ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

A CMDB Administrator is implementing Vulnerability Response or Security Incident Response and needs to ensure customers have enough context to estimate risk and set task priorities.

Which Get Well Playbook from the CSDM Data Foundations Dashboard helps with this?

A. Locations without a Parent Location

B. Application Services with Business Application Relationships

C. Named Product Models without Product Owners

D. Percentage of Custom Status Values for CI Life Cycle Stages

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, Vulnerability Response and Security Incident Response rely heavily on business context to accurately assess risk, prioritize remediation tasks, and communicate impact to stakeholders. From a CSDM (Common Service Data Model) perspective, this context is primarily delivered through properly modeled relationships between Application Services and Business Applications.

The "Application Services with Business Application Relationships" Get Well Playbook directly addresses this requirement. In CSDM, Application Services represent the technical, deployable services that run in the environment, while Business Applications represent the logical applications that support business capabilities. When these two are correctly related, security teams can clearly understand which business processes, customers, and revenue streams are affected by a vulnerability or security incident.

Without this relationship, vulnerabilities may still be detected, but they lack meaningful prioritization. For example, a critical vulnerability on an application service supporting a revenuegenerating or customer-facing business application should be addressed far more urgently than one

tied to a low-impact internal tool. This relationship is what enables risk-based prioritization, rather than purely technical severity-based prioritization.

The other options do not fulfill this need. Location hierarchy issues (Option A) and CI lifecycle status consistency (Option D) relate more to CMDB hygiene and governance, not security context. Product ownership gaps (Option C) affect accountability but do not directly enable risk estimation during security response.

Therefore, Option B is the correct and CSDM-aligned Get Well Playbook for ensuring sufficient business context in Vulnerability Response and Security Incident Response workflows.

Question: 2

A customer's CMDB is aligned to the CSDM Walk stage.

What benefit is provided by the CMDB?

A. Allows for additional stratification of technical teams' support structure along the lines of OLAs and commitments

B. Improves the implementation velocity of APM Foundation for future business application rationalization

C. Enables impact assessments for incident, problem, and change on Business Services

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In the CSDM Walk stage, an organization has moved beyond basic data hygiene (Crawl) and has established foundational service models, especially Business Services and their relationships to underlying technical components. One of the most important and immediate benefits of reaching this stage is the ability to perform reliable impact analysis across ITSM processes.

When Business Services are correctly defined and related to Application Services, applications, and infrastructure CIs, the CMDB becomes a decision-support system rather than just a data repository. This enables impact assessments for Incident, Problem, and Change Management, which is exactly what Option C describes. For example, when an incident is logged against a CI, ServiceNow can automatically determine which Business Services are impacted and who the affected stakeholders are. Similarly, during Change Management, planners can assess downstream risk by identifying which business-facing services could be disrupted.

Option A is more aligned with advanced operational governance and support model optimization, which typically appears later as organizations mature toward the Run stage. Option B relates to Application Portfolio Management (APM) acceleration, which benefits more from accurate

application ownership and lifecycle data rather than core Walk-stage service modeling.

Therefore, the correct and CSDM-aligned benefit at the Walk stage is enabling impact assessments for incident, problem, and change on Business Services, making Option C the verified answer.

Question: 3

A CMDB Administrator needs to import external data into the CMDB. To reduce the risk of creating duplicates and prevent updates from unauthorized sources, it must be ensured that the Identification and Reconciliation Engine (IRE) is not bypassed.

What is the recommended method to import data into the CMDB utilizing the Identification and Reconciliation API?

A. IntegrationHub ETL

B. Table API (REST API or SOAP API)

C. Import Sets and Transform Maps

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, protecting CMDB data quality during ingestion is a core Data Foundations principle. The Identification and Reconciliation Engine (IRE) is designed to ensure that CI records are uniquely identified, merged correctly, and protected from unauthorized overwrites. Any ingestion method that bypasses IRE introduces a high risk of duplicates and data corruption.

IntegrationHub ETL is the recommended method because it is natively designed to work with the Identification and Reconciliation API. When properly configured, IntegrationHub ETL ensures that incoming data is processed through IRE, applying identification rules, reconciliation rules, and source precedence. This allows multiple data sources to coexist safely while maintaining CMDB integrity.

Option B (Table API) is explicitly discouraged for CMDB ingestion because it writes directly to CMDB tables and bypasses IRE entirely, making it one of the most common causes of duplicate and conflicting CI records. While REST and SOAP APIs are powerful, they are not safe for CMDB ingestion unless they explicitly invoke the IRE API, which most generic table integrations do not.

