level security that references the Salesforce User object. Its permissions determine whether Analytics can access the User object and the specific User fields required to evaluate that security predicate. These separate identities let Analytics perform scheduled ingestion and enforce user-based dataset security without relying on the permissions of whichever person is viewing the dashboard. Configure and validate access for both internal users when dataflows or User-based row-level security fail because of insufficient access. Salesforce's guidance on data security and dataset security is available in the [CRM Analytics security documentation](#) and the [CRM Analytics dashboard and dataset security documentation](#).

QUESTION NO: 12

Which permission gives you view-only access to Analytics?

- A. Use Analytics
- B. Manage Analytics
- C. Create and Edit Analytics Dashboard
- D. None of above

ANSWER: A

Explanation:

Use Analytics gives a user access to CRM Analytics (formerly Tableau CRM/Einstein Analytics) for consuming analytics content. In a view-only scenario, this permission lets users open the Analytics experience and view dashboards, lenses, and other content that has been shared with them, subject to the relevant app, folder, and underlying Salesforce data-access permissions. It is the foundational user permission for people who need to interact with published analytics rather than administer the Analytics environment or build and modify analytics assets.

Access remains governed by sharing and security. For example, a user must have access to a dashboard's app or folder and must be permitted to see the source-record data returned by the dashboard. This supports distributing dashboards broadly while preserving Salesforce's normal record-level security model. Administrators assign Analytics permissions through permission sets or profiles, alongside the appropriate CRM Analytics license or permission set license, depending on the organization's licensing configuration. Salesforce's CRM Analytics setup guidance and permission documentation describe **Use Analytics** as the permission that provides access for Analytics users: [Salesforce Help: CRM Analytics User Permissions](#) and [Salesforce Help: Set Up CRM Analytics](#).

QUESTION NO: 13

A Senior Sales Business Analyst asks for a dashboard that contains fiscal year product opportunities. Information is maintained in a spreadsheet which is comprehensive, but its contents often need to be explained. After sketching the dashboard, the consultant is ready to start building.

Which concept should be applied to this dashboard?

- A. Limit widget actions and exploration so users will focus only on the high-level information.
- B. Use the different chart types to make the dashboard interesting and appealing.
- C. Design the dashboard with the desktop in mind and use the same layout for mobile devices for consistency.
- D. Use charts to help users ask questions, not illustrate a conclusion.

ANSWER: D

Explanation:

Use charts to help users ask questions, not illustrate a conclusion. is the appropriate dashboard-design concept. The source spreadsheet is comprehensive but requires explanation, which indicates that simply reproducing its details in a dashboard would not make the information readily understandable or actionable. The dashboard should instead present

opportunity data in visual forms that help sales stakeholders identify patterns, notice exceptions, compare products, and explore the drivers behind results for the fiscal year.

For example, a chart showing pipeline by product and fiscal period can prompt useful questions about which products are growing, where opportunity volume is declining, or whether a product's pipeline has sufficient value for its sales target. Well-designed interactive dashboard components then allow users to investigate those questions with filters, selections, and drill-downs as appropriate. This supports analysis and informed decision-making rather than attempting to use visuals to assert a predetermined narrative.

Salesforce dashboard guidance emphasizes designing visualizations around the questions users need to answer and making dashboards focused, clear, and actionable. See Salesforce's [dashboard build tips](#) and the [CRM Analytics dashboards Trailhead module](#).

QUESTION NO: 14

A consultant is creating a Churn Prediction model to identify customers who are not likely to renew their contract.

What is the appropriate action to take?

- A. Replace nulls in the Churn Reason field as 'No reason given',
- B. Exclude active customers (customer who have not churned) from the training dataset since their Churned) from the training dataset since their Churn Reason field are nulls.
- C. For the Churn Reason field, enable the "This variable contains sensitive data" box because it may contain sensitive customer behavior.
- D. Exclude the Churn reason field from the dataset.

