Number: Data Architect
Passing Score: 800
Time Limit: 120
File Version: 5.0

Exam Code: Data Architect
Exam Name: Salesforce Certified Data Architect

Exam A

QUESTION 1
A manager at Cloud Kicks is importing Leads into Salesforce and needs to avoid creating duplicate records.
Which two approaches should the manager take to achieve this goal? (Choose two.)

A. Acquire an AppExchange Lead de-duplication application.
B. Implement Salesforce Matching and Duplicate Rules.
C. Run the Salesforce Lead Mass de-duplication tool.
D. Create a Workflow Rule to check for duplicate records.

Correct Answer: A, B
Section:
Explanation:
Acquiring an AppExchange Lead de-duplication application and implementing Salesforce Matching and Duplicate Rules are two approaches that the manager at Cloud Kicks should take to avoid creating duplicate records when importing Leads into Salesforce. An AppExchange Lead de-duplication application can provide additional features and functionality for finding and preventing duplicate Leads during import, such as fuzzy matching, custom rules, mass merge, etc. Salesforce Matching and Duplicate Rules can allow the manager to define how Salesforce identifies duplicate Leads based on various criteria and how users can handle them during import, such as blocking, allowing, or alerting them. The other options are not feasible or effective for avoiding duplicate records, as they would either not work during import, not provide de-duplication capabilities, or require additional customization.

QUESTION 2
DreamHouse Realty has an integration that creates records in a Salesforce Custom Object. The Custom Object has a field marked as required on the page layout.
DreamHouse Realty has noticed that many of the records coming from the external system are missing data in this field.
The Architect needs to ensure this field always contains data coming from the source system.
Which two approaches should the Architect take? Choose 2 answers

A. Set up a Validation Rule to prevent blank values.
B. Create a Workflow to default a value into this field.
C. Mark the field required in setup at the field level.
D. Blame the customer's external system for bad data.

Correct Answer: A, C
Section:
Explanation:
Setting up a Validation Rule to prevent blank values and marking the field required in setup at the field level are two approaches that the Architect should take to ensure that the field always contains data coming from the source system. A Validation Rule can display an error message when a record is created or edited with a blank value for the field, and prevent it from being saved. Marking the field required in setup at the field level can enforce the requirement for all records, regardless of the page layout or the source system.The other options are not effective or recommended for ensuring data quality, as they would either not prevent blank values, not apply to all records, or not address the root cause of the problem

QUESTION 3
Universal Containers has two systems. Salesforce and an on -premise ERP system. An architect has been tasked with copying Opportunity records to the ERP once they reach a Closed/Won Stage. The Opportunity record in the ERP system will be read-only for all fields copied in from Salesforce. What is the optimal real-time approach that achieves this solution?

A. Implement a Master Data Management system to determine system of record.
B. Implement a workflow rule that sends Opportunity data through Outbound Messaging.

C. Have the ERP poll Salesforce nightly and bring in the desired Opportunities.
D. Implement an hourly integration to send Salesforce Opportunities to the ERP system.

Correct Answer: B
Section:
Explanation:
Implementing a workflow rule that sends Opportunity data through Outbound Messaging is the optimal real-time approach that achieves the solution of copying Opportunity records to the ERP once they reach a Closed/Won Stage. A workflow rule can trigger an Outbound Message when an Opportunity record meets certain criteria, such as having a Closed/Won Stage. An Outbound Message can send data from Salesforce to an external system via SOAP API, without requiring any code. The external system can then process the data and create a read-only record in the ERP system.The other options are not optimal or real-time, as they would either require additional systems or tools, not provide real-time data synchronization, or not meet the frequency requirement

QUESTION 4
Universal Containers (UC) has three systems: Salesforce, a cloud -based ERP system, and an on -premise Order Management System (OMS). An architect has been tasked with creating a solution that uses Salesforce as the system of record for Leads and the OMS as the system of record for Account and Contacts. UC wants Accounts and Contacts to be able to maintain their names in each system (i.e., 'John Doe' in the OMS and 'Johnny Doe' in Salesforce), but wants to have a consolidated data store which links referenced records across the systems. What approach should an architect suggest so the requirements are met?

A. Have Salesforce poll the OMS nightly and bring in the desired Accounts and Contacts.
B. Implement an integration tool to send OMS Accounts and Contacts to Salesforce.
C. Implement a Master Data Management strategy to reconcile Leads, Accounts, and Contacts.
D. Use the Streaming API to send Account and Contact data from Salesforce to the OMS.

Correct Answer: C
Section:
Explanation:
Implementing a Master Data Management strategy to reconcile Leads, Accounts, and Contacts is the approach that the architect should suggest so that the requirements are met for UC. A Master Data Management strategy can provide a consistent and unified view of data across multiple systems, by defining and enforcing rules for data quality, governance, and integration. A Master Data Management strategy can also handle complex scenarios where different systems have different systems of record for different entities or fields, and where data can be modified in different systems with different values. A Master Data Management strategy can also provide a consolidated data store that links referenced records across systems and enables cross-system reporting and analysis.The other options are not suitable or sufficient for meeting the requirements, as they would either not provide a consolidated data store, not handle different systems of record, or not allow data modification in different system

QUESTION 5
An architect has been asked by a client to develop a solution that will integrate data and resolve duplicates and discrepancies between Salesforce and one or more external systems. What two factors should the architect take into consideration when deciding whether or not to use a Master Data Management system to achieve this solution?
Choose 2 answers

A. Whether the systems are cloud -based or on -premise.
B. Whether or not Salesforce replaced a legacy CRM.
C. Whether the system of record changes for different tables.
D. The number of systems that are integrating with each other.

Correct Answer: C, D
Section:
Explanation:
Whether the system of record changes for different tables and the number of systems that are integrating with each other are two factors that the architect should take into consideration when deciding whether or not to use a Master Data Management system to achieve the solution of integrating data and resolving duplicates and discrepancies between Salesforce and one or more external systems. The system of record is the authoritative source of truth for a given entity or field in a given context. If different systems have different systems of record for different tables, then a Master Data Management system can help to manage and synchronize the data across systems and ensure data quality and consistency. The number of systems that are integrating with each other is another factor that affects the complexity and scalability of the integration solution. If there are many systems that need to integrate with each other, then a Master Data Management system can provide a centralized and standardized way of integrating data and resolving duplicates and discrepancies across systems. The

other factors are not relevant or important for deciding whether or not to use a Master Data Management system, as they do not affect the data quality or integration challenges that a Master Data Management system can address.

QUESTION 6
Get Cloud Consulting needs to integrate two different systems with customer records into the Salesforce Account object. So that no duplicate records are created in Salesforce, Master Data Management will be used.
An Architect needs to determine which system is the system of record on a field level.
What should the Architect do to achieve this goal?

A. Master Data Management systems determine system of record, and the Architect doesn't have to think about what data is controlled by what system.
B. Key stakeholders should review any fields that share the same purpose between systems to see how they will be used in Salesforce.
C. The database schema for each external system should be reviewed, and fields with different names should always be separate fields in Salesforce.
D. Any field that is an input field in either external system will be overwritten by the last record integrated and can never have a system of record.

Correct Answer: B
Section:
Explanation:
Key stakeholders from both systems should collaborate with the Architect to determine which system is the system of record on a field level is what the Architect should do to achieve this goal of integrating two different systems with customer records into the Salesforce Account object using Master Data Management. The system of record is the authoritative source of truth for a given entity or field in a given context. Different systems may have different levels of accuracy, completeness, timeliness, or relevance for different fields. Therefore, it is important to involve key stakeholders from both systems who have knowledge and expertise about their data quality and business needs to decide which system should be the system of record for each field. The Architect should facilitate this collaboration and document the decisions and rationale for each field. The other options are not correct or feasible, as they would either delegate or abdicate the responsibility of determining the system of record, ignore or disregard the input from key stakeholders, or assume or impose a default system of record without considering the data quality and business needs.

QUESTION 7
Salesforce is being deployed in Ursa Major Solar's disparate, multi-system ERP environment. Ursa major Solar wants to maintain data synchronization between systems.
Which two techniques should be used to achieve this goal? (Choose two.)

A. Integrate Salesforce with the ERP environment.
B. Utilize workbench to update files within systems.
C. Utilize an MDM strategy to outline a single source of truth.
D. Build synchronization reports and dashboards.

Correct Answer: A, C
Section:
Explanation:
Option A is correct because integrating Salesforce with the ERP environment is a technique to maintain data synchronization between systems1.Option C is correct because utilizing an MDM strategy to outline a single source of truth is another technique to ensure data quality and consistency across systems2.Option B is not correct because utilizing workbench to update files within systems is not a technique to maintain data synchronization, but a tool to perform data manipulation tasks3.Option D is not correct because building synchronization reports and dashboards is not a technique to maintain data synchronization, but a way to monitor and analyze data4.

QUESTION 8
All accounts and opportunities are created in Salesforce. Salesforce is integrated with three systems:
* An ERP system feeds order data into Salesforce and updates both Account and Opportunity records.
* An accounting system feeds invoice data into Salesforce and updates both Account and Opportunity records.
* A commission system feeds commission data into Salesforce and updates both Account and Opportunity records.
How should the architect determine which of these systems is the system of record?

A. Account and opportunity data originates in Salesforce, and therefore Salesforce is the system of record.
B. Whatever system updates the attribute or object should be the system of record for that field or object.

C. Whatever integration data flow runs last will, by default, determine which system is the system of record.
D. Data flows should be reviewed with the business users to determine the system of record per object or field.

Correct Answer: D
Section:
Explanation:
Option D is correct because data flows should be reviewed with the business users to determine the system of record per object or field1.The system of record is the authoritative source of data for a given entity or attribute2. It may vary depending on the business context and requirements.Option A is not correct because account and opportunity data originates in Salesforce, but it may not be the system of record for all attributes or objects if they are updated by other systems2.Option B is not correct because whatever system updates the attribute or object may not be the system of record for that field or object if there are conflicting or overlapping updates from other systems2.Option C is not correct because whatever integration data flow runs last may not determine which system is the system of record if there are different business rules or logic applied by different systems2.

QUESTION 9
A company wants to document the data architecture of a Salesforce organization.
What are two valid metadata types that should be included? (Choose two.)

A. RecordType
B. Document
C. CustomField
D. SecuritySettings

Correct Answer: A, C
Section:
Explanation:
Option A is correct because RecordType is a valid metadata type that should be included in documenting the data architecture of a Salesforce organization1.RecordType defines different business processes, picklist values, and page layouts for different users2.Option C is correct because CustomField is another valid metadata type that should be included in documenting the data architecture of a Salesforce organization1.CustomField defines custom attributes for standard or custom objects3.Option B is not correct because Document is not a valid metadata type, but a standard object that stores documents in folders4. Option D is not correct because SecuritySettings is not a valid metadata type, but a setup menu that allows administrators to configure various security features such as password policies, network access, session settings, etc.

QUESTION 10
An Architect needs to document the data architecture for a multi-system, enterprise Salesforce implementation.
Which two key artifacts should the Architect use? (Choose two.)

A. User stories
B. Data model
C. Integration specification
D. Non-functional requirements

Correct Answer: B, C
Section:
Explanation:
Option B is correct because data model is a key artifact that an architect should use to document the data architecture for a multi-system, enterprise Salesforce implementation1.Data model describes the structure and relationship of data entities within an organization2.Option C is correct because integration specification is another key artifact that an architect should use to document the data architecture for a multi-system, enterprise Salesforce implementation1.Integration specification defines the scope, requirements, design, testing, and deployment of integration solutions between Salesforce and other systems3.Option A is not correct because user stories are not key artifacts for documenting the data architecture, but agile development tools that capture the features and functionalities that users want from a system4. Option D is not correct because non-functional requirements are not key artifacts for documenting the data architecture, but quality attributes that specify how well a system performs its functions.

QUESTION 11
As part of a phased Salesforce rollout. there will be 3 deployments spread out over the year. The requirements have been carefully documented. Which two methods should an architect use to trace back configuration

changes to the detailed requirements? Choose 2 answers

A. Review the setup audit trail for configuration changes.
B. Put the business purpose in the Description of each field.
C. Maintain a data dictionary with the justification for each field.
D. Use the Force.com IDE to save the metadata files in source control.

Correct Answer: B, D
Section:
Explanation:
Option B is correct because putting the business purpose in the Description of each field is a method that an architect can use to trace back configuration changes to the detailed requirements1.The Description of each field provides a brief explanation of what the field is used for and why it is needed2.Option D is correct because using the Force.com IDE to save the metadata files in source control is another method that an architect can use to trace back configuration changes to the detailed requirements1.The Force.com IDE is an integrated development environment that allows developers to work with Salesforce metadata files and Apex code3.Source control is a system that tracks and manages changes to code and configuration files4. Option A is not correct because reviewing the setup audit trail for configuration changes is not a method to trace back configuration changes to the detailed requirements, but a way to monitor and audit the changes made in the setup area. Option C is not correct because maintaining a data dictionary with the justification for each field is not a method to trace back configuration changes to the detailed requirements, but a document that provides information about the data entities and attributes in a system.

QUESTION 12
How can an architect find information about who is creating, changing, or deleting certain fields within the past two months?

A. Remove 'customize application' permissions from everyone else.
B. Export the metadata and search it for the fields in question.
C. Create a field history report for the fields in question.
D. Export the setup audit trail and find the fields in question.

Correct Answer: D
Section:
Explanation:
The setup audit trail tracks changes made in your org's setup area for the past 180 days1.You can export the setup audit trail and find the fields in question by filtering by action type, user, or date2.

QUESTION 13
Universal Containers wants to automatically archive all inactive Account data that is older than 3 years. The information does not need to remain accessible within the application. Which two methods should be recommended to meet this requirement? Choose 2 answers

A. Use the Force.com Workbench to export the data.
B. Schedule a weekly export file from the Salesforce UI.
C. Schedule jobs to export and delete using an ETL tool.
D. Schedule jobs to export and delete using the Data Loader.

Correct Answer: C, D
Section:
Explanation:
Both C and D are valid methods to automatically archive and delete inactive Account data that is older than 3 years1.You can use an ETL tool or the Data Loader to schedule jobs to export and delete data based on certain criteria3. Option A is not recommended because the Force.com Workbench is a web-based tool that does not support scheduling or automation. Option B is not suitable because the weekly export file from the Salesforce UI does not delete data from Salesforce.

QUESTION 14
Cloud Kicks needs to purge detailed transactional records from Salesforce. The data should be aggregated at a summary level and available in Salesforce.

What are two automated approaches to fulfill this goal? (Choose two.)

A. Third-party Integration Tool (ETL)
B. Schedulable Batch Apex
C. Third-party Business Intelligence system
D. Apex Triggers

Correct Answer: A, B
Section:
Explanation:
Both A and B are automated approaches to purge detailed transactional records from Salesforce and aggregate them at a summary level1. You can use a third-party integration tool (ETL) or schedulable batch Apex to perform these tasks. Option C is not correct because a third-party business intelligence system does not purge data from Salesforce, but only analyzes it. Option D is not correct because Apex triggers are not automated, but execute when a record is inserted, updated, deleted, or undeleted.

QUESTION 15
Universal Containers (UC) is concerned that data is being corrupted daily either through negligence or maliciousness. They want to implement a backup strategy to help recover any corrupted data or data mistakenly changed or even deleted. What should the data architect consider when designing a field -level audit and recovery plan?

A. Reduce data storage by purging old data.
B. Implement an AppExchange package.
C. Review projected data storage needs.
D. Schedule a weekly export file.

Correct Answer: C
Section:
Explanation:
Option C is the best answer because reviewing projected data storage needs is an important step in designing a field-level audit and recovery plan1. You need to estimate how much data storage you will need in the future and plan accordingly. Option A is not correct because reducing data storage by purging old data may not be sufficient or desirable for backup purposes. Option B is not correct because implementing an AppExchange package may not be customized or compatible with your org's requirements. Option D is not correct because scheduling a weekly export file may not be frequent or granular enough for field-level audit and recovery.

QUESTION 16
Every year, Ursa Major Solar has more than 1 million orders. Each order contains an average of 10 line items. The Chief Executive Officer (CEO) needs the Sales Reps to see how much money each customer generates year-over-year. However, data storage is running low in Salesforce.
Which approach for data archiving is appropriate for this scenario?

A. 1. Annually export and delete order line items. 2. Store them in a zip file in case the data is needed later.
B. 1. Annually aggregate order amount data to store in a custom object. 2. Delete those orders and order line items.
C. 1. Annually export and delete orders and order line items. 2. Store them in a zip file in case the data is needed later.
D. 1. Annually delete orders and order line items. 2. Ensure the customer has order information in another system.

Correct Answer: B
Section:
Explanation:
Option B is the most appropriate approach for data archiving in this scenario1. By aggregating order amount data at a summary level and storing it in a custom object, you can reduce data storage and still provide visibility to the Sales Reps on how much money each customer generates year-over-year. Option A is not correct because storing data in a zip file does not make it available in Salesforce. Option C is not correct because exporting and deleting orders and order line items may lose important details that are needed for analysis or reporting. Option D is not correct because deleting orders and order line items without exporting them may cause data loss or inconsistency if the customer does not have order information in another system.

QUESTION 17
Get Cloudy Consulting monitors 15,000 servers, and these servers automatically record their status every 10 minutes. Because of company policy, these status reports must be maintained for 5 years. Managers at Get Cloudy Consulting need access to up to one week's worth of these status reports with all of their details.
An Architect is recommending what data should be integrated into Salesforce and for how long it should be stored in Salesforce.
Which two limits should the Architect be aware of? (Choose two.)

A. Data storage limits
B. Workflow rule limits
C. API Request limits
D. Webservice callout limits

Correct Answer: A, C
Section:
Explanation:
Data storage limits and API request limits are two important factors that affect the data integration and storage in Salesforce. Data storage limits determine how much data can be stored in Salesforce, and API request limits determine how many API calls can be made to Salesforce in a 24-hour period. Both of these limits depend on the edition and license type of the Salesforce org. Workflow rule limits and webservice callout limits are not directly related to data integration and storage, but rather to business logic and external services.

QUESTION 18
A Salesforce customer has plenty of data storage. Sales Reps are complaining that searches are bringing back old records that aren't relevant any longer. Sales Managers need the data for their historical reporting. What strategy should a data architect use to ensure a better user experience for the Sales Reps?

A. Create a Permission Set to hide old data from Sales Reps.
B. Use Batch Apex to archive old data on a rolling nightly basis.
C. Archive and purge old data from Salesforce on a monthly basis.
D. Set data access to Private to hide old data from Sales Reps.

Correct Answer: C
Section:
Explanation:
Archiving and purging old data from Salesforce on a monthly basis is a good strategy to improve the user experience for the Sales Reps, as it will reduce the clutter and improve the search performance. Creating a permission set or setting data access to private are not effective ways to hide old data from Sales Reps, as they will still consume data storage and affect search results. Using Batch Apex to archive old data on a rolling nightly basis is also not a good option, as it will consume API requests and processing time, and may not comply with the data retention policy.

QUESTION 19
Universal Containers (UC) is implementing a formal, cross -business -unit data governance program As part of the program, UC will implement a team to make decisions on enterprise -wide data governance. Which two roles are appropriate as members of this team? Choose 2 answers

A. Analytics/BI Owners
B. Data Domain Stewards
C. Salesforce Administrators
D. Operational Data Users

Correct Answer: A, B
Section:
Explanation:
Analytics/BI Owners and Data Domain Stewards are appropriate roles as members of a team that makes decisions on enterprise-wide data governance. Analytics/BI Owners are responsible for defining the business requirements and metrics for data analysis and reporting, and Data Domain Stewards are responsible for defining and enforcing the data quality standards and rules for specific data domains. Salesforce Administrators and

Operational Data Users are not suitable roles for this team, as they are more focused on the operational aspects of data management, such as configuration, maintenance, and usage.

QUESTION 20
Universal Containers (UC) has a complex system landscape and is implementing a data governance program for the first time Which two first steps would be appropriate for UC to initiate an assessment of data architecture? Choose 2 answers

A. Engage with IT program managers to assess current velocity of projects in the pipeline.
B. Engage with database administrators to assess current database performance metrics.
C. Engage with executive sponsorship to assess enterprise data strategy and goals.
D. Engage with business units and IT to assess current operational systems and data models.

Correct Answer: C, D
Section:
Explanation:
Engaging with executive sponsorship to assess enterprise data strategy and goals, and engaging with business units and IT to assess current operational systems and data models are two first steps that would be appropriate for UC to initiate an assessment of data architecture. These steps will help to understand the current state of data management, the business needs and expectations, and the gaps and opportunities for improvement. Engaging with IT program managers to assess current velocity of projects in the pipeline, and engaging with database administrators to assess current database performance metrics are not relevant steps for assessing data architecture, as they are more related to project management and technical performance.

QUESTION 21
A data architect has been tasked with optimizing a data stewardship engagement for a Salesforce instance Which three areas of Salesforce should the architect review before proposing any design recommendation? Choose 3 answers

A. Review the metadata xml files for redundant fields to consolidate.
B. Determine if any integration points create records in Salesforce.
C. Run key reports to determine what fields should be required.
D. Export the setup audit trail to review what fields are being used.
E. Review the sharing model to determine impact on duplicate records.