Option C (Import Sets and Transform Maps) can be configured to call IRE, but this requires additional scripting and strict governance. Because of this complexity and higher risk of misconfiguration, it is not the recommended approach when safer, purpose-built options exist.

Therefore, IntegrationHub ETL is the verified and best-practice answer, making Option A correct.

Question: 4

(Choose 2 options)

A CMDB Administrator has built a number of Technology Management Service Offerings (Technical Service Offerings) based on Dynamic CI Groups to better maintain group alignment for the member CIs.

Which groups are synced to CIs from the offering that has a relationship to a Dynamic CI Group?

A. Approval Group

B. Managed by Group

C. Owned by Group

D. Support Group

Answer: BD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, Dynamic CI Groups are a core Data Foundations capability used to automatically manage CI membership based on rules rather than manual maintenance. When Technology Management Service Offerings (Technical Service Offerings) are related to Dynamic CI Groups, ServiceNow uses those relationships to synchronize operational support attributes to the member CIs.

The two CI attributes that are intentionally designed to sync in this model are the Managed by Group and the Support Group. These groups directly influence operational ownership and support routing, which is why they are automatically aligned when Dynamic CI Groups are used. This ensures that incidents, changes, problems, and operational tasks are routed consistently as CI membership changes over time.

The Support Group defines who provides day-to-day operational support and is critical for Incident and Request Management workflows. The Managed by Group represents the team responsible for the technical lifecycle and operational health of the CI. Synchronizing these attributes eliminates manual updates and reduces misrouted tickets, which is a key goal of Configuration Management maturity.

The Approval Group (Option A) is not synced because approvals are process-driven and often context-specific rather than CI-driven. Similarly, the Owned by Group (Option C) represents accountability or financial ownership, which is intentionally decoupled from dynamic operational

grouping to avoid unintended governance changes.

Therefore, the correct answers are B (Managed by Group) and D (Support Group).

Question: 5

Which is a purpose or requirement of CMDB Data Manager in ServiceNow?

A. Encrypts archived records for enhanced security

B. Automates the enforcement of relationship rules between CIs in the CMDB

C. Automates the archival and deletion of records based on retention policies

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Manager capability in ServiceNow is designed to support CMDB governance, specifically around data lifecycle management. Its primary purpose is to ensure that CI records are retained, archived, and deleted in accordance with defined retention policies, regulatory requirements, and organizational data governance standards.

As CMDBs mature, they naturally accumulate obsolete, retired, or decommissioned CIs. If these records are not properly managed, they negatively impact CMDB health, reporting accuracy, discovery reconciliation, and performance. CMDB Data Manager addresses this by automating the archival and deletion of records once lifecycle conditions and retention thresholds are met.

Option A is incorrect because encryption of archived records is handled by platform-level security and data protection features, not CMDB Data Manager. Option B is also incorrect because relationship rule enforcement is managed through CSDM guidance, CMDB relationship rules, and identification/reconciliation logic—not by CMDB Data Manager.

By automating retention-based archival and cleanup, CMDB Data Manager helps organizations maintain a lean, compliant, and high-quality CMDB, which directly supports CMDB Health metrics such as correctness and compliance.

Therefore, the correct and verified answer is Option C.

Question: 6

A CMDB Administrator identifies duplicate CIs. One was created by a manual import, and the other was created by automated discovery. The discovered CI has the latest IP address, while the manually

imported CI has an accurate relationship to a critical business application.

How does the Administrator use the Duplicate CI Remediator to resolve this issue?

A. Merge the two CIs automatically, retaining all attributes from the discovered CI

B. Retain the manually imported CI and delete the discovered CI

C. Retain the discovered CI, but merge the relationship from the manually imported CI

D. Retain the discovered CI and delete the manually imported CI

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, the Duplicate CI Remediator is designed to resolve duplicate records while preserving the most authoritative data from each source. Data Foundations guidance clearly states that automated discovery is the system of record for technical attributes, such as IP address, hostname, and operational status, while manually maintained records often contain valuable business context, such as relationships to business applications or services.

In this scenario, the discovered CI contains the most accurate and up-to-date technical data, making it the correct CI to retain as the primary record. However, the manually imported CI has a critical relationship to a business application, which is essential for impact analysis, incident prioritization, and CSDM alignment. Deleting this CI without preserving the relationship would result in loss of business context and reduced CMDB value.