ANSWER: D

Explanation:

Exclude the Churn reason field from the dataset. is appropriate because a churn-reason value is generally captured only after a customer has already churned or decided not to renew. Using that field as a predictor would introduce target leakage: the model would learn from information that is unavailable when the business needs a prediction for an active customer. This can produce deceptively strong training and validation metrics but will not support reliable, actionable predictions before renewal decisions occur.

For a churn model, input fields must represent information known at the point at which the prediction will be generated, such as account attributes, product usage, support interactions, contract details, and historical engagement measures. The outcome should represent whether the customer churned, while post-churn information must be excluded from the modeling data. This aligns with Einstein Discovery data-preparation practices: select relevant predictors and ensure the data used to train a model is available at prediction time. See Salesforce's [Einstein Discovery data preparation guidance](#) and the [Einstein Discovery basics module](#).

QUESTION NO: 15

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:

The statement `q = cogroup Opportunity by 'Id' left, Activities by 'OpportunityId';` is correct because it performs a left cogroup using Opportunity as the dataset on the preserved, or left, side. A cogroup combines streams based on their grouping keys, and the `left` keyword ensures that every group from the Opportunity stream remains in the result. Therefore, an Opportunity is returned even when no Activity rows have a matching opportunity lookup value. For Opportunities that do have matching Activities, the corresponding Activity group is available for subsequent SAQL operations, such as `foreach`, `flatten`, aggregations, and projections.

This join direction directly matches the reporting requirement: Opportunities are the complete population to retain, while Activities are optional related data. The grouping key uses the Opportunity record identifier on the Opportunity stream and the Activity-side opportunity relationship field on the Activities stream. Salesforce documents `cogroup` as the SAQL operation used to group multiple streams by related keys, including support for join modifiers that control which unmatched groups are retained. See the [Salesforce SAQL cogroup reference](#) and the [SAQL Developer Guide](#).

QUESTION NO: 16

A consultant is preparing a dataflow that enriches individual Opportunity records with Account attributes and also combines two transaction datasets that have the same schema.

Which two statements correctly describe the appropriate dataflow transformations? Choose 2 answers

- A. Use augment to add Account attributes to Opportunity records.
- B. Use append to combine transaction datasets with matching schemas.
- C. Use append to add Account attributes to each Opportunity record.
- D. Use augment to stack all transaction rows from both datasets.

ANSWER: A B

Explanation:

Use augment to add Account attributes to Opportunity records because augment performs a lookup-style enrichment. It retains the grain of the left stream, which in this case is the individual Opportunity record, while adding matching fields from the Account stream.

Use append to combine transaction datasets with matching schemas because append combines the records from compatible streams into one output stream. It is used when datasets represent similar row types and the requirement is to stack their rows rather than look up fields through a relationship.

Using augment for two transaction datasets would require a lookup relationship and would not append all rows. Using append to add Account fields would not match the Account records to the related Opportunities. See [CRM Analytics Dataflow Transformations](#).

QUESTION NO: 17

Max story creations per org per month

- A. 1000 (can purchase more)
- B. 500 (can purchase more)
- C. 200 (can purchase more)
- D. 500

ANSWER: B

Explanation:

500 (can purchase more) is correct. An Einstein Discovery story is the analytical model-building process that examines a dataset to identify patterns, predictions, and improvement recommendations. Salesforce applies a monthly, organization-level allocation to story creations so that Einstein Discovery usage can be governed as a shared org resource. The included allocation is 500 story creations per org per month. This is a count of story creation activity, rather than a limit on the number of stories that can remain available in the org at one time. Organizations with a business need for additional modeling capacity can purchase more story creations, allowing the monthly allocation to be increased through the applicable Salesforce entitlement or add-on. This distinction is important when planning an Einstein Discovery implementation: teams should estimate how frequently analysts will create or recreate stories while iterating on datasets, outcomes, and model configurations, then account for the base monthly allocation and any necessary purchased capacity. Salesforce documents Einstein Discovery capabilities and administration in its [Tableau CRM and Einstein Discovery documentation](#); related learning material is available in the [Einstein Discovery Basics Trailhead module](#).