Correct Answer: B, C, E
Section:
Explanation:
Determining if any integration points create records in Salesforce, running key reports to determine what fields should be required, and reviewing the sharing model to determine impact on duplicate records are three areas of Salesforce that the architect should review before proposing any design recommendation. These areas will help to identify the sources and quality of data, the business rules and validations for data entry, and the access and visibility of data across users and roles. Reviewing the metadata xml files for redundant fields to consolidate, and exporting the setup audit trail to review what fields are being used are not necessary steps for optimizing a data stewardship engagement, as they are more related to metadata management and audit tracking.

QUESTION 22
To avoid creating duplicate Contacts, a customer frequently uses Data Loader to upsert Contact records into Salesforce. What common error should the data architect be aware of when using upsert?

A. Errors with duplicate external Id values within the same CSV file.
B. Errors with records being updated and inserted in the same CSV file.
C. Errors when a duplicate Contact name is found cause upsert to fail.
D. Errors with using the wrong external Id will cause the load to fail.

Correct Answer: A
Section:
Explanation:

Data Loader uses external Id fields to match records in the CSV file with records in Salesforce during an upsert operation. If the CSV file contains duplicate external Id values within the same file, Data Loader will throw an error saying ''Duplicate Id Specified'' and will not process those records. Therefore, it is important to ensure that the CSV file does not have any duplicate external Id values before using Data Loader to upsert records.

QUESTION 23
Universal Containers has deployed Salesforce for case management The company is having difficulty understanding what percentage of cases are resolved from the initial call to their support organization. What first step is recommended to implement a reporting solution to measure the support reps case closure rates?

A. Enable field history tracking on the Case object.
B. Create a report on Case analytic snapshots.
C. Install AppExchange packages for available reports.
D. Create Contact and Opportunity Reports and Dashboards.

Correct Answer: A
Section:
Explanation:
Enabling field history tracking on the Case object is the first step to implement a reporting solution to measure the support reps case closure rates. Field history tracking allows you to track changes to certain fields on the Case object, such as Status, Owner, or Priority. You can then create reports based on the field history data to analyze how long it took to close a case, how many times the case owner changed, or how many cases were escalated.

QUESTION 24
The head of sales at Get Cloudy Consulting wants to understand key relevant performance figures and help managers take corrective actions where appropriate.
What is one reporting option Get Cloudy Consulting should consider?

A. Case SLA performance report
B. Sales KPI Dashboard
C. Opportunity analytic snapshot
D. Lead conversion rate report

Correct Answer: B
Section:
Explanation:
A Sales KPI Dashboard is one reporting option that Get Cloudy Consulting should consider to understand key relevant performance figures and help managers take corrective actions where appropriate. A Sales KPI Dashboard can display various metrics that indicate the health and effectiveness of the sales team, such as quota attainment, pipeline value, win rate, average deal size, sales cycle length, and more. A Sales KPI Dashboard can also help identify trends, patterns, and areas for improvement.

QUESTION 25
In their legacy system. Universal Containers has a monthly accounts receivable report that compiles data from Accounts, Contacts, Opportunities, Orders. and Order Line Items. What difficulty will an architect run into when implementing this in Salesforce?

A. Salesforce allows up to four objects in a single report type.
B. Salesforce does not support Orders or Order Line Items.
C. A report cannot contain data from Accounts and Contacts.
D. Custom report types cannot contain Opportunity data.

Correct Answer: A
Section:
Explanation:
Salesforce allows up to four objects in a single report type. This means that if an architect wants to create a report that compiles data from Accounts, Contacts, Opportunities, Orders, and Order Line Items, they will run into a

difficulty because that would require five objects in a single report type. To overcome this limitation, the architect can either create two separate report types and join them in a joined report, or use an external reporting tool that can access data from multiple objects.

QUESTION 26
Universal Containers keeps its Account data in Salesforce and its Invoice data in a third -party ERP system. They have connected the Invoice data through a Salesforce external object. They want data from both Accounts and Invoices visible in one report in one place. What two approaches should an architect suggest for achieving this solution? Choose 2 answers

A. Create a report in an external system combining Salesforce Account data and Invoice data from the ERP.
B. Create a report combining data from the Account standard object and the Invoices external object.
C. Create a Visualforce page combining Salesforce Account data and Invoice external object data.
D. Create a separate Salesforce report for Accounts and Invoices and combine them in a dashboard.

Correct Answer: A, C
Section:
Explanation:
Creating a report in an external system combining Salesforce Account data and Invoice data from the ERP, and creating a Visualforce page combining Salesforce Account data and Invoice external object data are two approaches that an architect can suggest for achieving this solution. Both of these approaches can display data from both Accounts and Invoices in one place, using either an external reporting tool or a custom web page. Creating a report combining data from the Account standard object and the Invoices external object is not possible because Salesforce does not support reporting on external objects. Creating a separate Salesforce report for Accounts and Invoices and combining them in a dashboard is not ideal because it will not show the data in one report, but rather in two separate components.

QUESTION 27
Universal Containers wishes to maintain Lead data from Leads even after they are deleted and cleared from the Recycle Bin. What approach should be implemented to achieve this solution?

A. Use a Lead standard report and filter on the IsDeleted standard field.
B. Use a Converted Lead report to display data on Leads that have been deleted.
C. Query Salesforce with the queryAll API method or using the ALL ROWS SOQL keywords.
D. Send data to a Data Warehouse and mark Leads as deleted in that system.

Correct Answer: C
Section:
Explanation:
According to the exam guide, one of the objectives is to 'describe how to use queryAll() or ALL ROWS keywords to access deleted records in Apex and SOQL'1. This implies that option C is the correct way to access deleted records in Salesforce. Option D is not correct because sending data to a data warehouse does not maintain the data in Salesforce.Options A and B are not correct because they do not apply to deleted records2.

QUESTION 28
Universal Containers (UC) has deployed Salesforce to manage Marketing. Sales, and Support efforts in a multi -system ERP environment After reaching the limits of native reports & dashboards. UC leadership is looking to understand what options can be used to provide more analytical insights. What two approaches should an architect recommend? Choose 2 answers

A. AppExchange Apps
B. Wave Analytics
C. Weekly Snapshots
D. Setup Audit Trails

Correct Answer: A, B
Section:
Explanation:
According to the exam guide, one of the objectives is to 'describe the use cases and considerations for using AppExchange apps and Wave Analytics'1. This implies that options A and B are both valid approaches to provide more analytical insights.Option C is not correct because weekly snapshots are used to track changes over time, not to provide advanced analytics3. Option D is not correct because setup audit trails are used to monitor

changes in the setup menu, not to provide analytical insights.

QUESTION 29
Universal Containers is setting up an external Business Intelligence (BI) system and wants to extract 1,000,000 Contact records. What should be recommended to avoid timeouts during the export process?

A. Use the SOAP API to export data.
B. Utilize the Bulk API to export the data.
C. Use GZIP compression to export the data.
D. Schedule a Batch Apex job to export the data.

Correct Answer: C
Section:
Explanation:
According to the exam guide, one of the objectives is to 'describe the use cases and considerations for using various tools and techniques for data migration (for example, Data Loader, Bulk API)'1. This implies that option B is the correct way to extract large volumes of data from Salesforce. The Bulk API is designed to handle large-scale data operations and avoid timeouts. Option A is not correct because the SOAP API is not optimized for large data sets and may encounter limits. Option C is not correct because GZIP compression does not prevent timeouts, but rather reduces the size of the data transferred. Option D is not correct because Batch Apex is used to process records asynchronously in Salesforce, not to export data to an external system.

QUESTION 30
Universal Containers (UC) is a business that works directly with individual consumers (B2C). They are moving from a current home-grown CRM system to Salesforce. UC has about one million consumer records. What should the architect recommend for optimal use of Salesforce functionality and also to avoid data loading issues?

A. Create a Custom Object Individual Consumer c to load all individual consumers.
B. Load all individual consumers as Account records and avoid using the Contact object.
C. Load one Account record and one Contact record for each individual consumer.
D. Create one Account and load individual consumers as Contacts linked to that one Account.

Correct Answer: D
Section:
Explanation:
According to the exam guide, one of the objectives is to 'describe best practices for implementing a single-org strategy in a B2C scenario'1. This implies that option D is the best practice for loading individual consumers as contacts in Salesforce. This approach avoids creating unnecessary accounts and reduces data duplication. Option C is not correct because it creates one account per contact, which increases data volume and complexity. Options A and B are not correct because they do not leverage the standard contact object, which provides native functionality and integration with other Salesforce features.

QUESTION 31
DreamHouse Realty has a data model as shown in the image. The Project object has a private sharing model, and it has Roll-Up summary fields to calculate the number of resources assigned to the project, total hours for the project, and the number of work items associated to the project.
There will be a large amount of time entry records to be loaded regularly from an external system into Salesforce.

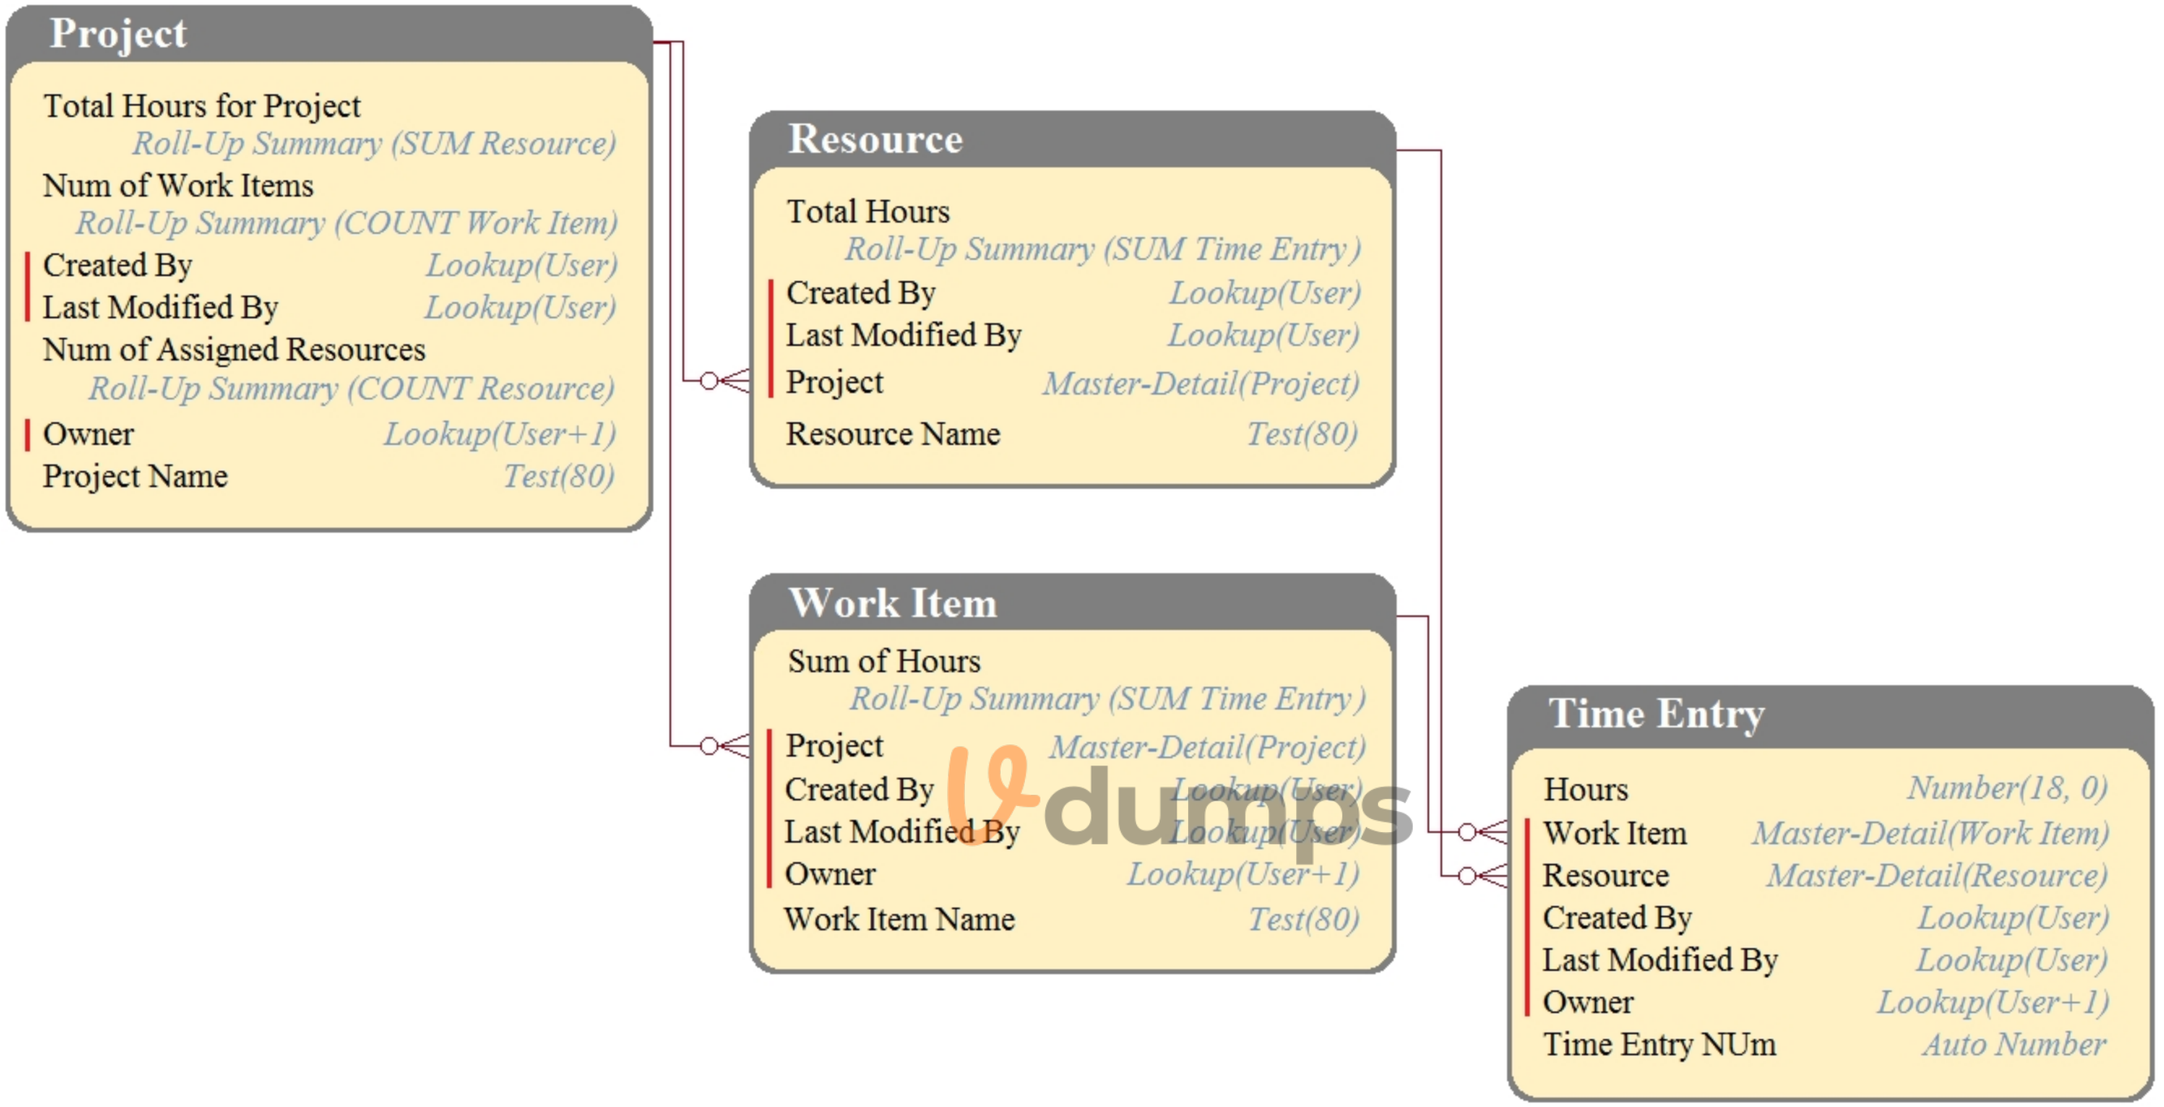

What should the Architect consider in this situation?

A. Load all data after deferring sharing calculations.
B. Calculate summary values instead of Roll-Up by using workflow.
C. Calculate summary values instead of Roll-Up by using triggers.
D. Load all data using external IDs to link to parent records.

Correct Answer: A
Section:
Explanation:
According to the exam guide, one of the objectives is to 'describe the use cases and considerations for deferring sharing calculations'1. This implies that option A is the correct way to load large amounts of data into Salesforce without affecting performance and data integrity.Deferring sharing calculations allows the data to be loaded first and then the sharing rules to be applied later2.Option B is not correct because workflows are not recommended for calculating summary values, as they can cause performance issues and data skew3. Option C is not correct because triggers are also not recommended for calculating summary values, as they can cause governor limit errors and data inconsistency. Option D is not correct because external IDs are used to link records from different systems, not to improve data loading performance.

QUESTION 32
Get Cloudy Consulting is migrating their legacy system's users and data to Salesforce. They will be creating 15,000 users, 1.5 million Account records, and 15 million Invoice records. The visibility of these records is controlled by a 50 owner and criteria-based sharing rules.
Get Cloudy Consulting needs to minimize data loading time during this migration to a new organization.
Which two approaches will accomplish this goal? (Choose two.)

A. Create the users, upload all data, and then deploy the sharing rules.
B. Contact Salesforce to activate indexing before uploading the data.
C. First, load all account records, and then load all user records.
D. Defer sharing calculations until the data has finished uploading.

Correct Answer: A, D
Section:
Explanation:
Creating the users, uploading all data, and then deploying the sharing rules will reduce the number of sharing recalculations that occur during the data load. Deferring sharing calculations until the data has finished uploading will also improve the performance by postponing the sharing rule evaluation.These are the recommended best practices for loading large data sets into Salesforce

QUESTION 33
Cloud Kicks has the following requirements:
- Data needs to be sent from Salesforce to an external system to generate invoices from their Order Management System (OMS).
- A Salesforce administrator must be able to customize which fields will be sent to the external system without changing code.
What are two approaches for fulfilling these requirements? (Choose two.)

A. A set<sobjectFieldset> to determine which fields to send in an HTTP callout.
B. An Outbound Message to determine which fields to send to the OMS.
C. A Field Set that determines which fields to send in an HTTP callout.
D. Enable the field -level security permissions for the fields to send.

Correct Answer: B, C
Section:
Explanation:
An Outbound Message is a native Salesforce feature that allows sending data to an external system without code.It can be configured to include any fields from the source object3. A Field Set is a collection of fields that can be used in Visualforce pages or Apex classes to dynamically determine the fields to send in an HTTP callout. Both of these approaches meet the requirements of Cloud Kicks.

QUESTION 34
The architect is planning a large data migration for Universal Containers from their legacy CRM system to Salesforce. What three things should the architect consider to optimize performance of the data migration? Choose 3 answers

A. Review the time zones of the User loading the data.
B. Remove custom indexes on the data being loaded.
C. Determine if the legacy system is still in use.
D. Defer sharing calculations of the Salesforce Org.
E. Deactivate approval processes and workflow rules.

Correct Answer: B, D, E
Section:
Explanation:
Removing custom indexes on the data being loaded will prevent unnecessary index maintenance and improve the data load speed. Deferring sharing calculations of the Salesforce Org will avoid frequent sharing rule

evaluations and reduce the load time. Deactivating approval processes and workflow rules will prevent triggering any automation logic that might slow down or fail the data load.

QUESTION 35
Universal Containers has a large volume of Contact data going into Salesforce.com. There are 100,000 existing contact records. 200,000 new contacts will be loaded. The Contact object has an external ID field that is unique and must be populated for all existing records. What should the architect recommend to reduce data load processing time?

A. Load Contact records together using the Streaming API via the Upsert operation.
B. Delete all existing records, and then load all records together via the Insert operation.
C. Load all records via the Upsert operation to determine new records vs. existing records.
D. Load new records via the Insert operation and existing records via the Update operation.

Correct Answer: D
Section:
Explanation:
Loading new records via the Insert operation and existing records via the Update operation will allow using the external ID field as a unique identifier and avoid any duplication or overwriting of records. This is faster and safer than deleting all existing records or using the Upsert operation, which might cause conflicts or errors .

QUESTION 36
An architect is planning on having different batches to load one million Opportunities into Salesforce using the Bulk API in parallel mode. What should be considered when loading the Opportunity records?

A. Create indexes on Opportunity object text fields.
B. Group batches by the AccountId field.
C. Sort batches by Name field values.
D. Order batches by Auto -number field.

Correct Answer: D
Section:
Explanation:
Ordering batches by Auto-number field will ensure that the records are processed in a sequential order and avoid any locking issues that might occur when loading related records in parallel mode. Creating indexes, grouping batches by AccountId, or sorting batches by Name field values are not necessary or beneficial for loading Opportunity records using the Bulk API.

QUESTION 37
Universal Containers wants to develop a dashboard in Salesforce that will allow Sales Managers to do data exploration using their mobile device (i.e., drill down into sales-related data) and have the possibility of adding ad-hoc filters while on the move. What is a recommended solution for building data exploration dashboards in Salesforce?

A. Create a Dashboard in an external reporting tool, export data to the tool, and add link to the dashboard in Salesforce.
B. Create a Dashboard in an external reporting tool, export data to the tool, and embed the dashboard in Salesforce using the Canval toolkit.
C. Create a standard Salesforce Dashboard and connect it to reports with the appropriate filters.
D. Create a Dashboard using Analytics Cloud that will allow the user to create ad-hoc lenses and drill down.

Correct Answer: D
Section:
Explanation:
Creating a Dashboard using Analytics Cloud that will allow the user to create ad-hoc lenses and drill down is a recommended solution for building data exploration dashboards in Salesforce. Analytics Cloud is a powerful data analysis tool that enables users to explore data using interactive dashboards, charts, graphs, and tables on any device. Users can also create lenses, which are ad-hoc data queries that can be saved and reused, and drill down into data details using filters and facets. Creating a Dashboard in an external reporting tool, exporting data to the tool, and adding link to the dashboard in Salesforce will not provide a seamless user experience and may require additional data integration and security considerations. Creating a Dashboard in an external reporting tool, exporting data to the tool, and embedding the dashboard in Salesforce using the Canval toolkit will not provide a native Salesforce solution and may require additional data integration and security considerations. Creating a standard Salesforce Dashboard and connecting it to reports with the appropriate filters will not allow the

user to create ad-hoc lenses and drill down into data details on their mobile device.

QUESTION 38
DreamHouse Realty has a legacy system that captures Branch Offices and Transactions. DreamHouse Realty has 15 Branch Offices. Transactions can relate to any Branch Office. DreamHouse Realty has created hundreds of thousands of Transactions per year.
A Data Architect needs to denormalize this data model into a single Transaction object with a Branch Office picklist.
What are two important considerations for the Data Architect in this scenario? (Choose two.)

A. Standard list view in-line editing.
B. Limitations on Org data storage.
C. Bulk API limitations on picklist fields.
D. Limitations on master-detail relationships.

Correct Answer: B, C
Section:
Explanation:
The Data Architect should consider the limitations on Org data storage and the Bulk API limitations on picklist fields when denormalizing the data model into a single Transaction object with a Branch Office picklist.The Org data storage limit is the total amount of data that can be stored in a Salesforce Org, and it depends on the edition and license type of the Org1.The Bulk API limit on picklist fields is the maximum number of values that can be imported or exported using the Bulk API, and it is 1,000 values per picklist field2. These limitations could affect the performance and scalability of the data model, and the Data Architect should plan accordingly.

QUESTION 39
Cloud Kicks is launching a Partner Community, which will allow users to register shipment requests that are then processed by Cloud Kicks employees. Shipment requests contain header information, and then a list of no more than 5 items being shipped.
First, Cloud Kicks will introduce its community to 6,000 customers in North America, and then to 24,000 customers worldwide within the next two years. Cloud Kicks expects 12 shipment requests per week per customer, on average, and wants customers to be able to view up to three years of shipment requests and use Salesforce reports.
What is the recommended solution for the Cloud Kicks Data Architect to address the requirements?

A. Create an external custom object to track shipment requests and a child external object to track shipment items. External objects are stored off-platform in Heroku's Postgres database.
B. Create an external custom object to track shipment requests with five lookup custom fields for each item being shipped. External objects are stored off-platform in Heroku's Postgres database.
C. Create a custom object to track shipment requests and a child custom object to track shipment items. Implement an archiving process that moves data off-platform after three years.
D. Create a custom object to track shipment requests with five lookup custom fields for each item being shipped Implement an archiving process that moves data off-platform after three years.

Correct Answer: C
Section:
Explanation:
The recommended solution for the Cloud Kicks Data Architect to address the requirements is to create a custom object to track shipment requests and a child custom object to track shipment items. Implement an archiving process that moves data off-platform after three years. This solution would allow Cloud Kicks to store and manage their shipment data on Salesforce, and use Salesforce reports to analyze it.However, since Cloud Kicks expects a large volume of data over time, they should implement an archiving process that moves data off-platform after three years to avoid hitting the Org data storage limit and maintain optimal performance3.External objects are not a good option for this scenario, because they are stored off-platform in an external system, such as Heroku's Postgres database, and they have limited functionality and performance compared to custom objects

QUESTION 40
Universal Containers has successfully migrated 50 million records into five different objects multiple times in a full copy sandbox. The Integration Engineer wants to re-run the test again a month before it goes live into Production. What is the recommended approach to re-run the test?

A. Truncate all 5 objects quickly and re-run the data migration test.
B. Refresh the full copy sandbox and re-run the data migration test.
C. Hard delete all 5 objects' data and re-run the data migration test.
D. Truncate all 5 objects and hard delete before running the migration test.

Correct Answer: B
Section:
Explanation:
The recommended approach to re-run the test is to refresh the full copy sandbox and re-run the data migration test. A full copy sandbox is a replica of the production Org, including all data, metadata, and attachments. Refreshing a full copy sandbox means creating a new copy of the production Org and replacing the existing sandbox. This would ensure that the test is run on a clean and up-to-date environment, without any leftover data or configuration from previous tests. Truncating or hard deleting objects would not be sufficient, because they would not remove all the data or metadata from the sandbox, and they could also affect other dependent objects or processes.

QUESTION 41
Universal Containers is integrating a new Opportunity engagement system with Salesforce. According to their Master Data Management strategy, Salesforce is the system of record for Account, Contact, and Opportunity dat a. However, there does seem to be valuable Opportunity data in the new system that potentially conflicts with what is stored in Salesforce. What is the recommended course of action to appropriately integrate this new system?

A. The MDM strategy defines Salesforce as the system of record, so Salesforce Opportunity values prevail in all conflicts.
B. A policy should be adopted so that the system whose record was most recently updated should prevail in conflicts.
C. The Opportunity engagement system should become the system of record for Opportunity records.
D. Stakeholders should be brought together to discuss the appropriate data strategy moving forward.

Correct Answer: D
Section:
Explanation:
The recommended course of action to appropriately integrate the new Opportunity engagement system with Salesforce is to bring the stakeholders together to discuss the appropriate data strategy moving forward. This is because there may be valuable data in both systems that need to be reconciled and harmonized, and the Master Data Management (MDM) strategy may need to be revised or updated to accommodate the new system.The other options are not recommended, as they may result in data loss, inconsistency, or duplication

QUESTION 42
For a production cutover, a large number of Account records will be loaded into Salesforce from a legacy system. The legacy system does not have enough information to determine the Ownership for these Accounts upon initial load. Which two recommended options assign Account ownership to mitigate potential performance problems?

A. Let a ''system user'' own all the Account records without assigning any role to this user in Role Hierarchy.
B. Let a ''system user'' own the Account records and assign this user to the lowest-level role in the Role Hierarchy.
C. Let the VP of the Sales department, who will report directly to the senior VP, own all the Account records.
D. Let a ''system user'' own all the Account records and make this user part of the highest-level role in the Role Hierarchy.

Correct Answer: A, B
Section:
Explanation:
The two recommended options to assign Account ownership to mitigate potential performance problems are to let a ''system user'' own all the Account records without assigning any role to this user in Role Hierarchy, or to let a ''system user'' own the Account records and assign this user to the lowest-level role in the Role Hierarchy.This is because these options would reduce the number of sharing calculations and rules that need to be applied to the Account records, and improve the performance and scalability of the system34. The other options are not recommended, as they would increase the sharing complexity and overhead, and potentially expose sensitive data to unauthorized users.

QUESTION 43
Universal Containers (UC) is implementing its new Internet of Things technology, which consists of smart containers that provide information on container temperature and humidity updated every 10 minutes back to UC. There are roughly 10,000 containers equipped with this technology with the number expected to increase to 50,000 across the next five years. It is essential that Salesforce user have access to current and historical temperature and humidity data for each container. What is the recommended solution?

A. Create new custom fields for temperature and humidity in the existing Container custom object, as well as an external ID field that is unique for each container. These custom fields are updated when a new measure is

received.

B. Create a new Container Reading custom object, which is created when a new measure is received for a specific container. The Container Reading custom object has a master-detail relationship to the container object.
C. Create a new Lightning Component that displays last humidity and temperature data for a specific container and can also display historical trends obtaining relevant data from UC's existing data warehouse.
D. Create a new Container Reading custom object with a master-detail relationship to Container which is created when a new measure is received for a specific container. Implement an archiving process that runs every hour.

Correct Answer: D
Section:
Explanation:
The recommended solution for Universal Containers (UC) to implement its new Internet of Things technology is to create a new Container Reading custom object with a master-detail relationship to Container which is created when a new measure is received for a specific container. Implement an archiving process that runs every hour. This solution would allow UC to store and access current and historical temperature and humidity data for each container on Salesforce, and use reports and dashboards to analyze it. However, since UC expects a large volume of data over time, they should implement an archiving process that moves data off-platform after a certain period of time to avoid hitting the Org data storage limit and maintain optimal performance. The other options are not recommended, as they would either not store the historical data on Salesforce, or create too many custom fields on the Container object that could impact performance and usability.

QUESTION 44
Universal Containers is planning out their archiving and purging plans going forward for their custom objects Topic__c and Comment__c. Several options are being considered, including analytics snapshots, offsite storage, scheduled purges, etc. Which three questions should be considered when designing an appropriate archiving strategy?

A. How many fields are defined on the custom objects that need to be archived?
B. Which profiles and users currently have access to these custom object records?
C. If reporting is necessary, can the information be aggregated into fewer, summary records?
D. Will the data being archived need to be reported on or accessed in any way in the future?
E. Are there any regulatory restrictions that will influence the archiving and purging plans?

Correct Answer: C, D, E
Section:
Explanation:
The three questions that should be considered when designing an appropriate archiving strategy are: If reporting is necessary, can the information be aggregated into fewer, summary records? Will the data being archived need to be reported on or accessed in any way in the future? Are there any regulatory restrictions that will influence the archiving and purging plans?These questions are important because they help determine the scope, frequency, and method of archiving and purging data from Salesforce. For example, if reporting is necessary, then summary records or analytics snapshots can be used to store aggregated data and reduce the number of records that need to be archived1.If the data being archived needs to be accessed in the future, then offsite storage or external objects can be used to retain the data and make it available on demand2.If there are any regulatory restrictions, such as GDPR or HIPAA, then the archiving and purging plans need to comply with them and ensure data security and privacy

QUESTION 45
Which three characteristics of a skinny table help improve report and query performance?

A. Skinny tables can contain frequently used fields and thereby help avoid joins.
B. Skinny tables can be used to create custom indexes on multi-select picklist fields.
C. Skinny tables provide a view across multiple objects for easy access to combined data.
D. Skinny tables are kept in sync with changes to data in the source tables.
E. Skinny tables do not include records that are available in the recycle bin.

Correct Answer: A, D, E
Section:
Explanation:
The three characteristics of a Skinny table that help improve report and query performance are: Skinny tables can contain frequently used fields and thereby help avoid joins. Skinny tables are kept in sync with changes to data in the source tables. Skinny tables do not include records that are available in the recycle bin. These characteristics are beneficial because they reduce the query complexity and execution time, and improve the data accuracy

and freshness.For example, skinny tables can contain frequently used fields from multiple objects, such as Account and Contact, and thereby help avoid joins that can slow down queries4.Skinny tables are updated automatically when the source tables are modified, so they always reflect the latest data5. Skinny tables do not include records that are available in the recycle bin, so they only contain active records that are relevant for reports and queries.

QUESTION 46
Universal Containers (UC) has a multi-level master-detail relationship for opportunities, a custom opportunity line item object, and a custom discount request. UC has opportunity as master and custom line item object as detail in master-detail relationship. UC also has a custom line item object as master and a custom discount request object as detail in another master-detail relationship. UC has a requirement to show all sums of discounts across line items at an opportunity level. What is the recommended solution to address these requirements?

A. Use roll-up for the line-item-level summary and a trigger for the opportunity amount summary, as only one level roll-up is allowed.
B. Update the master-detail relationships to lookup relationships in order to allow the discount amount to roll up.
C. Remove the master-detail relationships and rely completely on workflow/triggers to summarize the discount amount.
D. Roll-up discount request amount at the line-item-level and line-item-level summary discount at the opportunity level.

Correct Answer: D
Section:
Explanation:
The recommended solution to address the requirements of Universal Containers (UC) is to roll-up discount request amount at the line-item-level and line-item-level summary discount at the opportunity level. This solution would allow UC to show all sums of discounts across line items at an opportunity level using standard roll-up summary fields. Since UC has a multi-level master-detail relationship for opportunities, a custom opportunity line item object, and a custom discount request object, they can use two levels of roll-up summary fields to aggregate the discount amounts from the child objects to the parent objects. For example, they can create a roll-up summary field on the custom line item object to sum up the discount request amounts from the custom discount request object. Then, they can create another roll-up summary field on the opportunity object to sum up the line-item-level summary discounts from the custom line item object.

QUESTION 47
Which two best practices should be followed when using SOSL for searching?

A. Use searches against single Objects for greater speed and accuracy.
B. Keep searches specific and avoid wildcards where possible.
C. Use SOSL option to ignore custom indexes as search fields are pre-indexed.
D. Use Find in ''ALL FIELDS'' for faster searches.

Correct Answer: B, D
Section:
Explanation:
The two best practices that should be followed when using SOSL for searching are: Keep searches specific and avoid wildcards where possible. Use Find in ''ALL FIELDS'' for faster searches. These best practices are helpful because they optimize the search performance and relevance. For example, keeping searches specific and avoiding wildcards where possible can reduce the number of results returned and improve the search accuracy. Using Find in ''ALL FIELDS'' for faster searches can leverage Salesforce's pre-defined indexes on common fields and return results more quickly than specifying individual fields.

QUESTION 48
Universal Containers (UC) maintains a collection of several million Account records that represent business in the United Sates. As a logistics company, this list is one of the most valuable and important components of UC's business, and the accuracy of shipping addresses is paramount. Recently it has been noticed that too many of the addresses of these businesses are inaccurate, or the businesses don't exist. Which two scalable strategies should UC consider to improve the quality of their Account addresses?

A. Contact each business on the list and ask them to review and update their address information.
B. Build a team of employees that validate Accounts by searching the web and making phone calls.
C. Integrate with a third-party database or services for address validation and enrichment.
D. Leverage Data.com Clean to clean up Account address fields with the D&B database.

Correct Answer: C, D
Section:
Explanation:
Integrating with a third-party database or service for address validation and enrichment is a scalable strategy that can improve the quality of the Account addresses by comparing them with a reliable source of data1.Leveraging Data.com Clean to clean up Account address fields with the D&B database is another scalable strategy that can automatically update and enrich Account records with verified information from Data.com2.

QUESTION 49
Universal Containers (UC) loads bulk leads and campaigns from third-party lead aggregators on a weekly and monthly basis. The expected lead record volume is 500K records per week, and the expected campaign records volume is 10K campaigns per week. After the upload, Lead records are shared with various sales agents via sharing rules and added as Campaign members via Apex triggers on Lead creation. UC agents work on leads for 6 months, but want to keep the records in the system for at least 1 year for reference. Compliance requires them to be stored for a minimum of 3 years. After that, data can be deleted. What statement is true with respect to a data archiving strategy for UC?

A. UC can store long-term lead records in custom storage objects to avoid counting against storage limits.
B. UC can leverage the Salesforce Data Backup and Recovery feature for data archival needs.
C. UC can leverage recycle bin capability, which guarantees record storage for 15 days after deletion.
D. UC can leverage a ''tier''-based approach to classify the record storage need.

Correct Answer: D
Section:
Explanation:
Leveraging a ''tier''-based approach to classify the record storage need is a true statement with respect to a data archiving strategy for UC.This approach involves defining different tiers of data based on their usage, value, and retention policies, and then applying appropriate storage and archiving solutions for each tier.

QUESTION 50
Universal Containers (UC) is implementing Salesforce Sales Cloud and Service Cloud. As part of their implementation, they are planning to create a new custom object (Shipments), which will have a lookup relationship to Opportunities. When creating shipment records, Salesforce users need to manually input a customer reference, which is provided by customers, and will be stored in the Customer_Reference__c text custom field. Support agents will likely use this customer reference to search for Shipment records when resolving shipping issues. UC is expecting to have around 5 million shipment records created per year. What is the recommended solution to ensure that support agents using global search and reports can quickly find shipment records?

A. Implement an archiving process for shipment records created after five years.
B. Implement an archiving process for shipment records created after three years.
C. Set Customer-Reference_c as an External ID (non-unique).
D. Set Customer-Reference_c as an External ID (unique).

Correct Answer: D
Section:
Explanation:
Setting Customer-Reference_c as an External ID (unique) is the recommended solution to ensure that support agents using global search and reports can quickly find shipment records.This is because external ID fields are indexed and can be used as filters in reports and list views4. Making the field unique also ensures that there are no duplicate values that could confuse the search results.

QUESTION 51
Universal Containers has a custom object with millions of rows of data.
When executing SOQL queries, which three options prevent a query from being selective? (Choose three.)

A. Using leading % wildcards.
B. Using trailing % wildcards.
C. Performing large loads and deletions.
D. Using NOT and != operators.

E. Using a custom index on a deterministic formula field.

Correct Answer: A, C, D
Section:
Explanation:
Using leading % wildcards, performing large loads and deletions, and using NOT and != operators are three options that prevent a query from being selective.These options can cause the query to scan a large number of records or indexes, which can negatively impact the performance and scalability of the query5.Using a custom index on a deterministic formula field, on the other hand, can improve the query selectivity by allowing the query to use the index instead of scanning all the records6.

QUESTION 52
Universal Containers (UC) wants to store product data in Salesforce, but the standard Product object does not support the more complex hierarchical structure which is currently being used in the product master system. How can UC modify the standard Product object model to support a hierarchical data structure in order to synchronize product data from the source system to Salesforce?

A. Create a custom lookup filed on the standard Product to reference the child record in the hierarchy.
B. Create a custom lookup field on the standard Product to reference the parent record in the hierarchy.
C. Create a custom master-detail field on the standard Product to reference the child record in the hierarchy.
D. Create an Apex trigger to synchronize the Product Family standard picklist field on the Product object.

Correct Answer: B
Section:
Explanation:
Creating a custom lookup field on the standard Product to reference the parent record in the hierarchy is the correct way to modify the standard Product object model to support a hierarchical data structure. This allows UC to create a self-relationship on the Product object and define parent-child relationships among products.

QUESTION 53
UC is having issues using Informatica Cloud Louder to export +10MOrder records. Each Order record has 10 Order Line Items. What two steps can you take to help correct this? Choose two answers.

A. Export in multiple batches
B. Export Bulk API in parallel mode
C. Use PK Chunking
D. Limit Batch to 10K records

Correct Answer: A, C
Section:
Explanation:
Exporting in multiple batches and using PK Chunking are two steps that can help correct the issues with exporting large volumes of Order records using Informatica Cloud Loader. Exporting in multiple batches can reduce the load on the system and avoid timeouts or errors. Using PK Chunking can split a large data set into smaller chunks based on the record IDs and enable parallel processing of each chunk.

QUESTION 54
Universal Containers would like to have a Service-Level Agreement (SLA) of 1 day for any data loss due to unintentional or malicious updates of records in Salesforce. What approach should be suggested to address this requirement?

A. Build a daily extract job and extract data to on-premise systems for long-term backup and archival purposes.
B. Schedule a Weekly Extract Service for key objects and extract data in XL sheets to on-premise systems.
C. Store all data in shadow custom objects on any updates and deletes, and extract them as needed .
D. Evaluate a third-party AppExchange app, such as OwnBackup or Spanning, etc., for backup and archival purposes.

Correct Answer: D

Section:
Explanation:
Evaluating a third-party AppExchange app, such as OwnBackup or Spanning, etc., for backup and archival purposes is the best approach to address the requirement of having a one-day SLA for any data loss. These apps can provide automated backups, granular restores, and compliance features that can help UC recover from any data loss or corruption scenarios.

QUESTION 55
Universal Containers (UC) has 1,000 accounts and 50,000 opportunities. UC has an enterprise security requirement to export all sales data outside of Salesforce on a weekly basis. The security requirement also calls for exporting key operational data that includes events such as file downloads, logins, logouts, etc. Which two recommended approaches would address the above requirement?

A. Use Field Audit History to capture operational data and extract it to on-premise systems.
B. Use Weekly Export to extract transactional data to on-premise systems.
C. Use a custom built extract job to extract operational data to on-premise systems.
D. Use Event Monitoring to extract event data to on-premise systems.