QUESTION NO: 18

In what order does Einstein Discovery present the insights that it uncovers?

- A. Alphabetical, in ascending order
- B. Alphabetical, in descending order
- C. Insights that explain the most variation in the outcome variable, in ascending order
- D. Insights that explain the most variation in the outcome variable, in descending order

ANSWER: D

Explanation:

Einstein Discovery presents insights in descending order based on how much variation in the outcome variable each insight explains. This prioritization helps users focus first on the factors with the greatest demonstrated impact on the predicted or analyzed outcome, such as customer churn, revenue, or case resolution time. The ordering is driven by the statistical strength and explanatory contribution of the insight rather than by the names of fields or values. In practical terms, the first insights shown are those that account for the largest portion of the observed differences in outcome values across the data, making them the most useful starting point for understanding key drivers and taking action. Einstein Discovery is designed to surface these data-backed drivers automatically, allowing business users to quickly identify the conditions associated with better or worse outcomes and explore the relevant comparisons. Salesforce describes Einstein Discovery as analyzing data to identify meaningful patterns and explain outcomes; its insight presentation emphasizes the most influential explanations first. See Salesforce's [Einstein Discovery documentation](#) and the [Einstein Discovery Basics Trailhead module](#).

QUESTION NO: 19

The Tableau CRM team at a company creates three recipes.

- 1) myRecipeOne: this recipe takes 2 hours to run.
- 2) mvRedpeTwo: this recipe takes 1 hour and 30 minutes to run.
- 3) myRecipeThree: this recipe takes 1 minute and 30 seconds to run.

If all three recipes run, how many count towards the 24 hour rolling recipe run limit.

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

ANSWER: A**Explanation:**

2 is correct because Tableau CRM recipe usage is subject to a rolling 24-hour recipe runtime limit, but recipe runs that complete in less than two minutes are excluded from that limit. myRecipeOne runs for two hours and mvRedpeTwo runs for one hour and 30 minutes, so each is a qualifying recipe run and counts toward the organization's rolling limit. myRecipeThree runs for one minute and 30 seconds, which is below the two-minute threshold. Consequently, it does not count toward the rolling recipe run limit, even though the recipe did execute successfully.

This distinction is important when planning scheduled data preparation. The limit concerns the aggregate runtime of qualifying recipe executions within the preceding 24-hour window, rather than simply counting every recipe invocation. Administrators can use this behavior to understand why a short recipe execution does not consume runtime allowance, while longer recipe runs do. Salesforce documents the runtime threshold and rolling-limit behavior in its recipe limits guidance: [CRM Analytics Recipe Limits](#). For additional context on creating and running recipes in CRM Analytics, see [Prepare Data for CRM Analytics](#).

QUESTION NO: 20

A consultant is reviewing an Einstein Discovery story before deployment. The business requires a prediction to be shown for every new record, even when some predictor values are unavailable.

Which two actions should the consultant take? Choose 2 answers

- A. Review the story's missing-value handling.
- B. Test predictions using records that contain missing predictor values.
- C. Exclude every predictor that can ever contain a null value.
- D. Use the target field as a replacement value for missing predictors.

ANSWER: A B**Explanation:**

Review the story's missing-value handling is required because prediction records can contain null values even when the training data was prepared carefully. The consultant must understand how the story handles missing predictor values and verify that the resulting prediction behavior supports the business requirement.

Test predictions using records that contain missing predictor values is also appropriate. Testing representative incomplete records confirms that the deployed model can return predictions under realistic operating conditions and helps the consultant identify fields whose missingness creates unacceptable prediction quality or coverage.

Removing all records with missing values from the training dataset does not ensure that future records will contain every needed value. Salesforce recommends reviewing data quality and validating story behavior before deployment. See the [Einstein Discovery Basics](#) Trailhead module and Salesforce Help for [preparing data for Einstein Discovery](#).

QUESTION NO: 21

What happens if you first disable Analytics, and then you re-enable Analytics later? Select 2

- A. User permissions are removed from each defined permission set if Analytics is disabled. (Missed)
- B. User permissions are not removed from each defined permission set if Analytics is disabled.
- C. You must define the permission sets again if Analytics is re-enabled. (Missed)
- D. You must not define the permission sets again if Analytics is re-enabled.

ANSWER: A C

Explanation:

When Analytics is disabled, the Analytics-specific user permissions that had been included in defined permission sets are removed. Disabling the feature therefore does more than merely prevent users from opening Analytics; it removes the feature permissions that grant access to Analytics capabilities. This is important for administrators planning a temporary shutdown, because permission-set access does not automatically resume simply by turning the feature back on.

After Analytics is re-enabled, the administrator must define the required permission sets again and assign the appropriate Analytics permissions before users can access and administer the feature as intended. This approach ensures that access is deliberately re-established after the feature is restored, consistent with Salesforce's permission-based security model. Administrators should review the required Analytics permission sets, recreate or update them with the relevant Analytics permissions, and then assign them to the appropriate users. Salesforce's setup guidance for enabling Analytics and managing access is available in [Salesforce Help: Enable CRM Analytics](#). For an overview of CRM Analytics access and setup concepts, see [Trailhead: Get Started with CRM Analytics](#).

QUESTION NO: 22

As part of their analysis, Yasmine, Tonya, and Hadiqa need to create and save lenses to their project app. What is the minimum app role they need to do this?

- A. Manager
- B. Viewer
- C. Editor
- D. Designer

ANSWER: C

Explanation:

Editor is the minimum CRM Analytics app role that permits users to create and save lenses in an app. A lens is an interactive exploration of a dataset; after configuring charts, filters, groupings, and other analysis elements, the user must be able to save that work to make it available in the app. The Editor role provides the required ability to create and modify app content, including lenses, while also allowing collaborators to continue working with the shared analytics assets. This makes it the appropriate least-privilege role for Yasmine, Tonya, and Hadiqa: they can perform and preserve their analysis without being granted the broader administrative and content-management capabilities associated with managing the entire app. Salesforce documents CRM Analytics app roles and the access they grant in its app-sharing guidance, including the distinction between viewing content and editing or saving app assets. See [Share CRM Analytics Apps](#) and the [CRM Analytics Lenses overview](#) for details on lenses and app collaboration permissions.

QUESTION NO: 23

A customer is reviewing a story that is set to maximize the daily sales quantity of consumer products in stores and sees this chart. The visualized tooltip belongs to the blue bar for San Francisco, reflecting, November daily sales quantities in that city specifically.

What two conclusions can be drawn from this insight?

- A. The average daily in SAN Francisco stores in November as 1601 items higher than the global average of 335.
- B. November sales are higher than in other months. This November-effect is the strongest in San Francisco.
- C. The average daily quantity in San Francisco stores in November was 1239 items higher than the average of all other months in San Francisco.
- D. The average daily quantity in San Francisco stores in November was 1239 items higher than the average of all November sales in the country

ANSWER: A C**Explanation:**

The tooltip supports two distinct comparisons for San Francisco's November daily sales quantity. First, it reports that the average daily quantity for San Francisco stores in November is 1,601 items higher than the global average of 335 items. This is a comparison between the selected city-and-month segment and the overall benchmark represented in the insight. It establishes that the observed November result for San Francisco is substantially above the global average daily quantity.

Second, the insight compares November with San Francisco's own performance during all other months. It states that San Francisco's average daily quantity in November was 1,239 items higher than its average for the other months. This identifies a within-city seasonal difference: the November value is not merely high relative to a broad benchmark, but is also elevated relative to San Francisco's usual monthly pattern.