Correct Answer: B, D
Section:
Explanation:
Using Weekly Export to extract transactional data to on-premise systems and using Event Monitoring to extract event data to on-premise systems are two recommended approaches that would address the security requirement of UC. Weekly Export is a built-in feature that allows UC to download a zip file of their Salesforce data on a weekly basis. Event Monitoring is a tool that allows UC to access detailed information about various events that occur in their Salesforce org, such as logins, logouts, file downloads, etc.

QUESTION 56
Universal Containers has a public website with several forms that create Lead records in Salesforce using the REST API. When designing these forms, which two techniques will help maintain a high level of data quality?

A. Do client-side validation of phone number and email field formats.
B. Prefer picklist form fields over free text fields, where possible.
C. Ensure the website visitor is browsing using an HTTPS connection.
D. Use cookies to track when visitors submit multiple forms.

Correct Answer: A, B
Section:
Explanation:
Doing client-side validation of phone number and email field formats and preferring picklist form fields over free text fields, where possible, are two techniques that will help maintain a high level of data quality when designing forms that create Lead records in Salesforce using the REST API. Client-side validation can ensure that the input data is in the correct format and prevent errors or invalid values from being submitted. Picklist form fields can provide a predefined set of values for the users to choose from and avoid typos or inconsistencies in the data.

QUESTION 57
Universal Containers is exporting 40 million Account records from Salesforce using Informatica Cloud. The ETL tool fails and the query log indicates a full table scan time-out failure. What is the recommended solution?

A. Modify the export job header to specify Export-in-Parallel.
B. Modify the export job header to specify Sforce-Enable-PKChunking.
C. Modify the export query that includes standard index fields(s).
D. Modify the export query with LIMIT clause with Batch size 10,000.

Correct Answer: B
Section:
Explanation:
The Sforce-Enable-PKChunking header enables you to extract large data sets from Salesforce by using a technique called primary key (PK) chunking. PK chunking splits your query into multiple queries based on the record IDs

of the queried object.This reduces the query time and avoids the query timeout issues that can occur with large data volumes

QUESTION 58
Universal Containers (UC) wants to capture information on how data entities are stored within the different applications and systems used within the company. For that purpose, the architecture team decided to create a data dictionary covering the main business domains within UC. Which two common techniques are used building a data dictionary to store information on how business entities are defined?

A. Use Salesforce Object Query Language.
B. Use a data definition language.
C. Use an entity relationship diagram.
D. Use the Salesforce Metadata API.

Correct Answer: C, D
Section:
Explanation:
A data dictionary is a document that describes the structure, format, and meaning of data entities and attributes. A common technique to build a data dictionary is to use an entity relationship diagram (ERD), which shows the logical relationships between objects and fields in a graphical way.Another technique is to use the Salesforce Metadata API, which allows you to retrieve and deploy the metadata that defines your Salesforce org

QUESTION 59
Universal Containers has received complaints that customers are being called by multiple Sales Reps where the second Sales Rep that calls is unaware of the previous call by their coworker. What is a data quality problem that could cause this?

A. Missing phone number on the Contact record.
B. Customer phone number has changed on the Contact record.
C. Duplicate Contact records exist in the system.
D. Duplicate Activity records on a Contact.

Correct Answer: C
Section:
Explanation:
A data quality problem that could cause customers to be called by multiple Sales Reps is having duplicate Contact records in the system. Duplicate records can result from data entry errors, data imports, or integrations with other systems.Duplicate records can lead to confusion, inefficiency, and customer dissatisfaction

QUESTION 60
Universal Containers (UC) is a major supplier of office supplies. Some products are produced by UC and some by other manufacturers. Recently, a number of customers have complained that product descriptions on the invoices do not match the descriptions in the online catalog and on some of the order confirmations (e.g., 'ballpoint pen' in the catalog and 'pen' on the invoice, and item color labels are inconsistent: 'what vs. 'White' or 'blk' vs. 'Black'). All product data is consolidated in the company data warehouse and pushed to Salesforce to generate quotes and invoices. The online catalog and webshop is a Salesforce Customer Community solution. What is a correct technique UC should use to solve the data inconsistency?

A. Change integration to let product master systems update product data directly in Salesforce via the Salesforce API.
B. Add custom fields to the Product standard object in Salesforce to store data from the different source systems.
C. Define a data taxonomy for product data and apply the taxonomy to the product data in the data warehouse.
D. Build Apex Triggers in Salesforce that ensure products have the correct names and labels after data is loaded into salesforce.

Correct Answer: C
Section:
Explanation:
A correct technique UC should use to solve the data inconsistency is to define a data taxonomy for product data and apply the taxonomy to the product data in the data warehouse. A data taxonomy is a hierarchical classification of data entities and attributes that defines their meaning, format, and relationships.A data taxonomy can help ensure consistency, accuracy, and completeness of product data across different systems and

channels

QUESTION 61
What makes Skinny tables fast? Choose three answers.

A. They do not include soft-deleted records
B. They avoid resource intensive joins
C. Their tables are kept in sync with their source tables when the source tables are modified
D. They can contain fields from other objects
E. They support up to a max of 100 of columns

Correct Answer: A, B, C
Section:
Explanation:
Skinny tables are custom tables that contain frequently used fields from a standard or custom object. They are used to improve performance by reducing the number of database joins required for queries.Skinny tables have the following characteristics1:
They do not include soft-deleted records, which means they only contain active records and save space.
They avoid resource intensive joins by storing data from multiple objects in one table, which reduces the query time and complexity.
Their tables are kept in sync with their source tables when the source tables are modified, which ensures data consistency and accuracy.

QUESTION 62
Universal Containers is creating a new B2C service offering for consumers to ship goods across continents. This is in addition to their well-established B2B offering. Their current Salesforce org uses the standard Account object to track B2B customers. They are expecting to have over 50,000,000 consumers over the next five years across their 50 business regions. B2C customers will be individuals. Household data is not required to be stored. What is the recommended data model for consumer account data to be stored in Salesforce?

A. Use the Account object with Person Accounts and a new B2C page layout.
B. Use the Account object with a newly created Record Type for B2C customers.
C. Create a new picklist value for B2C customers on the Account Type field.
D. Use 50 umbrella Accounts for each region, with customers as associated Contacts.

Correct Answer: A
Section:
Explanation:
The recommended data model for consumer account data to be stored in Salesforce is to use the Account object with Person Accounts and a new B2C page layout.Person Accounts are a special type of accounts that allow you to store information about individual consumers who are not associated with a business account2.Person Accounts have the following advantages3:
They allow you to use the same standard objects and features that you use for business accounts, such as contacts, opportunities, cases, etc.
They enable you to create different page layouts and record types for B2C and B2B customers, which allows you to customize the user interface and business logic for each segment.
They support large data volumes and can scale up to 50 million records or more, which meets the expected growth of consumers over the next five years.

QUESTION 63
Universal Containers (UC) is implementing a new customer categorization process where customers should be assigned to a Gold, Silver, or Bronze category if they've purchased UC's new support service. Customers are expected to be evenly distributed across all three categories. Currently, UC has around 500,000 customers, and is expecting 1% of existing non-categorized customers to purchase UC's new support service every month over the next five years. What is the recommended solution to ensure long-term performance, bearing in mind the above requirements?

A. Implement a new global picklist custom field with Gold, Silver, and Bronze values and enable it in Account.
B. Implement a new picklist custom field in the Account object with Gold, Silver, and Bronze values.
C. Implement a new Categories custom object and a master-detail relationship from Account to Category.
D. Implement a new Categories custom object and create a lookup field from Account to Category.

Correct Answer: B
Section:
Explanation:
The recommended solution to ensure long-term performance for the customer categorization process is to implement a new picklist custom field in the Account object with Gold, Silver, and Bronze values. A picklist field is a simple and efficient way to store a predefined set of values for a record. A picklist field has the following benefits:
It allows you to easily filter, sort, and group records by their category values, which can help you analyze and segment your customers.
It does not require creating any additional objects or relationships, which can reduce the data model complexity and maintenance overhead.
It supports large data volumes and can handle millions of records without affecting performance or scalability.

QUESTION 64
Ursa Major Solar's legacy system has a quarterly accounts receivable report that compiles data from the following:
- Accounts
- Contacts
- Opportunities
- Orders
- Order Line Items

Which issue will an architect have when implemented this in Salesforce?

A. Custom report types CANNOT contain Opportunity data.
B. Salesforce does NOT support Orders or Order Line Items.
C. Salesforce does NOT allow more than four objects in a single report type.
D. A report CANNOT contain data from Accounts and Contacts.

Correct Answer: C
Section:
Explanation:
The issue that an architect will have when implementing the quarterly accounts receivable report in Salesforce is that Salesforce does not allow more than four objects in a single report type. A report type defines the set of records and fields available to a report based on the relationships between a primary object and up to four related objects. A report type has the following limitations:
It cannot include more than four objects in a single report type, which means that the report cannot compile data from five objects (Accounts, Contacts, Opportunities, Orders, and Order Line Items) at once.
It cannot include objects that are more than two relationships away from each other, which means that the report cannot access fields from Order Line Items through Opportunities and Orders.
It cannot include objects that have a many-to-many relationship with each other, which means that the report cannot access fields from Contacts and Opportunities through the junction object Opportunity Contact Role.

QUESTION 65
Two million Opportunities need to be loaded in different batches into Salesforce using the Bulk API in parallel mode.
What should an Architect consider when loading the Opportunity records?

A. Use the Name field values to sort batches.
B. Order batches by Auto-number field.
C. Create indexes on Opportunity object text fields.
D. Group batches by the AccountId field.

Correct Answer: D
Section:
Explanation:
Grouping batches by the AccountId field can improve the performance and avoid locking issues when loading Opportunity records using the Bulk API in parallel mode1.This is because the AccountId field is indexed and can be used to distribute the records evenly across batches

QUESTION 66
Ursa Major Solar has defined a new Data Quality Plan for their Salesforce data.

Which two approaches should an Architect recommend to enforce the plan throughout the organization? (Choose two.)

A. Ensure all data is stored in an external system and set up an integration to Salesforce for view-only access.
B. Schedule reports that will automatically catch duplicates and merge or delete the records every week.
C. Enforce critical business processes by using Workflow, Validation Rules, and Apex code.
D. Schedule a weekly dashboard displaying records that are missing information to be sent to managers for review.

Correct Answer: C, D
Section:
Explanation:
Enforcing critical business processes by using Workflow, Validation Rules, and Apex code can help ensure data quality and consistency by applying rules and logic to the data entry and update3.Scheduling a weekly dashboard displaying records that are missing information to be sent to managers for review can help identify and fix data quality issues by providing visibility and accountability4.

QUESTION 67
DreamHouse Realty has a Salesforce deployment that manages Sales, Support, and Marketing efforts in a multi-system ERP environment. The company recently reached the limits of native reports and dashboards and needs options for providing more analytical insights.
What are two approaches an Architect should recommend? (Choose two.)

A. Weekly Snapshots
B. Einstein Analytics
C. Setup Audit Trails
D. AppExchange Apps

Correct Answer: B, D
Section:
Explanation:
Einstein Analytics can provide more analytical insights than native reports and dashboards by allowing users to explore data from multiple sources, create interactive visualizations, and apply AI-powered features5.AppExchange Apps can also provide more analytical insights by offering pre-built solutions or integrations with external tools that can enhance the reporting and analytics capabilities of Salesforce6.

QUESTION 68
Cloud Kicks currently has a Public Read/Write sharing model for the company's Contacts. Cloud Kicks management team requests that only the owner of a contact record be allowed to delete that contact.
What should an Architect do to meet these requirements?

A. Set the profile of the users to remove delete permission from the Contact object.
B. Check if the current user is NOT the owner by creating a 'before delete' trigger.
C. Set the Sharing settings as Public Read Only for the Contact object.
D. Check if the current user is NOT the owner by creating a validation rule on the Contact object.

Correct Answer: B
Section:
Explanation:
Checking if the current user is NOT the owner by creating a ''before delete'' trigger can meet the requirement of allowing only the owner of a contact record to delete that contact. A trigger is a piece of Apex code that can execute before or after a record is inserted, updated, deleted, or undeleted. A ''before delete'' trigger can prevent the deletion of a record by using theaddError()method.

QUESTION 69
An Architect needs information about who is creating, changing, or deleting certain fields within the past four months.
How can the Architect access this information?

A. Create a field history report for the fields in question.
B. After exporting the setup audit trail, find the fields in question.
C. After exporting the metadata, search it for the fields in question.
D. Remove 'customize application' permissions from everyone else.

Correct Answer: B
Section:
Explanation:
Exporting the setup audit trail can provide information about who is creating, changing, or deleting certain fields within the past four months. The setup audit trail tracks the recent setup changes that administrators and other users have made to the organization. The setup audit trail history shows up to 20 most recent changes in the Setup area, but administrators can download a report (in CSV format) of up to six months of setup history.

QUESTION 70
Universal Containers has more than 10 million records in the Order_c object. The query has timed out when running a bulk query. What should be considered to resolve query timeout?

A. Tooling API
B. PK Chunking
C. Metadata API
D. Streaming API

Correct Answer: B
Section:
Explanation:
PK Chunking can resolve query timeout when running a bulk query on an object with more than 10 million records. PK Chunking is a feature of the Bulk API that splits a query into multiple batches based on the record IDs (primary keys) of the queried object. This can improve the query performance and avoid timeouts by reducing the number of records processed in each batch.

QUESTION 71
Universal Containers (UC) has a data model as shown in the image. The Project object has a private sharing model, and it has Roll -Up summary fields to calculate the number of resources assigned to the project, total hours for the project, and the number of work items associated to the project. What should the architect consider, knowing there will be a large amount of time entry records to be loaded regularly from an external system into Salesforce.com?

A. Load all data using external IDs to link to parent records.
B. Use workflow to calculate summary values instead of Roll -Up.
C. Use triggers to calculate summary values instead of Roll -Up.
D. Load all data after deferring sharing calculations.

Correct Answer: D
Section:
Explanation:
Loading all data after deferring sharing calculations can improve the performance and avoid locking issues when loading a large amount of time entry records into Salesforce.com.This is because deferring sharing calculations can temporarily suspend the calculation of sharing rules until all the data is loaded, and then recalculate them in one operation

QUESTION 72
Which two aspects of data does an Enterprise data governance program aim to improve?

A. Data integrity
B. Data distribution
C. Data usability

D. Data modeling

Correct Answer: A, C
Section:
Explanation:
Data integrity and data usability are two aspects of data that an Enterprise data governance program aims to improve.Data integrity refers to the accuracy, consistency, and validity of the data across the enterprise2.Data usability refers to the ease of access, analysis, and interpretation of the data by the end users

QUESTION 73
Universal Containers (UC) has over 10 million accounts with an average of 20 opportunities with each account. A Sales Executive at UC needs to generate a daily report for all opportunities in a specific opportunity stage.
Which two key considerations should be made to make sure the performance of the report is not degraded due to large data volume?

A. Number of queries running at a time.
B. Number of joins used in report query.
C. Number of records returned by report query.
D. Number of characters in report query.

Correct Answer: B, C
Section:
Explanation:
The number of joins used in report query and the number of records returned by report query are two key considerations to make sure the performance of the report is not degraded due to large data volume.The number of joins used in report query affects the complexity and execution time of the query, especially when joining multiple large objects4. The number of records returned by report query affects the amount of data that needs to be processed and displayed by the report engine.

QUESTION 74
A health care provider wishes to use salesforce to track patient care. The following actions are in Salesforce
1. Payment Providers: Orgas who pay for the care 2 patients.
2. Doctors: They provide care plan for patients and need to support multiple patients, they are provided access to patient information.
3. Patients: They are individuals who need care.
A data architect needs to map the actor to Sf objects. What should be the optimal selection by the data architect?

A. Patients as Contacts, Payment providers as Accounts, & Doctors as Accounts
B. Patients as Person Accounts, Payment providers as Accounts, & Doctors as Contacts
C. Patients as Person Accounts, Payment providers as Accounts, & Doctors as Person Account
D. Patients as Accounts, Payment providers as Accounts, & Doctors as Person Accounts

Correct Answer: C
Section:
Explanation:
Patients as Person Accounts, Payment providers as Accounts, & Doctors as Person Accounts is the optimal selection by the data architect to map the actor to Salesforce objects. This is because Person Accounts are a special type of accounts that can store both business and personal information for individual customers. Payment providers are organizations that pay for the care of patients, so they can be modeled as Accounts. Doctors are also individuals who provide care plans for patients and need access to patient information, so they can also be modeled as Person Accounts.

QUESTION 75
NTO (Northern Trail Outlets) has a complex Salesforce org which has been developed over past 5 years. Internal users are complaining abt multiple data issues, including incomplete and duplicate data in the org. NTO has decided to engage a data architect to analyze and define data quality standards.
Which 3 key factors should a data architect consider while defining data quality standards? Choose 3 answers:

A. Define data duplication standards and rules
B. Define key fields in staging database for data cleansing
C. Measure data timeliness and consistency
D. Finalize an extract transform load (ETL) tool for data migration
E. Measure data completeness and accuracy

Correct Answer: A, C, E
Section:
Explanation:
Defining data duplication standards and rules, measuring data timeliness and consistency, and measuring data completeness and accuracy are three key factors that a data architect should consider while defining data quality standards. Defining data duplication standards and rules can help prevent or reduce duplicate records in the org by specifying criteria and actions for identifying and merging duplicates. Measuring data timeliness and consistency can help ensure that the data is up-to-date, reliable, and synchronized across different sources. Measuring data completeness and accuracy can help ensure that the data is sufficient, relevant, and correct for the intended purposes.

QUESTION 76
Universal Containers (UC) requires 2 years of customer related cases to be available on SF for operational reporting. Any cases older than 2 years and upto 7 years need to be available on demand to the Service agents. UC creates 5 million cases per yr.
Which 2 data archiving strategies should a data architect recommend? Choose 2 options:

A. Use custom objects for cases older than 2 years and use nightly batch to move them.
B. Sync cases older than 2 years to an external database, and provide access to Service agents to the database
C. Use Big objects for cases older than 2 years, and use nightly batch to move them.
D. Use Heroku and external objects to display cases older than 2 years and bulk API to hard delete from Salesforce.

Correct Answer: C, D
Section:
Explanation:
The best data archiving strategies for UC are to use Big objects and Heroku with external objects. Big objects allow storing large amounts of data on the Salesforce platform without affecting performance or storage limits. They also support point-and-click tools, triggers, and Apex code. Heroku is a cloud platform that can host external databases and integrate with Salesforce using external objects. External objects enable on-demand access to external data sources via standard Salesforce APIs and user interfaces. Using bulk API to hard delete cases from Salesforce will free up storage space and improve performance.

QUESTION 77
NTO would like to retrieve their SF orgs meta data programmatically for backup within a various external. Which API is the best fit for accomplishing this task?

A. Metadata API
B. Tooling API
C. Bulk API in serial mode
D. SOAP API

Correct Answer: A
Section:
Explanation:
The best API for retrieving Salesforce org metadata programmatically is the Metadata API. The Metadata API provides access to the metadata that defines the structure and configuration of an org, such as custom objects, fields, workflows, security settings, etc. It also supports deploying, retrieving, creating, updating, and deleting metadata components. The Metadata API can be used with various tools, such as Ant, Workbench, or IDEs.

QUESTION 78
A customer is operating in a highly reputated industry and is planning to implement SF. The customer information maintained in SF, includes the following:
Personally, identifiable information (PII)

IP restrictions on profiles organized by Geographic location
Financial records that need to be private and accessible only by the assigned Sales associate.
User should not be allowed to export information from Salesforce.
Enterprise security has mandate access to be restricted to users within a specific geography and detail monitoring of user activity. Which 3 Salesforce shield capabilities should a data architect recommend? Choose 3 answers:

A. Event monitoring to monitor all user activities
B. Restrict access to SF from users outside specific geography
C. Prevent Sales users access to customer PII information
D. Transaction security policies to prevent export of SF Data.
E. Encrypt Sensitive Customer information maintained in SF.

Correct Answer: B, D, E
Section:
Explanation:
The best Salesforce Shield capabilities for the customer are to restrict access to SF from users outside specific geography, implement transaction security policies to prevent export of SF data, and encrypt sensitive customer information maintained in SF. Salesforce Shield is a set of security features that help protect enterprise data on the Salesforce platform. It includes three components: Event Monitoring, Platform Encryption, and Field Audit Trail. Restricting access to SF from users outside specific geography can be done using network-based security features, such as IP whitelisting or VPN. Transaction security policies can be used to define actions or notifications based on user behavior patterns, such as exporting data or logging in from an unknown device. Platform Encryption can be used to encrypt data at rest using a tenant secret key that is controlled by the customer.

QUESTION 79
NTO processes orders from its website via an order management system (OMS). The OMS stores over 2 million historical records and is currently not integrated with SF. The Sales team at NTO using Sales cloud and would like visibility into related customer orders yet they do not want to persist millions of records directly in Salesforce. NTO has asked the data architect to evaluate SF connect and the concept of data verification. Which 3 considerations are needed prior to a SF Connect implementation?
Choose 3 answers:

A. Create a 2nd system Admin user for authentication to the external source.
B. Develop an object relationship strategy.
C. Identify the external tables to sync into external objects
D. Assess whether the external data source is reachable via an ODATA endpoint.
E. Configure a middleware tool to poll external table data

Correct Answer: B, C, D
Section:
Explanation:
The three considerations needed prior to a SF Connect implementation are to develop an object relationship strategy, identify the external tables to sync into external objects, and assess whether the external data source is reachable via an ODATA endpoint. SF Connect is a feature that allows integrating external data sources with Salesforce using external objects. External objects are similar to custom objects, but they store metadata only and not data. They enable on-demand access to external data via standard Salesforce APIs and user interfaces. To implement SF Connect, a data architect needs to consider how the external objects will relate to other objects in Salesforce, which external tables will be exposed as external objects, and whether the external data source supports ODATA protocol for data access.

QUESTION 80
UC has been using SF for 10 years. Lately, users have noticed, that the pages load slowly when viewing Customer and Account list view.
To mitigate, UC will implement a data archive strategy to reduce the amount of data actively loaded.
Which 2 tasks are required to define the strategy? Choose 2 answers:

A. Identify the recovery point objective.
B. Identify how the archive data will be accessed and used.
C. Identify the recovery time objective.

D. Identify the data retention requirements

Correct Answer: B, D
Section:
Explanation:
The two tasks required to define the data archive strategy are to identify how the archive data will be accessed and used, and identify the data retention requirements. Data archiving is the process of moving infrequently used or historical data from active storage to a separate storage location for long-term retention. Data archiving can improve performance, reduce storage costs, and comply with legal or regulatory obligations. To define a data archive strategy, a data architect needs to consider how the archived data will be accessed and used by different users or processes in Salesforce or outside Salesforce, and how long the archived data needs to be retained based on business or legal requirements.

QUESTION 81
UC has a classic encryption for Custom fields and is leveraging weekly data reports for data backups. During the data validation of exported data UC discovered that encrypted field values are still being exported as part of data exported. What should a data architect recommend to make sure decrypted values are exported during data export?

A. Set a standard profile for Data Migration user, and assign view encrypted data
B. Create another field to copy data from encrypted field and use this field in export
C. Leverage Apex class to decrypt data before exporting it.
D. Set up a custom profile for data migration user and assign view encrypted data.

Correct Answer: B
Section:
Explanation:
The best solution to make sure decrypted values are exported during data export is to create another field to copy data from encrypted field and use this field in export. This is because classic encryption does not support exporting decrypted values of encrypted fields. The view encrypted data permission only allows users to view decrypted values in the user interface, but not in reports or data exports. Therefore, a workaround is to create a formula field or a workflow field update that copies the value of the encrypted field to another field, and use that field for data export. However, this solution has some drawbacks, such as exposing sensitive data in plain text and consuming extra storage space.A better solution would be to use Shield Platform Encryption, which supports exporting decrypted values of encrypted fields with the Export Encrypted Data permission

QUESTION 82
Universal containers is implementing Salesforce lead management. UC Procure lead data from multiple sources and would like to make sure lead data as company profile and location information. Which solution should a data architect recommend to make sure lead datahas both profile and location information? Option

A. Ask sales people to search for populating company profile and location data
B. Run reports to identify records which does not have company profile and location data
C. Leverage external data providers populate company profile and location data
D. Export data out of Salesforce and send to another team to populate company profile and location data

Correct Answer: C
Section:
Explanation:
The best solution to make sure lead data has both profile and location information is to leverage external data providers to populate company profile and location data. This is because external data providers can enrich lead data with additional information from third-party sources, such as Dun & Bradstreet, ZoomInfo, or Clearbit. This can help improve lead quality, segmentation, and conversion.Salesforce supports integrating with external data providers using Data.com Clean or other AppExchange solutions2. Asking sales people to search for populating company profile and location data is inefficient and prone to errors. Running reports to identify records which do not have company profile and location data is useful, but does not solve the problem of how to populate the missing data. Exporting data out of Salesforce and sending to another team to populate company profile and location data is cumbersome and time-consuming.

QUESTION 83
UC has millions of case records with case history and SLA dat
a. UC's compliance team would like historical cases to be accessible for 10 years for Audit purpose.
What solution should a data architect recommend?

A. Archive Case data using Salesforce Archiving process
B. Purchase more data storage to support case object
C. Use a custom object to store archived case data.
D. Use a custom Big object to store archived case data.

Correct Answer: D
Section:
Explanation:
The best solution to store historical cases for 10 years for audit purpose is to use a custom Big object to store archived case data. Big objects are a type of custom object that can store massive amounts of data on the Salesforce platform without affecting performance or storage limits. They also support point-and-click tools, triggers, and Apex code.Big objects can be used for archiving historical data that needs to be retained for compliance or analytics purposes3.Archiving case data using Salesforce Archiving process is not a good option because it only supports archiving cases that are closed for more than one year, and it does not allow customizing the archival criteria or accessing the archived data via Apex or APIs4. Purchasing more data storage to support case object is expensive and may impact performance. Using a custom object to store archived case data is not scalable and may consume a lot of storage space.

QUESTION 84
NTO need to extract 50 million records from a custom object everyday from its Salesforce org. NTO is facing query timeout issues while extracting these records.
What should a data architect recommend in order to get around the time out issue?

A. Use a custom auto number and formula field and use that to chunk records while extracting data.
B. The REST API to extract data as it automatically chunks records by 200.
C. Use ETL tool for extraction of records.
D. Ask SF support to increase the query timeout value.

Correct Answer: C
Section:
Explanation:
The best solution to extract 50 million records from a custom object everyday from Salesforce org without facing query timeout issues is to use an ETL tool for extraction of records. ETL stands for extract, transform, and load, and it refers to a process of moving data from one system to another. An ETL tool is a software application that can connect to various data sources, perform data transformations, and load data into a target destination.ETL tools can handle large volumes of data efficiently and reliably, and they often provide features such as scheduling, monitoring, error handling, and logging5. Using a custom auto number and formula field and use that to chunk records while extracting data is a possible workaround, but it requires creating additional fields and writing complex queries.The REST API can extract data as it automatically chunks records by 200, but it has some limitations, such as a maximum of 50 million records per query job6.Asking SF support to increase the query timeout value is not feasible because query timeout values are not configurable

QUESTION 85
UC has a variety of systems across its technology landscape, including SF, legacy enterprise resource planning (ERP) applications and homegrown CRM tools. UC has decided that they would like to consolidate all customer, opportunity and order data into Salesforce as part of its master data management (MDM) strategy.
What are the 3 key steps that a data architect should take when merging data from multiple systems into Salesforce? Choose 3 answers:

A. Create new fields to store additional values from all the systems.
B. Install a 3rd party AppExchange tool to handle the merger
C. Analyze each system's data model and perform gap analysis
D. Utilize an ETL tool to merge, transform and de-duplicate data.
E. Work with Stakeholders to define record and field survivorship rules

Correct Answer: C, D, E
Section:
Explanation:
The three key steps that a data architect should take when merging data from multiple systems into Salesforce are:

Analyze each system's data model and perform gap analysis. This step involves understanding the structure and meaning of the data in each system, identifying the common and unique data elements, and mapping the data fields between the systems. This step also involves assessing the quality and consistency of the data, and identifying any data cleansing or transformation needs.
Utilize an ETL tool to merge, transform, and de-duplicate data. This step involves using an ETL tool to connect to the source systems, extract the data, apply any data transformations or validations, and load the data into Salesforce. This step also involves applying de-duplication rules or algorithms to avoid creating duplicate records in Salesforce.
Work with stakeholders to define record and field survivorship rules. This step involves collaborating with the business users and owners of the data to determine which records and fields should be retained or overwritten in case of conflicts or discrepancies. This step also involves defining the criteria and logic for record and field survivorship, and implementing them in the ETL tool or in Salesforce.
Creating new fields to store additional values from all the systems is not a key step, but rather a possible outcome of the gap analysis. It may not be necessary or desirable to create new fields for every value from every system, as it may result in redundant or irrelevant data. Installing a 3rd party AppExchange tool to handle the merger is not a key step, but rather a possible option for choosing an ETL tool. It may not be the best option depending on the requirements, budget, and preferences of the organization.

QUESTION 86
UC has a roll-up summary field on Account to calculate the count of contacts associated with an account. During the account load, SF is throwing an ''Unable to lock a row'' error.
Which solution should a data architect recommend, to resolve the error?

A. Leverage data loader platform API to load data.
B. Perform Batch job in parallel mode and reduce Batch size
C. Perform Batch job in serial mode and reduce batch size
D. Defer roll-up summary fields calculation during data migration.

Correct Answer: C
Section:
Explanation:
The best solution to resolve the error of ''Unable to lock a row'' during the account load is to perform batch job in serial mode and reduce batch size. This is because roll-up summary fields are calculated synchronously when the parent record is updated, and asynchronously when the child record is updated. Therefore, updating many child records at once can cause locking issues on the parent record.To avoid this, it is recommended to use serial mode and smaller batch sizes when loading data using tools like Data Loader or Bulk API12. Leverage data loader platform API to load data is not a good option because it does not specify the mode or batch size. Perform batch job in parallel mode and reduce batch size is not a good option because parallel mode can still cause locking issues even with smaller batches.Defer roll-up summary fields calculation during data migration is not a good option because it is not possible to defer or disable roll-up summary fields calculation

QUESTION 87
UC is migrating individual customers (B2C) data from legacy systems to SF. There are millions of customers stored as accounts and contacts in legacy database.
Which object model should a data architect configure within SF?

A. Leverage person account object in Salesforce
B. Leverage custom person account object in SF
C. Leverage custom account and contact object in SF
D. Leverage standard account and contact object in SF

Correct Answer: A
Section:
Explanation:
The best object model to configure within SF for migrating individual customers (B2C) data from legacy systems is to leverage person account object in Salesforce.Person accounts are a special type of accounts that store information about individual people by combining certain account and contact fields into a single record4. Person accounts are useful for B2C scenarios where there is no need to associate a company name with a contact.Person accounts also support standard Salesforce features and functionality, such as leads, campaigns, reports, dashboards, etc5. Leverage custom person account object in SF is not a good option because there is no such thing as a custom person account object. Leverage custom account and contact object in SF is not a good option because it would require creating and maintaining additional objects and fields that may not be necessary or compatible with standard Salesforce features. Leverage standard account and contact object in SF is not a good option because it would require filling in dummy values for the account name field, which is mandatory for standard accounts.

QUESTION 88
NTO has 1 million customer records spanning 25 years. As part of its new SF project, NTO would like to create a master data management strategy to help preserve the history and relevance of its customer data.

Which 3 activities will be required to identify a successful master data management strategy? Choose 3 answers:

A. Identify data to be replicated
B. Create a data archive strategy
C. Define the systems of record for critical data
D. Install a data warehouse
E. Choose a Business Intelligence tool.

Correct Answer: A, B, C
Section:
Explanation:
The three activities that will be required to identify a successful master data management strategy are:
Identify data to be replicated: This activity involves determining which data elements need to be copied from one system to another, and how frequently the replication should occur. This can help ensure data consistency and availability across systems.
Create a data archive strategy: This activity involves defining how historical data will be stored, accessed, and deleted over time. This can help optimize data storage, performance, and compliance.
Define the systems of record for critical data: This activity involves identifying which system owns and maintains the authoritative version of each data element.This can help avoid data conflicts and duplication across systems67.
Install a data warehouse is not a required activity, but rather a possible option for consolidating data from multiple sources for analytics purposes. Choose a Business Intelligence tool is not a required activity, but rather a possible option for visualizing and reporting on data from various sources.

QUESTION 89
UC is migrating data from legacy system to SF. UC would like to preserve the following information on records being migrated:
Date time stamps for created date and last modified date.
Ownership of records belonging to inactive users being migrated to Salesforce.
Which 2 solutions should a data architect recommends to preserve the date timestamps and ownership on records? Choose 2 answers.

A. Log a case with SF to update these fields
B. Enable update records with Inactive Owners Permission
C. Enable Set Audit fields upon Record Creation Permission
D. Enable modify all and view all permission.

Correct Answer: B, C
Section:
Explanation:
The two solutions that a data architect should recommend to preserve the date timestamps and ownership on records being migrated are:
Enable update records with Inactive Owners Permission: This permission allows users to update record owner and sharing-based records with inactive owners.This can help preserve the original ownership of records that belong to users who are no longer active in Salesforce8.
Enable Set Audit fields upon Record Creation Permission: This permission allows users to set audit fields (such as Created By or Last Modified By) when they create a record via API importing tools like Data Loader.This can help preserve the original date timestamps of records that were created or modified in another system9.
Log a case with SF to update these fields is not a good solution because it is not necessary or feasible to ask Salesforce support to update these fields manually or programmatically.Enable modify all and view all permission is not a good solution because it does not affect the ability to preserve the date timestamps and ownership on records, but rather grants users access to all records regardless of sharing settings

QUESTION 90
UC has migrated its Back-office data into an on-premise database with REST API access. UC recently implemented Sales cloud for its sales organization. But users are complaining about a lack of order data inside SF.
UC is concerned about SF storage limits but would still like Sales cloud to have access to the data.
Which design patterns should a data architect select to satisfy the requirement?

A. Migrate and persist the data in SF to take advantage of native functionality.

B. Use SF Connect to virtualize the data in SF and avoid storage limits.
C. Develop a bidirectional integration between the on-premise system and Salesforce.
D. Build a UI for the on-premise system and iframe it in Salesforce

Correct Answer: B
Section:
Explanation:
The best design pattern to satisfy the requirement of accessing order data from an on-premise database with REST API access without consuming SF storage limits is to use SF Connect to virtualize the data in SF and avoid storage limits. SF Connect is an integration tool that allows users to access and integrate data from external sources using external objects. External objects are similar to custom objects, except that the data resides in another system and is accessed in real time via web service callouts.SF Connect supports various adapters to connect to different types of external data sources, such as OData, cross-org, or Apex custom adapter11. Migrate and persist the data in SF to take advantage of native functionality is not a good option because it would consume SF storage limits and require data synchronization between systems. Develop a bidirectional integration between the on-premise system and Salesforce is not a good option because it would be complex and costly to implement and maintain, and it would also consume SF storage limits. Build a UI for the on-premise system and iframe it in Salesforce is not a good option because it would not provide a seamless user experience and it would not allow users to search, report, or perform actions on the external data.

QUESTION 91
NTO has decided to franchise its brand. Upon implementation, 1000 franchisees will be able to access BTO's product information and track large customer sales and opportunities through a portal. The Franchisees will also be able to run monthly and quarterly sales reports and projections as well as view the reports in dashboards.
Which licenses does NTO need to provide these features to the Franchisees?

A. Salesforce Sales Cloud license
B. Lightning Platform license
C. Customer Community license
D. Partner Community license

Correct Answer: D
Section:
Explanation:
The best license to provide these features to the franchisees is the Partner Community license. Partner Community licenses are designed for external users who collaborate with your sales team on deals, such as resellers, distributors, or brokers. Partner Community users can access standard CRM objects, such as accounts, contacts, leads, opportunities, campaigns, and reports.They can also access custom objects and run dashboards12. Salesforce Sales Cloud license is not a good option because it is intended for internal users who need full access to standard CRM and custom apps. Lightning Platform license is not a good option because it is intended for users who need access to custom apps but not to standard CRM functionality.Customer Community license is not a good option because it is intended for external users who need access to customer support features, such as cases and knowledge articles, but not to sales features

QUESTION 92
Northern trail Outfitters (NTO) uses Sales Cloud and service Cloud to manage sales and support processes. Some of NTOs team are complaining they see new fields on their page unsure of which values need be input. NTO is concerned about lack of governance in making changes to Salesforce.
Which governance measure should a data architect recommend to solve this issue?

A. Add description fields to explain why the field is used, and mark the field as required.
B. Create and manage a data dictionary and ups a governance process for changes made to common objects.
C. Create reports to identify which users are leaving blank, and use external data sources o agreement the missing data.
D. Create validation rules with error messages to explain why the fields is used

Correct Answer: B
Section:
Explanation:
To solve the issue of lack of governance in making changes to Salesforce, a data architect should recommend creating and managing a data dictionary and setting up a governance process for changes made to common objects. A data dictionary is a document that defines the metadata, structure, and relationship of each object and field in Salesforce. A governance process is a set of rules and procedures that govern how changes are

proposed, reviewed, approved, and deployed in Salesforce. These measures will help NTO to maintain consistency, quality, and clarity of their data model and avoid confusion and errors among users. Option A is incorrect because adding description fields to explain why the field is used, and marking the field as required will not prevent unauthorized or unnecessary changes to Salesforce. Option C is incorrect because creating reports to identify which users are leaving blank, and using external data sources to augment the missing data will not address the root cause of the issue, which is the lack of governance in making changes to Salesforce. Option D is incorrect because creating validation rules with error messages to explain why the fields are used will not stop users from seeing new fields on their page that they are unsure of.

QUESTION 93
Universal Containers (UC) has adopted Salesforce as its primary sales automated tool. UC has 100,00 customers with a growth rate of 10% a year, UC uses an on-premise web-based billing and invoice system that generates over 1 million invoices a year supporting a monthly billing cycle.
The UC sales team needs to be able to pull a customer record and view their account status, Invoice history, and opportunities without navigating outside of Salesforce.
What should a data architect use to provide the sales team with the required functionality?

A. Create a custom object and migrate the last 12 months of Invoice data into Salesforce so it can be displayed on the Account layout.
B. Write an Apex callout and populate a related list to display on the account record.
C. Create a mashup page that will present the billing system records within Salesforce.
D. Create a visual force tab with the billing system encapsulated within an iframe.

Correct Answer: C
Section:
Explanation:
To provide the sales team with the required functionality, a data architect should use a mashup page that will present the billing system records within Salesforce. A mashup page is a web page that combines data from multiple sources into a single integrated view. A mashup page can be created using Visualforce or Lightning Web Components, and can use Salesforce Connect or custom integrations to access external data from the on-premise web-based billing and invoice system. This will allow the sales team to pull a customer record and view their account status, invoice history, and opportunities without navigating outside of Salesforce. Option A is incorrect because creating a custom object and migrating the last 12 months of invoice data into Salesforce so it can be displayed on the account layout will consume a lot of storage space and may not reflect the latest data from the billing system. Option B is incorrect because writing an Apex callout and populating a related list to display on the account record will require additional development effort and may not be scalable or performant for large volumes of data. Option D is incorrect because creating a visual force tab with the billing system encapsulated within an iframe will not allow the sales team to view the billing system records within the customer record, but rather in a separate tab.

QUESTION 94
A large retail B2C customer wants to build a 360 view of its customer for its call center agents. The customer interaction is currently maintained in the following system: 1. Salesforce CRM
3. Customer Master Data management (MDM)
4. Contract Management system
5. Marketing solution
What should a data architect recommend that would help upgrade uniquely identify customer across multiple systems:

A. Store the salesforce id in all the solutions to identify the customer.
B. Create a custom object that will serve as a cross reference for the customer id.
C. Create a customer data base and use this id in all systems.
D. Create a custom field as external id to maintain the customer Id from the MDM solution.

Correct Answer: D
Section:
Explanation:
To help uniquely identify customer across multiple systems, a data architect should recommend creating a custom field as external ID to maintain the customer ID from the MDM solution. An external ID is a custom field that has the ''External ID'' attribute enabled, which means that it contains unique record identifiers from a system outside of Salesforce. By using the customer ID from the MDM solution as an external ID in Salesforce CRM, Contract Management system, and Marketing solution, the data architect can ensure that each customer can be easily identified and integrated across these systems. Option A is incorrect because storing the Salesforce ID in all the solutions to identify the customer will not work if the customer records are created or updated in other systems besides Salesforce CRM. Option B is incorrect because creating a custom object that will serve as a cross reference for the customer ID will require additional configuration effort and may not be consistent with the actual customer records in each system. Option C is incorrect because creating a customer database and using this ID in all systems will require additional infrastructure cost and maintenance effort.

QUESTION 95
Universal Containers has a requirement to store more than 100 million records in salesforce and needs to create a custom big object to support this business requirement.
Which two tools should a data architect use to build custom object?

A. Use DX to create big object.
B. Use Metadata API to create big object.
C. Go to Big Object in setup select new to create big object.
D. Go to Object Manager in setup and select new to create big object.

Correct Answer: B, C
Section:
Explanation:
To build a custom big object to support storing more than 100 million records in Salesforce, a data architect should use Metadata API or Big Object in setup. Metadata API is an API that allows developers to create, retrieve, update, or delete metadata components in Salesforce programmatically. Big Object in setup is a user interface that allows admins to create big objects declaratively without writing code. Both tools can be used to define custom big objects and their fields, indexes, and relationships in Salesforce. Option A is incorrect because DX (Developer Experience) is a set of tools that allows developers to create and manage applications on Salesforce Platform, but it does not support creating big objects directly. Option D is incorrect because Object Manager in setup is a user interface that allows admins to create and manage standard and custom objects in Salesforce, but it does not support creating big objects declaratively.