These conclusions follow the scope of the tooltip's displayed measures and comparison baselines. In Tableau CRM and Einstein Discovery, chart tooltips and generated insights can communicate segment-level measures alongside relevant comparison values, helping users interpret both overall and contextual performance. See Salesforce's [Tableau CRM story design Trailhead module](#) and [Einstein Discovery documentation](#).

QUESTION NO: 24

What does a configuration wizard do?

- A. Automatically creates multiple versions of the same app
- B. Lets you know the best way to set up your Salesforce data
- C. Asks you questions about how you store data in your org to help you create your app
- D. Customizes your template based on Salesforce settings

ANSWER: C**Explanation:**

A configuration wizard asks guided questions about how information is stored in the Salesforce org and uses the answers to create a configured analytics app from a template. During template installation, the wizard collects details such as the Salesforce objects, fields, relationships, and business-specific mappings that the app needs. This guided setup allows the template to adapt to the organization's existing data model without requiring the administrator to manually build every dataset, lens, dashboard, and binding from scratch.

In Tableau CRM, app templates are designed to be reusable starting points rather than one-size-fits-all implementations. The configuration wizard is the mechanism that turns that starting point into an app relevant to a particular org: it prompts the installer for the data choices required by the template, then applies those selections while creating the app. Salesforce describes this template-installation experience and the role of configuration in its Trailhead material on analytics templates. See [Use Analytics Templates](#) and [Salesforce Help: Analytics App Templates](#).

QUESTION NO: 25

The client is trying to create a SAQL step to predict sales in each sales region. They cannot get the query to return any results, but have identified that the error is in the time series statement. They have asked an Einstein Consultant to review the following query and fix any errors.

Which timeseries statement will fix the query'

- A. `q = timeseries q generate 'sum_Sales' as Forecasted_Sales' with (dateCols=('Date_Year, Date_Quarter', "Y-Q"), partition='Region', ignoreLast=true);`
- B. `q = timeseries q generate 'sum_Sales' as Forecasted_Sales' with (length = 12, dateCols=('Date_Year', 'Date_Month'. "Y-M"), partition='Region');`

C. `q = timeseries q generate 'sum_Sales' as 'Forecasted_Sales' with (dateCols=('Date_Year', 'Date_Quarter', 'Y-Q'), partition='Region', seasonality=4);`

D. `q = timeseries q generate 'sum_Sales' as 'Forecasted_Sales' with (dateCols=('Date_Year', 'Date_Quarter', 'Y-Q'), partition='Region');`

E. `q = timeseries q generate 'sum_Sales' as 'Forecasted_Sales' with (dateCols=('Date_Year', 'Date_Quarter', 'Y-Q'), partition='Region', seasonality=4);`

ANSWER: E

Explanation:

`q = timeseries q generate 'sum_Sales' as 'Forecasted_Sales' with (dateCols=('Date_Year', 'Date_Quarter', 'Y-Q'), partition='Region', seasonality=4);` correctly applies the SAQL `timeseries` function to produce a sales forecast for each region. The generated measure and output alias are both enclosed in single quotes, as required for SAQL identifiers. Its `dateCols` setting supplies the year field, quarter field, and the "Y-Q" format needed to construct a chronological quarterly time series. Using `partition='Region'` instructs Tableau CRM to create an independent forecast for every sales region rather than combining all regional sales into a single forecast. Finally, `seasonality=4` is appropriate for quarterly data because there are four observations in a yearly seasonal cycle. These settings create the date grain, grouping, and seasonal configuration that the forecasting operation needs. Salesforce documents the `timeseries` SAQL statement and its forecasting parameters in the [SAQL Timeseries Statement reference](#). Additional SAQL syntax guidance is available in the [Salesforce SAQL Developer Guide](#).

QUESTION NO: 26

An Einstein Analytics team wants to create a dashboard using two standard Salesforce objects. The dashboard should display data from the Case object, along with related data from the Account object. Which two processes can achieve this?

A. Extract data from the Case and Account objects using "sfcdigest", use "append" transformation to add Case data with Account data, store as a dataset, and then build a dashboard using the generated dataset.

B. Extract data from the Case and Account objects using an ETL tool, use the ETL tool to join Account and Case data, and then upload the data to Einstein Analytics. Build a dashboard using the ETL-generated dataset.

C. Extract data from the Case and Account objects using "sfcdigest", use "augment" transformation to add Account data to Case data, store as a dataset, and then build a dashboard using the generated dataset.

D. Extract data from the Case and Account objects using "dataloader", create a CSV file with Account and Case data, and then upload the CSV as a dataset to Einstein Analytics. Build a dashboard using that dataset.

ANSWER: B C

Explanation:

Extracting the Case and Account objects with an external ETL tool, joining the records in that tool, and uploading the resulting dataset can produce a dashboard-ready denormalized dataset. The ETL join should associate each Case with its related Account, ordinarily through the `Case.AccountId` relationship, so Account attributes are available alongside Case measures and dimensions in the uploaded dataset. CRM Analytics can then use that uploaded dataset as the dashboard's source.

Extracting both objects with `sfcdigest` and using `augment` is the native CRM Analytics dataflow approach. An `augment` transformation enriches a primary stream with fields from a related lookup stream using matching key fields. With Cases as the primary stream and Accounts as the lookup stream, the resulting output contains Case rows enriched with the appropriate Account data. Registering that output as a dataset makes the combined fields available for dashboard lenses and widgets. Salesforce's data-preparation guidance describes using dataflows to load Salesforce data and joining or augmenting related data to create analysis-ready datasets. See [Salesforce Trailhead: Prepare Data for CRM Analytics](#) and [Salesforce Help: Integrate Data with CRM Analytics](#).

QUESTION NO: 27

What are the 3 basic actions within a query in Analytics? Choose 3:

- A. Aggregate by a measure
- B. Count rows by measure you want to visualize.
- C. Filter to narrow down your results.
- D. Grouping by a dimension
- E. Display the results graphically

ANSWER: A C D

Explanation:

The core actions used to shape a CRM Analytics query are filtering the data, grouping records, and aggregating numeric values. "Filter to narrow down your results." is correct because a filter limits the records that proceed through the query, such as retaining only records within a particular date range, region, or status. "Grouping by a dimension" is correct because grouping organizes the retained records into categories based on a dimension, such as account, product, or month. This establishes the level of detail at which the result is returned. "Aggregate by a measure" is correct because aggregation calculates a summary value for each group, such as total revenue, average amount, minimum, maximum, or a record count. Together, these actions enable an analytics query to select the relevant population, organize it meaningfully, and produce the summarized values required for analysis. CRM Analytics uses these concepts in both its visual query tools and SAQL, where filtering, grouping, and aggregate functions are central to creating query results. See Salesforce Trailhead's [Introduction to SAQL](#) and [SAQL Aggregate Functions](#) for related query and aggregation guidance.

QUESTION NO: 28

The Vice President of Sales wants to create a "Pipeline Trending" dashboard in Einstein Analytics that displays specific colors depending on the values of the fields.

Which two techniques can a consultant use to help address this requirement? Choose 2 answers

- A. Use an XMD file and assign colors to each field value.
- B. Use SAQL queries to calculate a color field and assign it to each field value.
- C. Use a compute Expression to calculate a color field and assign it to each field value.
- D. Use the "fields" menu when exploring the dataset and assign colors to each field value.

ANSWER: A D

Explanation:

Use an XMD file and assign colors to each field value. is correct because an extended metadata (XMD) file provides dataset-level presentation metadata, including color mappings for dimension values. This lets a consultant define consistent colors for pipeline stages or other categorical field values so that the same values are presented consistently across lenses and dashboards that use the dataset. It is especially useful when the color convention must be reusable and governed centrally rather than configured repeatedly in individual visualizations.