QUESTION 96
NTO uses salesforce to manage relationships and track sales opportunities. It has 10 million customers and 100 million opportunities. The CEO has been complaining 10 minutes to run and sometimes failed to load, throwing a time out error.
Which 3 options should help improve the dashboard performance?
Choose 3 answers:

A. Use selective queries to reduce the amount of data being returned.
B. De-normalize the data by reducing the number of joins.
C. Remove widgets from the dashboard to reduce the number of graphics loaded.
D. Run the dashboard for CEO and send it via email.
E. Reduce the amount of data queried by archiving unused opportunity records.

Correct Answer: A, B, E
Section:
Explanation:
To improve the dashboard performance, the data architect should use selective queries to reduce the amount of data being returned, de-normalize the data by reducing the number of joins, and reduce the amount of data queried by archiving unused opportunity records. These options will help optimize the query performance, reduce the query complexity, and free up storage space. Option C is incorrect because removing widgets from the dashboard to reduce the number of graphics loaded will not affect the dashboard performance significantly, and may reduce the usability and functionality of the dashboard. Option D is incorrect because running the dashboard for CEO and sending it via email will not improve the dashboard performance, but rather shift the burden to another user.

QUESTION 97
Northern Trail Outfitters (NTO) wants to implement backup and restore for Salesforce data, Currently, it has data backup processes that runs weekly, which back up all Salesforce data to an enterprise data warehouse (EDW). NTO wants to move to daily backups and provide restore capability to avoid any data loss in case of outage.
What should a data architect recommend for a daily backup and restore solution?

A. Use AppExchange package for backup and restore.
B. Use ETL for backup and restore from EDW.
C. Use Bulk API to extract data on daily basis to EDW and REST API for restore.
D. Change weekly backup process to daily backup, and implement a custom restore solution.

Correct Answer: A
Section:
Explanation:
The data architect should recommend using AppExchange package for backup and restore. AppExchange is a marketplace for Salesforce apps and solutions that can be installed and configured in Salesforce orgs. There are several AppExchange packages that provide backup and restore functionality for Salesforce data, such as OwnBackup, Odaseva, or Spanning. These packages can perform daily backups of Salesforce data to a secure cloud storage, and provide restore capability to avoid any data loss in case of outage. Option B is incorrect because using ETL (Extract, Transform, Load) for backup and restore from EDW (Enterprise Data Warehouse) will require additional development effort and may not be reliable or secure. Option C is incorrect because using Bulk API to extract data on daily basis to EDW and REST API for restore will require additional integration effort and may not be scalable or performant. Option D is incorrect because changing weekly backup process to daily backup, and implementing a custom restore solution will require additional configuration effort and may not be robust or compliant.

QUESTION 98
A large Automobile company has implemented SF for its Sales Associates. Leads flow from its website to SF using a batch integration in SF. The Batch job connects the leads to Accounts in SF. Customer visiting their retail stores are also created in SF as Accounts.
The company has noticed a large number of duplicate accounts in SF. On analysis, it was found that certain customers could interact with its website and also visit the store. The Sales associates use Global Search to search for customers in Salesforce before they create the customers.
Which scalable option should a data Architect choose to implement to avoid duplicates?

A. Create duplicate rules in SF to validate duplicates during the account creation process
B. Implement a MDM solution to validate the customer information before creating Accounts in SF.
C. Build Custom search based on fields on Accounts which can be matched with customer when they visit the store
D. Customize Account creation process to search if customer exists before creating an Account.

Correct Answer: A
Section:
Explanation:
The data architect should choose to implement duplicate rules in SF (Salesforce) to validate duplicates during the account creation process. Duplicate rules are a feature in Salesforce that allow users to define criteria and actions for detecting and preventing duplicate records. By creating duplicate rules for accounts, the data architect can ensure that any leads from the website or customers from the retail stores that match existing accounts in Salesforce are flagged or blocked before they are created as new accounts. This will help avoid duplicate accounts in Salesforce and maintain data quality. Option B is incorrect because implementing a MDM (Master Data Management) solution to validate the customer information before creating accounts in SF will require additional infrastructure cost and maintenance effort. Option C is incorrect because building custom search based on fields on accounts which can be matched with customer when they visit the store will require additional development effort and may not be accurate or user-friendly. Option D is incorrect because customizing account creation process to search if customer exists before creating an account will require additional configuration effort and may not be consistent or scalable.

QUESTION 99
Universal Containers (UC) has implemented Salesforce, UC is running out of storage and needs to have an archiving solution, UC would like to maintain two years of data in Saleforce and archive older data out of Salesforce.
Which solution should a data architect recommend as an archiving solution?

A. Use a third-party backup solution to backup all data off platform.
B. Build a batch join move all records off platform, and delete all records from Salesforce.
C. Build a batch join to move two-year-old records off platform, and delete records from Salesforce.
D. Build a batch job to move all restore off platform, and delete old records from Salesforce.

Correct Answer: C
Section:
Explanation:
The data architect should recommend building a batch job to move two-year-old records off platform, and delete records from Salesforce as an archiving solution. A batch job is a process that runs in the background and performs operations on large volumes of data in Salesforce. By building a batch job that moves two-year-old records off platform to an external storage system, such as Amazon S3 or Google Cloud Storage, and deletes them from Salesforce, the data architect can reduce the storage consumption and improve the performance of Salesforce org. Option A is incorrect because using a third-party backup solution to backup all data off platform will not free up any storage space in Salesforce, unless the data is also deleted from Salesforce after backup. Option B is incorrect because building a batch job to move all records off platform, and delete all records from Salesforce will result in losing all the current data in Salesforce, which may not be desirable or feasible. Option D is incorrect because building a batch job to move all restore off platform, and delete old records from Salesforce does not

make sense, as restore implies restoring data back to Salesforce, not moving it off platform.

QUESTION 100
Northern trail Outfitters (NTO) runs its entire out of an enterprise data warehouse (EDW), NTD's sales team starting to use Salesforce after a recent implementation, but currently lacks data required to advanced and opportunity to the next stage.
NTO's management has research Salesforce Connect and would like to use It to virtualize and report on data from the EDW within Salesforce. NTO will be running thousands of reports per day across 10 to 15 external objects.
What should a data architect consider before implementing Salesforce Connect for reporting?

A. Maximum number for records returned
B. OData callout limits per day
C. Maximum page size for server-driven paging
D. Maximum external objects per org

Correct Answer: B
Section:
Explanation:
According to theSalesforce Connect Reportingblog post, one of the considerations for using Salesforce Connect for reporting is the OData callout limits per day. The blog post states that ''Salesforce Connect has a limit of 100,000 callouts per day. This limit is shared across all external data sources in your org. If you exceed this limit, you will receive an error message and no more callouts will be allowed until the next day.'' Therefore, a data architect should consider this limit before implementing Salesforce Connect for reporting.

QUESTION 101
Northern Trail Outfitters (NTO) has the following systems:
Customer master-source of truth for customer information
Service cloud-customer support
Marketing cloud-marketing support
Enterprise data warehouse---business reporting
The customer data is duplicated across all these systems and are not kept in sync. Customers are also complaining that they get repeated marketing emails and have to call into update their information.
NTO is planning to implement master data management (MDM) solution across the enterprise.
Which three data will an MDM tool solve?
Choose 3 answers

A. Data completeness
B. Data loss and recovery
C. Data duplication
D. Data accuracy and quality
E. Data standardization

Correct Answer: C, D, E
Section:
Explanation:
According to theWhat is Master Data Management (MDM)?article, some of the data challenges that an MDM tool can solve are data duplication, data accuracy and quality, and data standardization. The article states that ''MDM solutions comprise a broad range of data cleansing, transformation, and integration practices. As data sources are added to the system, MDM initiates processes to identify, collect, transform, and repair data. Once the data meets the quality thresholds, schemas and taxonomies are created to help maintain a high-quality master reference.'' Therefore, an MDM tool can help NTO eliminate data duplication across different systems, improve data accuracy and quality by removing errors and inconsistencies, and standardize data formats and definitions for better integration and analysis.

QUESTION 102
NTO has a loyalty program to reward repeat customers. The following conditions exists:
1. Reward levels are earned based on the amount spent during the previous 12 months.

2. The program will track every item a customer has bought and grant them points for discount.
3. The program generates 100 million records each month.
NTO customer support would like to see a summary of a customer's recent transaction and reward level(s) they have attained.
Which solution should the data architect use to provide the information within the salesforce for the customer support agents?

A. Create a custom object in salesforce to capture and store all reward program. Populate nightly from the point-of-scale system, and present on the customer record.
B. Capture the reward program data in an external data store and present the 12 months trailing summary in salesforce using salesforce connect and then external object.
C. Provide a button so that the agent can quickly open the point of sales system displaying the customer history.
D. Create a custom big object to capture the reward program data and display it on the contact record and update nightly from the point-of-scale system.

Correct Answer: D
Section:
Explanation:
According to theGet Started with Big Objectsunit on Trailhead, one of the use cases for custom big objects is to store and manage loyalty program data for customers. The unit states that ''From loyalty programs to transactions, order, and billing information, use a custom big object to keep track of every detail.'' Therefore, a custom big object can be used to capture the reward program data and display it on the contact record. Additionally, according to theBig Objects Implementation Guide, big objects can handle massive amounts of data (up to billions of records) and can be updated nightly from external systems using Bulk API or batch Apex. Therefore, a custom big object can meet the requirements of NTO's loyalty program scenario.

QUESTION 103
Universal Container (US) is replacing a home-grown CRM solution with Salesforce, UC has decided to migrate operational (Open and active) records to Salesforce, while keeping historical records in legacy system, UC would like historical records to be available in Salesforce on an as needed basis.
Which solution should a data architect recommend to meet business requirement?

A. Leverage real-time integration to pull records into Salesforce.
B. Bring all data Salesforce, and delete it after a year.
C. Leverage mashup to display historical records in Salesforce.
D. Build a chair solution to go the legacy system and display records.

Correct Answer: C
Section:
Explanation:
According to theUsing Mashupsarticle on Salesforce Developers, one of the techniques for deploying large data volumes is to use mashups to display historical records in Salesforce. The article states that ''Mashups are a way to display data from an external system within a Salesforce page without copying or synchronizing the data. Mashups use a combination of Visualforce, Apex callouts, and JavaScript code that runs in the browser. Mashups are useful when you want to display large amounts of read-only data that is stored outside of Salesforce.'' Therefore, a data architect should recommend this solution to meet the business requirement of UC.

QUESTION 104
UC has a legacy client server app that as a relational data base that needs to be migrated to salesforce.
What are the 3 key actions that should be done when data modeling in salesforce?
Choose 3 answers:

A. Identify data elements to be persisted in salesforce.
B. Map legacy data to salesforce objects.
C. Map legacy data to salesforce custom objects.
D. Work with legacy application owner to analysis legacy data model.
E. Implement legacy data model within salesforce using custom fields.

Correct Answer: A, B, E
Section:

Explanation:
According to theData Modelingunit on Trailhead, some of the key actions that should be done when data modeling in Salesforce are identifying data elements, mapping legacy data, and implementing legacy data model. The unit states that ''Before you start creating objects and fields in Salesforce, you need to identify the data elements that you want to store and work with. ... Next, you need to map your legacy data to Salesforce objects and fields. ... Finally, you need to implement your data model in Salesforce by creating custom objects and fields using declarative tools or Metadata API.'' Therefore, these are the correct actions for migrating a legacy client server app to Salesforce.

QUESTION 105
A custom pricing engine for a Salesforce customer has to be decided by factors with the following hierarchy:
State in which the customer is located
City in which the customer is located if available
Zip code in which the customer is located if available
Changes to this information should have minimum code changes
What should a data architect recommend to maintain this information for the custom pricing engine that is to be built in Salesforce?

A. Create a custom object to maintain the pricing criteria.
B. Assign the pricing criteria within customer pricing engine.
C. Maintain require pricing criteria in custom metadata types.
D. Configure the pricing criteria in price books.

Correct Answer: C
Section:
Explanation:
According to theGet Started with Custom Metadata Typesunit on Trailhead, one of the use cases for custom metadata types is to define custom charges for an accounting app. The unit states that ''Say that your org uses a standard accounting app. You can create a custom metadata type that defines custom charges, like duties and VAT rates. Then you can write some Apex code that calculates the total amount due for each invoice by using the metadata from your custom metadata type.'' Therefore, a similar approach can be used to maintain the pricing criteria for a custom pricing engine in Salesforce.

QUESTION 106
Universal Containers (UC) owns a complex Salesforce org with many Apex classes, triggers, and automated processes that will modify records if available. UC has identified that, in its current development state, UC runs change of encountering race condition on the same record.
What should a data architect recommend to guarantee that records are not being updated at the same time?

A. Embed the keywords FOR UPDATE after SOQL statements.
B. Disable classes or triggers that have the potential to obtain the same record.
C. Migrate programmatic logic to processes and flows.
D. Refactor or optimize classes and trigger for maximum CPU performance.

Correct Answer: A
Section:
Explanation:
According to theHow to avoid row lock or race condition in Apexblog post, one of the ways to prevent race condition in Apex is to use the FOR UPDATE keyword in SOQL statements. The blog post states that ''We need to lock the records on which we are working such that other batches or threads will not be having any effect on them. How can we lock a record, then? We need to make use of FOR UPDATE keyword in the SOQL query.'' Therefore, a data architect should recommend this solution to guarantee that records are not being updated at the same time by different processes.

QUESTION 107
UC migrating 100,000 Accounts from an enterprise resource planning (ERP) to salesforce and is concerned about ownership skew and performance.
Which 3 recommendations should a data architect provide to prevent ownership skew?
Choose 3 answers:

A. Assigned a default user as owner of accounts, and assign role in hierarchy.
B. Keep users out of public groups that can be used as the source for sharing rules.
C. Assign a default user as owner of account and do not assign any role to default user.
D. Assign ''view all'' permission on profile to give access to account.
E. Assign a default user as owner of accounts and assigned top most role in hierarchy.

Correct Answer: B, C, E
Section:
Explanation:
According to the Salesforce documentation1, ownership skew occurs when a large number of records (more than 10,000) are owned by a single user or queue. This can cause performance issues and lock contention when multiple users try to access or update those records. To prevent ownership skew, some of the recommended practices are:
Assign a default user as the owner of the records and do not assign any role to the default user (option C). This way, the records will not be visible to other users in the role hierarchy and will not cause sharing recalculations.
Keep users out of public groups that can be used as the source for sharing rules (option B). Sharing rules based on public groups can cause excessive sharing calculations and lock contention when many records are owned by a single user or queue.
Assign a default user as the owner of the records and assign the top most role in the hierarchy to the default user (option E). This way, the records will be visible to all users in the role hierarchy, but will not cause sharing recalculations or lock contention.
Assigning a default user as the owner of the records and assigning a role in the hierarchy (option A) is not a good practice, as it can cause sharing recalculations and lock contention when the role is updated or moved.
Assigning ''view all'' permission on profile to give access to the records (option D) is also not a good practice, as it can bypass the security and sharing model and expose sensitive data to unauthorized users.

QUESTION 108
Universal Containers has a rollup summary field on account to calculate the number of contacts associated with an account. During the account load, Salesforce is throwing an 'UNABLE _TO_LOCK_ROW' error.
Which solution should a data architect recommend to resolve the error?

A. Defer rollup summary field calculation during data migration.
B. Perform a batch job in serial mode and reduce the batch size.
C. Perform a batch job in parallel mode and reduce the batch size.
D. Leverage Data Loader's platform API to load data.

Correct Answer: B
Section:
Explanation:
According to the Salesforce documentation1, the ''UNABLE _TO_LOCK_ROW'' error occurs when a record is being updated or created, and another operation tries to access or update the same record at the same time. This can cause lock contention and timeout issues. To resolve the error, some of the recommended solutions are:
Perform a batch job in serial mode and reduce the batch size (option B). This means running the batch job one at a time and processing fewer records per batch. This can reduce the chances of concurrent updates and lock contention on the same records.
Use the FOR UPDATE keyword to lock records in Apex code or API calls. This means explicitly locking the records that are being accessed or updated by a transaction, and preventing other transactions from modifying them until the lock is released.This can avoid conflicts and errors between concurrent operations on the same records2.
Defer rollup summary field calculation during data migration (option A). This means disabling the automatic calculation of rollup summary fields on the parent object when child records are inserted or updated. This can improve performance and avoid locking issues on the parent records.However, this option is only available for custom objects, not standard objects3.
Performing a batch job in parallel mode and reducing the batch size (option C) is not a good solution, as it can still cause lock contention and errors if multiple batches try to access or update the same records at the same time. Leveraging Data Loader's platform API to load data (option D) is also not a good solution, as it can still encounter locking issues if other operations are modifying the same records at the same time.

QUESTION 109
A large insurance provider is looking to implement Salesforce. The following exist.
1. Multiple channel for lead acquisition
2. Duplication leads across channels
3. Poor customer experience and higher costs
On analysis, it found that there are duplicate leads that are resulting to mitigate the issues?

A. Build process is manually search and merge duplicates.
B. Standard lead information across all channels.
C. Build a custom solution to identify and merge duplicate leads.
D. Implement third-party solution to clean and event lead data.
E. Implement de-duplication strategy to prevent duplicate leads

Correct Answer: B, D, E
Section:
Explanation:
According to the Salesforce documentation2, duplicate leads are leads that have the same or similar information as other leads in Salesforce. Duplicate leads can cause poor customer experience, higher costs, and inaccurate reporting. To mitigate the issues caused by duplicate leads, some of the recommended practices are:
Standardize lead information across all channels (option B). This means using consistent formats, values, and fields for capturing lead data from different sources, such as web forms, email campaigns, or third-party vendors. This can help reduce data quality issues and make it easier to identify and prevent duplicate leads.
Implement a third-party solution to clean and enrich lead data (option D). This means using an external service or tool that can validate, correct, update, and enhance lead data before or after importing it into Salesforce. This can help improve data quality and accuracy, and reduce duplicate leads.
Implement a de-duplication strategy to prevent duplicate leads (option E). This means using Salesforce features or custom solutions that can detect and block duplicate leads from being created or imported into Salesforce.For example, using Data.com Duplicate Management3, which allows defining matching rules and duplicate rules for leads and other objects.
Building a process to manually search and merge duplicates (option A) is not a good practice, as it can be time-consuming, error-prone, and inefficient. Building a custom solution to identify and merge duplicate leads (option C) is also not a good practice, as it can be complex, costly, and difficult to maintain. It is better to use existing Salesforce features or third-party solutions that can handle duplicate leads more effectively.

QUESTION 110
Universal Containers (UC) uses the following Salesforce products:
Sales Cloud for customer management.
Marketing Cloud for marketing.
Einstein Analytics for business reporting.
UC occasionally gets a list of prospects from third-party source as comma-separated values (CSV) files for marketing purposes. Historically, UC would load contact Lead object in Salesforce and sync to Marketing Cloud to send marketing communications. The number of records in the Lead object has grown over time and has been consuming large amounts of storage in Sales Cloud, UC is looking for recommendations to reduce the storage and advice on how to optimize the marketing Cloud to send marketing communications. The number of records in the Lead object has grown over time and has been consuming large amounts of storage in Sales Cloud, UC is looking for recommendations to reduce the storage and advice on how to optimize the marketing process.
What should a data architect recommend to UC in order to immediately avoid storage issues in the future?

A. Load the CSV files in Einstein Analytics and sync with Marketing Cloud prior to sending marketing communications;
B. Load the CSV files in an external database and sync with Marketing Cloud prior to sending marketing communications.
C. Load the contacts directly to Marketing Cloud and have a reconciliation process to track prospects that are converted to customers.
D. Continue to use the existing process to use Lead object to sync with Marketing Cloud and delete Lead records from Sales after the sync is complete.

Correct Answer: C
Section:
Explanation:
According to the Salesforce documentation4, Marketing Cloud is a platform that allows creating and managing marketing campaigns across multiple channels, such as email, mobile, social media, web, etc. Marketing Cloud can integrate with Sales Cloud and other Salesforce products to share data and insights.One of the ways to integrate Marketing Cloud with Sales Cloud is using Marketing Cloud Connect5, which allows syncing data between the two platforms using synchronized data sources.
However, if UC occasionally gets a list of prospects from third-party sources as CSV files for marketing purposes, it may not be necessary or efficient to load them into Sales Cloud first and then sync them with Marketing Cloud.This can consume large amounts of storage in Sales Cloud, which has a limit based on the license type6. It can also cause data quality issues, such as duplicates or outdated information.
A better option for UC is to load the contacts directly to Marketing Cloud using Import Definition, which allows importing data from external files or databases into Marketing Cloud data extensions. Data extensions are custom tables that store marketing data in Marketing Cloud. This way, UC can avoid storage issues in Sales Cloud and optimize the marketing process by sending marketing communications directly from Marketing Cloud.
To track prospects that are converted to customers, UC can have a reconciliation process that compares the contacts in Marketing Cloud with the accounts or contacts in Sales Cloud. This can be done using SQL queries or API

calls to access and compare data from both platforms. Alternatively, UC can use Marketing Cloud Connect to sync the converted contacts from Sales Cloud to Marketing Cloud using synchronized data sources.
Loading the CSV files in Einstein Analytics and syncing with Marketing Cloud prior to sending marketing communications (option A) is not a good option, as it can add unnecessary complexity and latency to the process. Einstein Analytics is a platform that allows creating and analyzing data using interactive dashboards and reports. It is not designed for importing and exporting data for marketing purposes.
Loading the CSV files in an external database and syncing with Marketing Cloud prior to sending marketing communications (option B) is also not a good option, as it can incur additional costs and maintenance for the external database. It can also introduce data security and privacy risks, as the data may not be encrypted or protected by Salesforce.
Continuing to use the existing process to use Lead object to sync with Marketing Cloud and delete Lead records from Sales after the sync is complete (option D) is not a good option, as it can cause performance issues and data loss. Deleting Lead records from Sales can affect reporting and auditing, as well as trigger workflows and validations that may not be intended. It can also cause data inconsistency and synchronization errors between Sales Cloud and Marketing Cloud.