Use the "fields" menu when exploring the dataset and assign colors to each field value. is also correct because the Explorer interface allows a user to configure colors for values used in a visualization. This is a practical interactive approach when creating or refining a lens that will be added to the Pipeline Trending dashboard. The configured colors make significant business states, such as pipeline stages or trend categories, immediately recognizable to sales users. Salesforce documents XMD as the mechanism for extending dataset metadata and controlling display-related behavior in its [Extended Metadata \(XMD\) Developer Guide](#). For dashboard and lens authoring concepts, see Salesforce's [CRM Analytics Explorer documentation](#).

QUESTION NO: 29

A consultant is asked to configure a user to view a dashboard, within a designated app, without the ability to create and save lenses.

After adding the user to the app, which action should the consultant take to implement this requirement?

- A. Assign the Manager role to the user
- B. Ensure the user has the Analytics permission
- C. Assign the Editor role to the user
- D. Assign the Viewer role to the user

ANSWER: D

Explanation:

Assign the Viewer role to the user. In CRM Analytics, app-level roles control what a person can do with the content in a specific app. The Viewer role is intended for users who need to consume analytics content: they can open and interact with dashboards and view the app's available assets, while content-authoring capabilities remain unavailable. This meets the requirement to let the user access a dashboard in the designated app without allowing the user to create or save lenses.

CRM Analytics permissions and app membership work together. The required Analytics access lets a user enter CRM Analytics, but the role granted for a particular app determines the scope of that user's actions in that app. Assigning Viewer during app sharing applies the least-privilege approach: the person receives only dashboard-consumption access required for the business need, rather than permissions for modifying or creating analytics content. Salesforce's CRM Analytics learning material describes how app sharing and roles are used to control access to analytics assets; see [Share CRM Analytics Apps](#) and [CRM Analytics Apps](#).

QUESTION NO: 30

What are the two types of bindings? Choose 2:

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

ANSWER: A D

Explanation:

CRM Analytics dashboard bindings have two supported forms: **result bindings** and **selection bindings**. Result bindings use the output of a source step, such as a query or chart, as input for another step. This enables a dashboard to dynamically pass computed values, aggregates, or query results into another widget's query, filter, or display configuration. Selection bindings use the values a viewer actively selects in a source widget, such as choosing a bar in a chart or a row in a table. The selected value is then passed to a target step so that the dashboard responds interactively, commonly by filtering related charts, tables, or metrics. Together, these binding types make dashboards responsive without requiring separate static dashboards for every possible view of the data. Salesforce's CRM Analytics developer documentation identifies result and selection bindings as the available dashboard binding types and describes their use in connecting source and target steps. See the [Salesforce CRM Analytics Bindings Developer Guide](#) and the [CRM Analytics Dashboard Bindings Trailhead module](#).

QUESTION NO: 31

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 Analytics External Data API operation because each monthly CSV is a complete replacement snapshot: it includes targets for every sales representative for every month in the current year and all prior years. Replacing the dataset with that authoritative full file ensures that the target-attainment calculations use one consistent, current set of records. It also prevents duplicate historical rows from accumulating when the monthly process runs.

In the Analytics External Data API, the upload job's operation determines how the uploaded data is applied to the target dataset. The **Overwrite** operation replaces the dataset's existing data with the uploaded data, which directly matches a recurring full-refresh integration pattern. This is especially suitable where historical values may be corrected in the source file as part of a later monthly delivery, since the newly supplied file becomes the complete dataset state. Salesforce documents the available external-data upload operations and their behavior in the [CRM Analytics REST API external data resource documentation](#). For the broader CSV-based dataset upload workflow, see Salesforce's [dataset upload documentation](#).