QUESTION 111
Universal Containers (UC) is migrating from a legacy system to Salesforce CRM, UC is concerned about the quality of data being entered by users and through external integrations.
Which two solutions should a data architect recommend to mitigate data quality issues?

A. Leverage picklist and lookup fields where possible
B. Leverage Apex to validate the format of data being entered via a mobile device.
C. Leverage validation rules and workflows.
D. Leverage third-party- AppExchange tools

Correct Answer: A, C
Section:
Explanation:
According to the Salesforce documentation1, data quality is the measure of how well the data in Salesforce meets the expectations and requirements of the users and stakeholders. Data quality can be affected by various factors, such as data entry errors, data duplication, data inconsistency, data incompleteness, data timeliness, etc. To mitigate data quality issues, some of the recommended solutions are:
Leverage picklist and lookup fields where possible (option A). This means using fields that restrict the values or references that can be entered by the users or integrations. This can help reduce data entry errors, enforce data consistency, and improve data accuracy.
Leverage validation rules and workflows (option C). This means using features that allow defining rules and criteria to validate the data that is entered or updated by the users or integrations. This can help prevent invalid or incorrect data from being saved, and trigger actions or alerts to correct or improve the data.
Leveraging Apex to validate the format of data being entered via a mobile device (option B) is not a good solution, as it can be complex, costly, and difficult to maintain. It is better to use standard features or declarative tools that can handle data validation more effectively. Leveraging third-party AppExchange tools (option D) is also not a good solution, as it can incur additional costs and dependencies. It is better to use native Salesforce features or custom solutions that can handle data quality more efficiently.

QUESTION 112
Universal Containers (CU) is in the process of implementing an enterprise data warehouse (EDW). UC needs to extract 100 million records from Salesforce for migration to the EDW.
What data extraction strategy should a data architect use for maximum performance?

A. Install a third-party AppExchange tool.
B. Call the REST API in successive queries.
C. Utilize PK Chunking with the Bulk API.
D. Use the Bulk API in parallel mode.

Correct Answer: C
Section:
Explanation:
According to the Salesforce documentation2, extracting large amounts of data from Salesforce can be challenging and time-consuming, as it can encounter performance issues, API limits, timeouts, etc. To extract 100 million records from Salesforce for migration to an enterprise data warehouse (EDW), a data extraction strategy that can provide maximum performance is:
Utilize PK Chunking with the Bulk API (option C). This means using a feature that allows splitting a large query into smaller batches based on the record IDs (primary keys) of the queried object.This can improve performance and avoid timeouts by processing each batch asynchronously and in parallel using the Bulk API3.
Installing a third-party AppExchange tool (option A) is not a good solution, as it can incur additional costs and dependencies. It may also not be able to handle such a large volume of data efficiently. Calling the REST API in successive queries (option B) is also not a good solution, as it can encounter API limits and performance issues when querying such a large volume of data. Using the Bulk API in parallel mode (option D) is also not a good

solution, as it can still cause timeouts and errors when querying such a large volume of data without chunking.

QUESTION 113
A large multinational B2C Salesforce customer is looking to implement their distributor management application is Salesforce. the application has the following capabilities:
1.Distributor create sales order in salesforce
2.Sales order are based on product prices applicable to their region
3. Sales order are closed once they are fulfilled
4. It is decided to maintain the order in opportunity object
How should the data architect model this requirement?

A. Create lookup to Custom Price object and share with distributors.
B. Configure price Books for each region and share with distributors.
C. Manually update Opportunities with Prices application to distributors.
D. Add custom fields in Opportunity and use triggers to update prices.

Correct Answer: B
Section:
Explanation:
According to the Salesforce documentation, an opportunity is a standard object that represents a potential sale or deal with an account or contact. An opportunity can have products and prices associated with it using price books. A price book is a standard object that contains a list of products and their prices for different regions, currencies, segments, etc. A price book can be shared with different users or groups based on their visibility and access settings.
To model the requirement of implementing a distributor management application in Salesforce, where distributors create sales orders based on product prices applicable to their region, and sales orders are closed once they are fulfilled, a data architect should:
Configure price books for each region and share with distributors (option B). This means creating different price books for different regions with the appropriate products and prices, and sharing them with the distributors who belong to those regions. This way, distributors can create sales orders (opportunities) using the price books that are relevant to their region.
Creating a lookup to Custom Price object and sharing with distributors (option A) is not a good solution, as it can introduce unnecessary complexity and redundancy to the data model. It is better to use standard objects and features that are designed for managing products and prices in Salesforce. Manually updating opportunities with prices applicable to distributors (option C) is also not a good solution, as it can be time-consuming, error-prone, and inefficient. It is better to use automation tools or features that can update prices based on predefined criteria or logic. Adding custom fields in opportunity and using triggers to update prices (option D) is also not a good solution, as it can be complex, costly, and difficult to maintain. It is better to use standard fields and features that can handle prices more effectively.

QUESTION 114
North Trail Outfitters (NTO) operates a majority of its business from a central Salesforce org, NTO also owns several secondary orgs that the service, finance, and marketing teams work out of, At the moment, there is no integration between central and secondary orgs, leading to data-visibility issues.
Moving forward, NTO has identified that a hub-and-spoke model is the proper architect to manage its data, where the central org is the hub and the secondary orgs are the spokes.
Which tool should a data architect use to orchestrate data between the hub org and spoke orgs?

A. A middleware solution that extracts and distributes data across both the hub and spokes.
B. Develop custom APIs to poll the hub org for change data and push into the spoke orgs.
C. Develop custom APIs to poll the spoke for change data and push into the org.
D. A backup and archive solution that extracts and restores data across orgs.

Correct Answer: A
Section:
Explanation:
According to the Salesforce documentation, a hub-and-spoke model is an integration architecture pattern that allows connecting multiple Salesforce orgs using a central org (hub) and one or more secondary orgs (spokes). The hub org acts as the master data source and orchestrates the data flow between the spoke orgs. The spoke orgs act as the consumers or producers of the data and communicate with the hub org.
To orchestrate data between the hub org and spoke orgs, a data architect should use:
A middleware solution that extracts and distributes data across both the hub and spokes (option A). This means using an external service or tool that can connect to multiple Salesforce orgs using APIs or connectors, and perform data extraction, transformation, and distribution operations between the hub and spoke orgs. This can provide a scalable, flexible, and reliable way to orchestrate data across multiple orgs.

Developing custom APIs to poll the hub org for change data and push into the spoke orgs (option B) is not a good solution, as it can be complex, costly, and difficult to maintain. It may also not be able to handle large volumes of data or complex transformations efficiently. Developing custom APIs to poll the spoke orgs for change data and push into the hub org (option C) is also not a good solution, as it can have the same drawbacks as option B. It may also not be able to handle conflicts or errors effectively. Using a backup and archive solution that extracts and restores data across orgs (option D) is also not a good solution, as it can incur additional costs and dependencies. It may also not be able to handle real-time or near-real-time data orchestration requirements.

QUESTION 115
Universal Containers has 30 million case records. The Case object has 80 fields. Agents are reporting reports in the Salesforce org.
Which solution should a data architect recommend to improve reporting performance?

A. Create a custom object to store aggregate data and run reports.
B. Contact Salesforce support to enable skinny table for cases.
C. Move data off of the platform and run reporting outside Salesforce, and give access to reports.
D. Build reports using custom Lightning components.

Correct Answer: C
Section:
Explanation:
According to the Salesforce documentation1, reporting performance can be affected by various factors, such as the volume and complexity of data, the design and configuration of reports and dashboards, the number and type of users accessing the reports, etc. To improve reporting performance, some of the recommended solutions are:
Move data off of the platform and run reporting outside Salesforce, and give access to reports (option C). This means using an external service or tool that can extract, transform, and load (ETL) data from Salesforce to another system or database, such as a data warehouse or a business intelligence platform. This can improve reporting performance by reducing the load and latency on Salesforce, and enabling faster and more flexible reporting and analysis on the external system. Users can access the reports from the external system using a link or an embedded component in Salesforce.
Contact Salesforce support to enable skinny table for cases (option B). This means requesting Salesforce to create a custom table that contains a subset of fields from the Case object that are frequently used or queried.A skinny table can improve reporting performance by avoiding joins between standard and custom fields, omitting soft-deleted records, and leveraging indexes on the fields2.
Create a custom object to store aggregate data and run reports (option A). This means creating a custom object that contains summary or calculated data from the Case object, such as counts, sums, averages, etc. A custom object can improve reporting performance by reducing the number of records and fields that need to be queried and displayed.
Build reports using custom Lightning components (option D). This means creating custom components that use Lightning Web Components or Aura Components frameworks to display report data in Salesforce. A custom component can improve reporting performance by using client-side caching, pagination, lazy loading, or other techniques to optimize data rendering and interaction.

QUESTION 116
UC developers have created a new lightning component that uses an Apex controller using a SOQL query to populate a custom list view. Users are complaining that the component often fails to load and returns a time-out error.
What tool should a data architect use to identify why the query is taking too long?

A. Use Splunk to query the system logs looking for transaction time and CPU usage.
B. Enable and use the query plan tool in the developer console.
C. Use salesforce's query optimizer to analyze the query in the developer console.
D. Open a ticket with salesforce support to retrieve transaction logs to e analyzed for processing time.

Correct Answer: B
Section:
Explanation:
According to the Salesforce documentation1, the query plan tool is a tool that can be enabled and used in the developer console to analyze the performance of a SOQL query. The query plan tool shows the cost, cardinality, sObject type, and relative cost of each query plan that Salesforce considers for a query. The relative cost indicates how expensive a query plan is compared to the Force.com query optimizer threshold. A query plan with a relative cost above 1.0 is likely to cause a time-out error.
To identify why the query is taking too long, a data architect should use the query plan tool in the developer console (option B). This way, the data architect can see which query plan is chosen by Salesforce and how it affects the performance of the query. The data architect can also use the query plan tool to optimize the query by adding indexes, filters, or limits to reduce the cost and improve the efficiency of the query.
Using Splunk to query the system logs looking for transaction time and CPU usage (option A) is not a good solution, as it can be complex, costly, and difficult to integrate with Salesforce. It may also not provide enough information or insights to identify and optimize the query performance. Using Salesforce's query optimizer to analyze the query in the developer console (option C) is also not a good solution, as it is not a separate tool that

can be used in the developer console.The query optimizer is a feature that runs automatically when a SOQL query is executed and chooses the best query plan based on various factors2. Opening a ticket with Salesforce support to retrieve transaction logs to be analyzed for processing time (option D) is also not a good solution, as it can be time-consuming, dependent, and inefficient. It may also not provide enough information or insights to identify and optimize the query performance.

QUESTION 117
Northern Trail Outfitter has implemented Salesforce for its associates nationwide, Senior management is concerned that the executive dashboard is not reliable for their real-time decision-making. On analysis, the team the following issues with data entered in Salesforce.
Information in certain records is incomplete.
Incorrect entry in certain fields causes records to be excluded in report fitters.
Duplicate entries cause incorrect counts.
Which three steps should a data architect recommend to address the issues?

A. Periodically export data to cleanse data and import them back into Salesforce for executive reports.
B. Build a sales data warehouse with purpose-build data marts for dashboards and senior management reporting.
C. Explore third-party data providers to enrich and augment information entered in salesforce.
D. Leverage Salesforce features, such as validate rules, to avoid incomplete and incorrect records.
E. design and implement data-quality dashboard to monitor and act on records that are incomplete or incorrect

Correct Answer: B, C, D
Section:
Explanation:
According to the Salesforce documentation3, data quality is the measure of how well the data in Salesforce meets the expectations and requirements of the users and stakeholders. Data quality can be affected by various factors, such as data entry errors, data duplication, data inconsistency, data incompleteness, data timeliness, etc. To address the issues with data quality that affect the reliability of executive dashboards, a data architect should recommend:
Building a sales data warehouse with purpose-built data marts for dashboards and senior management reporting (option B). This means creating a separate database or system that stores and organizes sales data from Salesforce and other sources for analytical purposes. A data warehouse can provide a single source of truth for sales data and enable faster and more accurate reporting and analysis. A data mart is a subset of a data warehouse that focuses on a specific subject or business area, such as sales performance, customer segmentation, product profitability, etc. A data mart can provide tailored and relevant data for different users or groups based on their needs and interests.
Exploring third-party data providers to enrich and augment information entered in Salesforce (option C). This means using external services or tools that can validate, correct, update, and enhance the data that is entered or imported into Salesforce. This can help improve data quality and accuracy, and reduce data duplication and incompleteness.
Leveraging Salesforce features, such as validation rules, to avoid incomplete and incorrect records (option D). This means using features that allow defining rules and criteria to validate the data that is entered or updated by the users or integrations. This can help prevent invalid or incorrect data from being saved, and trigger actions or alerts to correct or improve the data.
Periodically exporting data to cleanse data and import them back into Salesforce for executive reports (option A) is not a good solution, as it can be time-consuming, error-prone, and inefficient. It may also cause data inconsistency and synchronization issues between Salesforce and other systems. Designing and implementing data-quality dashboard to monitor and act on records that are incomplete or incorrect (option E) is also not a good solution, as it can be complex, costly, and difficult to maintain. It may also not address the root causes of data quality issues or prevent them from occurring in the first place.

QUESTION 118
The data architect for UC has written a SOQL query that will return all records from the Task object that do not have a value in the WhatID field:
Select id, description, Subject from Task where WhatId!= NULL
When the data architect usages the query to select values for a process a time out error occurs.
What does the data architect need to change to make this query more performant?

A. Remove description from the requested field set.
B. Change query to SOSL.
C. Add limit 100 to the query.
D. Change the where clause to filter by a deterministic defined value.

Correct Answer: D

Section:
Explanation:
According to the Salesforce documentation, SOQL is a query language that allows querying data from Salesforce objects and fields. SOQL queries have various clauses and operators that can be used to filter and sort the results. However, some clauses and operators can affect the performance of SOQL queries by increasing the cost or complexity of executing them.
To make this query more performant, a data architect should change the where clause to filter by a deterministic defined value (option D). This means using a filter condition that specifies a concrete value or range of values for a field, such as WhatId = '001xx000003DGg3' or WhatId IN ('001xx000003DGg3', '001xx000003DGg4'). This can improve the performance of the query by reducing the number of records that need to be scanned and returned. A deterministic defined value can also leverage an index on the field, which can speed up the query execution.
Removing description from the requested field set (option A) is not a good solution, as it can affect the functionality or usability of the query. The description field may contain important or relevant information that is needed for the process. Changing the query to SOSL (option B) is also not a good solution, as SOSL is a different query language that allows searching text fields across multiple objects. SOSL queries have different syntax and limitations than SOQL queries, and may not return the same results or performance. Adding limit 100 to the query (option C) is also not a good solution, as it can affect the completeness or accuracy of the query. The limit clause specifies the maximum number of records that can be returned by the query, which may not include all the records that match the filter condition.

QUESTION 119
Universal Containers (UC) is in the process of migrating legacy inventory data from an enterprise resources planning (ERP) system into Sales Cloud with the following requirements:
Legacy inventory data will be stored in a custom child object called Inventory_c.
Inventory data should be related to the standard Account object.
The Inventory object should Invent the same sharing rules as the Account object.
Anytime an Account record is deleted in Salesforce, the related Inventory_c record(s) should be deleted as well.
What type of relationship field should a data architect recommend in this scenario?

A. Master-detail relationship filed on Account, related to Inventory_c
B. Master-detail relationship filed on Inventory_c, related to Account
C. Indirect lookup relationship field on Account, related to Inventory_c
D. Lookup relationship fields on Inventory related to Account

Correct Answer: B
Section:
Explanation:
According to the Salesforce documentation, a relationship field is a field that allows linking one object to another object in Salesforce. There are different types of relationship fields that have different characteristics and behaviors, such as master-detail, lookup, indirect lookup, external lookup, etc.
To recommend a type of relationship field for this scenario, where legacy inventory data will be stored in a custom child object called Inventory__c, inventory data should be related to the standard Account object, the Inventory__c object should inherit the same sharing rules as the Account object, and anytime an Account record is deleted in Salesforce, the related Inventory__c record(s) should be deleted as well, a data architect should recommend:
Master-detail relationship field on Inventory__c, related to Account (option B). This means creating a field on the Inventory__c object that references the Account object as its parent. A master-detail relationship field establishes a parent-child relationship between two objects, where the parent object controls certain behaviors of the child object. For example, a master-detail relationship field can:
Inherit the sharing and security settings from the parent object to the child object. This means that the users who can access and edit the parent record can also access and edit the related child records.
Cascade delete from the parent object to the child object. This means that when a parent record is deleted, all the related child records are also deleted.
Roll up summary fields from the child object to the parent object. This means that the parent object can display aggregated information from the child records, such as count, sum, min, max, or average.
Master-detail relationship field on Account, related to Inventory__c (option A) is not a good solution, as it reverses the direction of the relationship. This means creating a field on the Account object that references the Inventory__c object as its parent. This is not possible, as a standard object cannot be on the detail side of a master-detail relationship. Indirect lookup relationship field on Account, related to Inventory__c (option C) is also not a good solution, as it is a special type of relationship field that allows linking a custom object to a standard object on an external system using an indirect reference. This is not applicable for this scenario, as both objects are in Salesforce and do not need an external reference. Lookup relationship field on Inventory__c related to Account (option D) is also not a good solution, as it establishes a looser relationship between two objects than a master-detail relationship. A lookup relationship field does not inherit sharing and security settings from the parent object to the child object, does not cascade delete from the parent object to the child object, and does not roll up summary fields from the child object to the parent object.

QUESTION 120
A customer wants to maintain geographic location information including latitude and longitude in a custom object. What would a data architect recommend to satisfy this requirement?

A. Create formula fields with geolocation function for this requirement.

B. Create custom fields to maintain latitude and longitude information
C. Create a geolocation custom field to maintain this requirement
D. Recommend app exchange packages to support this requirement.

Correct Answer: C
Section:
Explanation:
The correct answer is C, create a geolocation custom field to maintain this requirement. A geolocation custom field is a compound field that can store both latitude and longitude information in a single field. It also supports geolocation functions and distance calculations. Creating formula fields or custom fields for latitude and longitude separately would be inefficient and redundant. Recommending app exchange packages would not be a direct solution to the requirement.

QUESTION 121
As part of addressing general data protection regulation (GDPR) requirements, UC plans to implement a data classification policy for all its internal systems that stores customer information including salesforce.
What should a data architect recommend so that UC can easily classify consumer information maintained in salesforce under both standard and custom objects?

A. Use App Exchange products to classify fields based on policy.
B. Use data classification metadata fields available in field definition.
C. Create a custom picklist field to capture classification of information on customer.
D. Build reports for customer information and validate.

Correct Answer: B
Section:
Explanation:
The correct answer is B, use data classification metadata fields available in field definition. Data classification metadata fields are standard fields that allow you to classify the sensitivity level of your data based on your organization's policies. You can use these fields to indicate whether a field contains confidential, restricted, or general data. These fields are available for both standard and custom objects in Salesforce. Using app exchange products, creating a custom picklist field, or building reports would not be as effective or consistent as using data classification metadata fields.

QUESTION 122
Northern Trail Outfitters has these simple requirements for a data export process:
File format should be in CSV.
Process should be scheduled and run once per week.
The expert should be configurable through the Salesforce UI.
Which tool should a data architect leverage to accomplish these requirements?

A. Bulk API
B. Data export wizard
C. Third-party ETL tool
D. Data loader

Correct Answer: B
Section:
Explanation:
The correct answer is B, data export wizard. The data export wizard is a tool that allows you to export your data in CSV format, schedule the export process to run once per week, and configure the export settings through the Salesforce UI. The data export wizard can handle up to 51 million records per export. The bulk API, third-party ETL tools, and data loader are also tools that can export data, but they are not as simple or user-friendly as the data export wizard.

QUESTION 123
UC is planning a massive SF implementation with large volumes of dat

a. As part of the org's implementation, several roles, territories, groups, and sharing rules have been configured. The data architect has been tasked with loading all of the required data, including user data, in a timely manner.
What should a data architect do to minimize data load times due to system calculations?