QUESTION NO: 32

The Universal Containers company used Einstein Analytics to create two datasets:

Dataset A: contains a list of activities with an "activityID" dimension and a "userID" dimension
Dataset B: contains a list of users with a "userID" dimension

The team wants to delete from Dataset A all activities related to users in Dataset B.

How can an Einstein Consultant help them achieve this?

- A. Use the dataflow transformation "delete" and set "userID" as the deletion ID.
- B. Use a combination of dataflow transformations: "augment" and "filter."
- C. Use an external ETL tool to extract both datasets and delete records.
- D. Use the recipe operation "delete" and set "userID" as the deletion ID.

ANSWER: B

Explanation:

Use a combination of dataflow transformations: "augment" and "filter." This is the appropriate approach because the desired result is an anti-join: retain activity records only when their associated user does not appear in the user dataset. First, configure an augment transformation with the activities as the primary stream and join it to the users stream using the shared `userID` field. The augment adds a field from the matching user record to each activity record while preserving the activity stream for subsequent processing. Next, apply a filter transformation that keeps only records for which the augmented user field is null or otherwise indicates that no matching user was found. The resulting output excludes all activities belonging to users listed in the user dataset, without requiring source-level record deletion. This pattern is useful when CRM Analytics data preparation needs to compare datasets and produce a curated output dataset for analysis. Salesforce documents

augment as the dataflow operation for joining streams and enriching records: [Salesforce Developer Guide: Augment](#). Additional dataflow transformation guidance is available in the [Salesforce Help: CRM Analytics Dataflows](#).

QUESTION NO: 33

Which three things can be done with the tableauCRM Dashboard Inspector?

Choose 3 answers

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

ANSWER: A B D

Explanation:

Dashboard Inspector is a diagnostic and performance-analysis tool for Tableau CRM dashboards. “Get a list of recommendations on how to improve the performance of the dashboard” is correct because the inspector analyzes dashboard design and query behavior, then presents guidance that can help an author reduce dashboard load time and improve efficiency. “View the total time required to run all queries” is also correct because the inspector provides an overall query-duration measure, helping consultants assess the query-processing portion of a dashboard’s performance. “View allqueries and the time it took to run each one” is correct because its query-level details identify the individual queries executed by the dashboard and their elapsed times. This information lets an author focus optimization work on the queries that contribute most to dashboard latency. Dashboard Inspector therefore supports evidence-based troubleshooting: it combines overall timing, individual query timings, and actionable recommendations rather than making performance changes automatically. Salesforce’s dashboard-performance guidance and Tableau CRM learning materials describe using dashboard inspection and query analysis to diagnose slow dashboards and target improvements. See [Salesforce Help: Dashboard Inspector](#) and [Trailhead: Optimize Dashboard Performance](#).

QUESTION NO: 34

Upper and lower limits on rows in discovery.

- A. 1000 row minimum, 15,000,000 max
- B. 3000 row minimum, 20,000,000 max
- C. 00 row minimum, 20,000,000 max
- D. 500 row minimum, 25,000,000 max

ANSWER: B

Explanation:

3000 row minimum, 20,000,000 max is correct for an Einstein Discovery analysis. Einstein Discovery requires a sufficiently large data population to build and evaluate predictive models meaningfully; 3,000 rows is the documented minimum number of rows for creating a discovery. This threshold helps ensure that the service has enough records for training, validation, and statistically useful insights. At the upper end, a discovery can use up to 20 million rows. When preparing data for discovery, consultants should therefore verify the row count after recipes, joins, filters, and dataflow transformations, rather than relying only on the source-system record total. This is especially important where date filters or security predicates substantially reduce the usable dataset. The row range applies to the data supplied to the discovery and should be considered alongside other Einstein Discovery data-preparation and field requirements. Salesforce documents these discovery-specific capacity

boundaries in its Einstein Discovery limits guidance; see [Einstein Discovery Limits](#) and the [Salesforce Help guidance for creating discoveries](#).