A. Enable defer sharing calculations, and suspend sharing rule calculations
B. Load the data through data loader, and turn on parallel processing.
C. Leverage the Bulk API and concurrent processing with multiple batches
D. Enable granular locking to avoid ''UNABLE _TO_LOCK_ROW'' error.

Correct Answer: A
Section:
Explanation:
The correct answer is A, enable defer sharing calculations, and suspend sharing rule calculations. Defer sharing calculations and suspend sharing rule calculations are features that allow you to temporarily disable the automatic recalculation of sharing rules when you load large volumes of data. This can improve the performance and speed of your data load process by avoiding unnecessary system calculations. Loading the data through data loader, leveraging the bulk API, or enabling granular locking are also options that can help with data load times, but they do not directly address the system calculations issue.

QUESTION 124
UC has the following system:
Billing system.
Customer support system.
CRM system.
US has been having trouble with business intelligence across the different systems. Recently US implemented a master data management (MDM) solution that will be the system of truth for the customer records.
Which MDM data element is needed to allow reporting across these systems?

A. Global unique customer number.
B. Email address.
C. Phone number.
D. Full name.

Correct Answer: A
Section:
Explanation:
The correct answer is A, global unique customer number. A global unique customer number is a data element that can uniquely identify each customer across different systems. It can be used as a key to link customer records from different sources and enable reporting across these systems. Email address, phone number, and full name are not reliable or consistent identifiers for customers, as they can change over time or be shared by multiple customers.

QUESTION 125
Universal Container is using Salesforce for Opportunity management and enterprise resource planning (ERP) for order management. Sales reps do not have access to the ERP and have no visibility into order status.
What solution a data architect recommends to give the sales team visibility into order status?

A. Leverage Canvas to bring the order management UI in to the Salesforce tab.
B. Build batch jobs to push order line items to salesforce.
C. leverage Salesforce Connect top bring the order line item from the legacy system to Salesforce.
D. Build real-time integration to pull order line items into Salesforce when viewing orders.

Correct Answer: C
Section:
Explanation:

The correct answer is C, leverage Salesforce Connect to bring the order line item from the legacy system to Salesforce. Salesforce Connect is a feature that allows you to integrate external data sources with Salesforce and access them in real time without copying or synchronizing the data. This way, the sales team can view the order status from the ERP system without having access to it. Leveraging Canvas, building batch jobs, or building real-time integration are also possible solutions, but they are more complex and costly than using Salesforce Connect.

QUESTION 126
Universal Containers (UC) is transitioning from Classic to Lightning Experience.
What does UC need to do to ensure users have access to its notices and attachments in Lightning Experience?

A. Add Notes and Attachments Related List to page Layout in Lighting Experience.
B. Manually upload Notes in Lighting Experience.
C. Migrate Notes and Attachment to Enhanced Notes and Files a migration tool
D. Manually upload Attachments in Lighting Experience.

Correct Answer: C
Section:
Explanation:
The correct answer is C, migrate Notes and Attachment to Enhanced Notes and Files using a migration tool. Enhanced Notes and Files are the new features in Lightning Experience that replace the classic Notes and Attachments. They offer more functionality and security than the classic version. To access them in Lightning Experience, you need to migrate your existing Notes and Attachments using a migration tool provided by Salesforce. Adding Notes and Attachments Related List, manually uploading Notes or Attachments, or doing nothing are not valid solutions, as they will not enable you to use the enhanced features in Lightning Experience.

QUESTION 127
A large automobile manufacturer has decided to use Salesforce as its CRM. It needs to maintain the following dealer types in their CRM:
Local dealers
Regional distributor
State distributor
Service dealer
The attributes are different for each of the customer types. The CRM users should be allowed to enter only attributes related to the customer types. The processes and business rules for each of the customer types could be different.
How should the different dealers be maintained in Salesforce?

A. Use Accounts for dealers, and create record types for each of the dealer types.
B. Create dealers as Accounts, and build custom views for each of the dealer types.
C. Use Accounts for dealers and custom picklist field for each of the dealer types
D. Create Custom objects for each dealer types and custom fields for dealer attributes.

Correct Answer: A
Section:
Explanation:
According to Trailhead2, record types are a feature that allows administrators to create different business processes, page layouts, and picklist values for different types of records within an object. Record types can be used to maintain different dealer types in Salesforce by creating record types for each of the dealer types and assigning them appropriate attributes, processes, and rules. Option A is the correct answer because it suggests using Accounts for dealers, and creating record types for each of the dealer types. Option B is incorrect because creating dealers as Accounts, and building custom views for each of the dealer types does not allow CRM users to enter only attributes related to the customer types, nor does it enable different processes and business rules for each of the customer types. Option C is incorrect because using Accounts for dealers and custom picklist field for each of the dealer types does not allow CRM users to enter only attributes related to the customer types, nor does it enable different processes and business rules for each of the customer types. Option D is incorrect because creating Custom objects for each dealer types and custom fields for dealer attributes can be unnecessary and complex, as it requires creating multiple objects and relationships instead of using the standard Account object.

QUESTION 128
Universal Containers (UC) needs to move millions of records from an external enterprise resource planning (ERP) system into Salesforce.
What should a data architect recommend to be done while using the Bulk API in serial mode instead of parallel mode?

A. Placing 20 batches on the queue for upset jobs.
B. Inserting 1 million orders distributed across a variety of accounts with potential lock exceptions.
C. Leveraging a controlled feed load with 10 batches per job.
D. Inserting 1 million orders distributed across a variety of accounts with lock exceptions eliminated and managed.

Correct Answer: B
Section:
Explanation:
According tothis article, inserting 1 million orders distributed across a variety of accounts with potential lock exceptions is a scenario where using the Bulk API in serial mode can help to prevent the lock contention issue that can occur in parallel mode.

QUESTION 129
A data architect is working with a large B2C retailer and needs to model the consumer account structure in Salesforce.
What standard feature should be selected in this scenario?

A. Individual Accounts
B. Account Contact
C. Contacts
D. Person Accounts

Correct Answer: D
Section:
Explanation:
According tothis article, person accounts are designed to store information about individual people by combining certain account and contact fields into a single record. This is suitable for a large B2C retailer that needs to model the consumer account structure in Salesforce.

QUESTION 130
Universal Containers (UC) has built a custom application on Salesforce to help track shipments around the world. A majority of the shipping records are stored on premise in an external data source. UC needs shipment details to be exposed to the custom application, and the data needs to be accessible in real time. The external data source is not OData enabled, and UC does not own a middleware tool.
Which Salesforce Connect procedure should a data architect use to ensure UC's requirements are met?

A. Write an Apex class that makes a REST callout to the external API.
B. Develop a process that calls an inviable web service method.
C. Migrate the data to Heroku and register Postgres as a data source.
D. Write a custom adapter with the Apex Connector Framework.

Correct Answer: D
Section:
Explanation:
According tothis article, the Apex Connector Framework enables developers to create custom adapters for Salesforce Connect to access data from external systems that are not OData enabled. This can meet UC's requirements of exposing shipment details to the custom application and accessing the data in real time.

QUESTION 131
Universal Containers is establishing a call center that will use Salesforce. UC receives 10 million calls and creates 100 million cases every month. Cases are linked to a custom call object using lookup relationship. UC would like to run reports and dashboard to better understand different case types being created on calls to better serve customers.
What solution should a data architect recommend to meet the business requirement?

A. Archive records to a data warehouse and run analytics on the data warehouse.
B. Leverage big objects to archive records and Einstein Analytics to run reports.
C. Leverage custom objects to store aggregate data and run analytics.
D. Leverage out-of-the-box reports and dashboard on case and interactive voice response (IVR) custom object.

Correct Answer: B
Section:
Explanation:
According tothis article, big objects can store and manage massive amounts of data on the Salesforce platform. This can help UC to archive records from other objects or bring massive datasets from outside systems into a big object. According tothis article, Einstein Analytics can be used to create dashboards and lenses to analyze big object data and get insights from advanced AI-driven analytics.

QUESTION 132
Universal Containers (UC) owns several Salesforce orgs across a variety of business units. UC management has declared that it needs the ability to report on Accounts and Opportunities from each org in one place. Once the data is brought together into a global view, management would like to use advanced AI-driven analytics on the dataset.
Which tool should a data architect recommend to accomplish this reporting requirement?

A. Run standard reports and dashboards.
B. Install a third-party AppExchange tool for multi-org reporting.
C. Use Einstein Analytics for multi-org.
D. Write a Python script to aggregate and visualize the data.

Correct Answer: C
Section:
Explanation:
According tothis article, Einstein Analytics for multi-org can help UC to report on accounts and opportunities from each org in one place. It can also provide advanced AI-driven analytics on the dataset using Einstein Discovery.

QUESTION 133
Universal Containers (UC) is migrating data from legacy system to Salesforce. During data analysis it was discovered that data types of fields being migrated do not match with Salesforce data types.
Which solution should a data architect use to ensure successful data migrations?

A. Migrate legacy data to a staging database for mapping then leverage an ETL tool to transform the data and load into Salesforce.
B. Export legacy data into the staging database and leverage stored procedures to transform data types before loading into Salesforce.
C. Migrate the legacy data leveraging an ETL tool to transform data types and load data into Salesforce.
D. Export legacy data into CSV files and leverage data loader to load data into Salesforce.

Correct Answer: C
Section:
Explanation:
According tothis article, migrating the legacy data leveraging an ETL tool to transform data types and load data into Salesforce can be a way to ensure successful data migrations. The article states that ETL tools can help with data cleansing, mapping, transformation, and loading, and that they can handle different data types and formats. The article also provides some best practices for data migration, such as identifying the data to migrate, creating templates for the data, preparing the destination org, and validating the data.

QUESTION 134
Based on government regulations, a Salesforce customer plans to implement the following in Salesforce for compliance:
Access to customer information based on record ownership
Ability for customers to request removal of their information from Salesforce
Prevent users from accessing Salesforce from outside company network (virtual private network, or VPN)
What should a data architect recommend to address these requirements?

A. Contact Salesforce support to restrict access only with VPN and other requirements
B. Allow users access to Salesforce through a custom web application hosted within VPN.
C. Implement IP restrictions, sharing settings, and custom Apex to support customer requests.
D. Implement Salesforce shield with Event Monitoring to address the requirement.

Correct Answer: C
Section:
Explanation:
IP restrictions can be used to prevent users from accessing Salesforce from outside the company network or VPN. According tothis article, sharing settings can be used to control access to customer information based on record ownership. According tothis article, custom Apex can be used to support customer requests for removal of their information from Salesforce.

QUESTION 135
Universal Containers has multiple systems all containing and maintaining customer dat
a. Although point-to-point integrations are in place, customers are complaining about consistency in the data.
What solution should the data architect recommend?

A. Improve existing point-to-point integrations
B. Data cleanse each system
C. Perform a onetime synchronization to level set the built up inconsistencies
D. An MDM solution as the customer master, with centralized integrations to ensure consistency across all systems.

Correct Answer: D
Section:
Explanation:
Master data management (MDM) is a solution that helps organizations manage their master data across multiple systems and ensure consistency and quality. An MDM solution can act as the customer master, with centralized integrations to other systems, to avoid data duplication and inconsistency.

QUESTION 136
Northern Trail Outfitters (NTO) has multiple Salesforce orgs based on regions. Users need read-only access to customers across all Salesforce orgs.
Which feature in Salesforce can be used to provide access to customer records across all NTO orgs?

A. Salesforce Connect
B. Salesforce 2 Salesforce
C. Federated Search
D. External APIs

Correct Answer: A
Section:
Explanation:
Salesforce Connect is a feature that allows users to access data from external sources and multiple Salesforce orgs, using either clicks or code. Salesforce Connect can provide read-only access to customer records across all NTO orgs, without replicating or storing the data in Salesforce.

QUESTION 137
Northern Trail Outfitters (NTO) plans to maintain contact preferences for customers and employees. NTO has implemented the following:
1. Customers are Person Accounts for their retail business.
2. Customers are represented as Contacts for their commercial business.
3. Employees are maintained as Users.

4. Prospects are maintained as Leads.
NTO needs to implement a standard communication preference management model for Person Accounts, Contacts, Users, and Leads.
Which option should the data architect recommend NTO to satisfy this requirement?

A. Create custom fields for contact preferences in Lead, Person Account, and Users objects.
B. Create case for contact preferences, and use this to validate the preferences for Lead, Person Accounts, and Users.
C. Create a custom object to maintain preferences and build relationships to Lead, Person Account, and Users.
D. Use Individual objects to maintain the preferences with relationships to Lead, Person Account, and Users.

Correct Answer: D
Section:
Explanation:
The Individual object is a standard object that lets you store details about data privacy and protection preferences for person accounts, contacts, users, and leads. The Individual object can be used to implement a standard communication preference management model for NTO, with relationships to the other objects.

QUESTION 138
Universal Containers (UC) is building a Service Cloud call center application and has a multisystem support solution. UC would like or ensure that all systems have access to the same customer information.
What solution should a data architect recommend?

A. Make Salesforce the system of record for all data.
B. Implement a master data management (MDM) strategy for customer data.
C. Load customer data in all systems.
D. Let each system be an owner of data it generates.

Correct Answer: B
Section:
Explanation:
A master data management (MDM) strategy for customer data can help UC ensure that all systems have access to the same customer information, without loading or duplicating data in all systems. An MDM strategy can also help UC avoid data conflicts and inconsistencies that may arise from having multiple systems as owners of data.

QUESTION 139
Universal Containers uses Apex jobs to create leads in Salesforce. The business team has requested that lead creation failures should be reported to them.
Which option does Apex provide to report errors from this Apex job?

A. Use Custom Object to store leads, and allow unprocessed leads to be reported.
B. Save Apex errors in a custom object, and allow access to this object for reporting.
C. Use Apex services to email failures to business when error occurs.
D. Use AppExchange package to clean lead information before Apex job processes them.

Correct Answer: C
Section:
Explanation:
saving Apex errors in a custom object can be a way to report errors from an Apex job. The article provides an example of how to create a custom object called AsyncApexError__c and use a trigger to insert error records into it. The custom object can then be used for reporting or notification purposes.

QUESTION 140
Northern Trail Outfitters (NTO) wants to capture a list of customers that have bought a particular product. The solution architect has recommended to create a custom object for product, and to create a lookup relationship between its customers and its products.

Products will be modeled as a custom object (NTO_ Product__ c) and customers are modeled as person accounts. Every NTO product may have millions of customers looking up a single product, resulting in a lookup skew. What should a data architect suggest to mitigate Issues related to lookup skew?

A. Create multiple similar products and distribute the skew across those products.
B. Change the lookup relationship to master-detail relationship.
C. Create a custom object to maintain the relationship between products and customers.
D. Select Clear the value of this field option while configuring the lookup relationship.

Correct Answer: A
Section:
Explanation:
creating multiple similar products and distributing the skew across those products can be a way to mitigate issues related to lookup skew. The article explains that lookup skew happens when a very large number of records are associated with a single record in the lookup object, and this can cause record locking and performance issues. The article suggests creating multiple copies of the same product record and assigning different child records to each copy, so that the number of child records per parent record is reduced.

QUESTION 141
Northern Trail Outfitters would like to report on the type of customers. A custom field for customer type was created in Account object. Users need to be limited to the following defined choices when entering information in this field:
1. High Value
2. Medium Value
3. Low Value
Which strategy should a data architect recommend to configure customer type?

A. Lookup to a custom object with defined choices.
B. Single-select restricted picklist with defined choices.
C. Provide help text to guide users with defined choices.
D. Create a validation rule to limit entry to defined choices.

Correct Answer: B
Section:
Explanation:
single-select restricted picklist with defined choices can be a way to configure customer type. The article states that picklists are fields that allow users to select one or more predefined values from a list, and restricted picklists ensure that users can only select from the defined values. This can help to limit the choices for customer type and ensure data quality.

QUESTION 142
Northern Trail Outfitters needs to implement an archive solution for Salesforce dat
a. This archive solution needs to help NTO do the following:
1. Remove outdated Information not required on a day-to-day basis.
2. Improve Salesforce performance.
Which solution should be used to meet these requirements?

A. Identify a location to store archived data and use scheduled batch jobs to migrate and purge the aged data on a nightly basis,
B. Identify a location to store archived data, and move data to the location using a time-based workflow.
C. Use a formula field that shows true when a record reaches a defined age and use that field to run a report and export a report into SharePoint.
D. Create a full copy sandbox, and use it as a source for retaining archived data.

Correct Answer: A
Section:

Explanation:
Identifying a location to store archived data and using scheduled batch jobs to migrate and purge the aged data on a nightly basis can be a way to meet the requirements for an archive solution. The article provides a use case of how to use Heroku Connect, Postgres, and Salesforce Connect to archive old data, free up space in the org, and still retain the option to unarchive the data if needed. The article also explains how this solution can improve Salesforce performance and meet data retention policies.

QUESTION 143
Northern Trail Outfitters Is planning to build a consent form to record customer authorization for marketing purposes.
What should a data architect recommend to fulfill this requirement?

A. Use custom fields to capture the authorization details.
B. Create a custom object to maintain the authorization.
C. Utilize the Authorization Form Consent object to capture the consent.
D. Use AppExchange solution to address the requirement.

Correct Answer: C
Section:
Explanation:
The Authorization Form Consent object is a standard object that allows you to capture customer consent for marketing purposes. It has fields such as Consent Captured Date, Consent Captured Source, Consent Description, and Consent Status. You can use this object to create consent forms and track customer responses. This is the best option to fulfill the requirement, as it does not require any custom development or external solution.

QUESTION 144
A consumer products company has decided to use Salesforce for its contact center. The contact center agents need access to the following information in Service Console when a customer contacts them:
1. Customer browsing activity on its website stored on its on premise system
2. Customer interactions with sales associates at its retail stores maintained in Salesforce
3. Contact center interactions maintained in Salesforce
4. Email campaign activity to customer from its marketing systems.
What should a data architect do to fulfill these requirements with minimum development effort in Salesforce?

A. Create web tabs in Service Console to show website and marketing activities.
B. Build custom components in Service Console to bring information from the marketing and website information.
C. Use Salesforce Connect to integrate the website and the marketing system into Service Console using external objects.
D. Build customer view in Service Console with components that show website data and marketing data as mashup.

Correct Answer: D
Section:
Explanation:
Building a customer view in Service Console with components that show website data and marketing data as mashup is the best option to fulfill the requirement with minimum development effort in Salesforce. A mashup is a technique that combines data from different sources into a single user interface. You can use Visualforce pages or Lightning components to create mashups that display data from external systems such as your website and marketing system. This way, you can provide your contact center agents with a comprehensive view of the customer information they need.

QUESTION 145
Northern Trail Outfitters <NTO> is streaming IoT data from connected devices to a cloud database. Every 24 hours. 100,000 records are generated.
NIO employees will need to see these lol records within Salesforce and generate weekly reports on it. Developers may also need to write programmatic logic to aggregate the records and incorporate them into workflows.
Which data pattern will allow a data architect to satisfy these requirements, while also keeping limits in mind?

A. Bidirectional integration
B. Unidirectional integration
C. Virtualization

D. Persistence

Correct Answer: D
Section:
Explanation:
Persistence is the data pattern that will allow a data architect to satisfy the requirements, while also keeping limits in mind. Persistence means storing data from external sources in Salesforce objects, either standard or custom. This allows you to access the data within Salesforce and use it for reporting, analytics, workflows, and other features. Persistence also helps you avoid hitting API limits or performance issues when accessing large volumes of data from external systems. You can use various tools such as Data Loader, Bulk API, or Platform Events to persist IoT data from connected devices to a cloud database in Salesforce.

QUESTION 146
A national nonprofit organization is using Salesforce to recruit members. The recruitment process requires a member to be matched with a volunteer opportunity. Given the following:
1. A record is created in Project__ c and used to track the project through completion.
2. The member may then start volunteering and is required to track their volunteer hours, which is stored in VTOTime_c object related to the project.
3. Ability to view or edit the VTOTime__c object needs to be the same as the Project__ c record.
4. Managers must see total hours volunteered while viewing the Project__ c record.
Which data relationship should the data architect use to support this requirement when creating the custom VTOTime__c object?

A. Lookup Field on Project_c to VTOTime_c displaying a list of VTOTime__c in a related list.
B. Lookup field on VTOTime_c to Project_c with formula filed on Project__ c showing Sum of hours from VTOTime__c records.
C. Master Detail Field on VTOTime_ c to Project_c with rollup summary field on Project __c showing sum of hours from VTOTime_c records.
D. Master Detail field on Project _c to VTOTime _c showing a list of VTOTime_c Records in a related list.

Correct Answer: C
Section:
Explanation:
A master-detail field on VTOTime__c to Project__c is the data relationship that the data architect should use to support the requirement when creating the custom VTOTime__c object. A master-detail relationship creates a parent-child relationship between two objects, where the master record controls certain behaviors of the detail record, such as security, ownership, deletion, and roll-up summary fields. By using a master-detail field on VTOTime__c to Project__c, you can ensure that the ability to view or edit the VTOTime__c object is the same as the Project__c record, and that managers can see the total hours volunteered while viewing the Project__c record using a roll-up summary field.