The Duplicate CI Remediator supports selective merging, allowing administrators to retain one CI while merging specific attributes or relationships from the duplicate. Option C reflects this best practice by retaining the discovered CI and merging the relationship from the manually imported CI, ensuring both technical accuracy and business relevance are preserved.

Options A and D would result in the loss of important relationship data, while Option B would discard the discovered CI, violating the principle that discovery should be the authoritative source for technical attributes.

Therefore, Option C is the correct and Data Foundations–aligned answer.

Question: 7

What is the difference between Data Certification and Attestation policies when managing a CI?

A. Attestation requires correcting specific attributes of a CI, while Data Certification tracks acknowledgement the CI still exists

B. Attestation can be scheduled, while Data Certification cannot be scheduled

C. Attestation can be assigned to a group or an individual, while Data Certification can only be assigned to an individual

D. Attestation tracks acknowledgement the CI still exists, while Data Certification requires validating specific attributes of a CI

Answer: D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Within ServiceNow CMDB governance, Attestation and Data Certification serve distinct but complementary purposes. The key difference lies in what is being validated.

Attestation is focused on existence and ownership confirmation. When a CI is attested, the assigned user or group is asked to confirm that the CI still exists, is still relevant, and is still owned or managed by the appropriate team. No detailed attribute-level validation is required. This lightweight process is commonly used to prevent "ghost CIs" from lingering in the CMDB.

Data Certification, on the other hand, is more rigorous. It requires the certifier to validate specific attributes of the CI, such as lifecycle status, support group, environment, or service relationships. Certification ensures data correctness and completeness, which directly impacts CMDB Health scores and downstream processes like Change and Incident Management.

Options A, B, and C incorrectly describe these mechanisms or their assignment and scheduling capabilities. Both attestation and certification can be scheduled and assigned flexibly, but their validation depth is what truly differentiates them.

Therefore, Option D correctly describes the distinction: attestation confirms existence, while data certification validates CI attributes.

Question: 8

CMDB class owners are receiving tasks under the "My Work" tab in the CMDB Workspace.

Which CMDB management tool is generating those tasks?

A. De-duplication templates

B. CMDB Data Manager

C. CMDB Health Dashboard

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Manager is the ServiceNow capability responsible for generating actionable governance tasks and assigning them to CI class owners and data stewards. These tasks appear directly in the "My Work" tab within the CMDB Workspace, enabling proactive and role-based CMDB governance.

CMDB Data Manager focuses on data lifecycle management, including archival, retirement, and cleanup of CIs based on defined policies. When lifecycle rules or retention thresholds are met—or when human validation is required—the Data Manager creates tasks to prompt responsible owners to take action. This ensures that CMDB data remains accurate, compliant, and lean over time.

The CMDB Health Dashboard (Option C) provides visibility into health metrics such as completeness, correctness, and compliance, but it does not generate tasks. Similarly, De-duplication templates (Option A) support duplicate identification and remediation workflows, but they do not create ongoing governance tasks in the CMDB Workspace.

By surfacing tasks in "My Work," CMDB Data Manager operationalizes governance and embeds accountability into daily workflows, which is a key principle of CMDB Data Foundations.

Therefore, the correct answer is Option B – CMDB Data Manager.

Question: 9

The CMDB Administrator group aims to display meaningful results on the CMDB Health Dashboard – Compliance Scorecard for server records that are not on the latest patch.

What must be configured to achieve this goal?

A. Certification Filter, Certification Template, Audit

B. Technical Service Offerings, Dynamic CI Groups, CMDB Groups

C. Stale, Orphan, Duplicate

D. Certification Policies, Data Filters, Scheduled Jobs

Answer: D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, the Compliance dimension of the CMDB Health Dashboard is driven by Data Certification. To surface meaningful compliance results—such as identifying servers that are not on the latest patch—the platform requires a combination of Certification Policies, Data Filters, and Scheduled Jobs.

Certification Policies define what data must be validated and which attributes are subject to compliance checks (for example, patch level, OS version, or last update date). Data Filters scope the population—such as server classes only—ensuring the compliance evaluation targets the correct CIs. Scheduled Jobs automate when certifications run, keeping compliance scores current and reflective of the latest state.

Options A and C are incorrect because audits and stale/orphan/duplicate checks relate to other health dimensions (correctness and completeness), not compliance. Option B focuses on service modeling and group alignment, which does not directly drive compliance scoring for patch currency.

Therefore, configuring Certification Policies, Data Filters, and Scheduled Jobs is required to accurately measure and display patch compliance on the CMDB Health Dashboard.

Question: 10

A CMDB Administrator needs to identify which attributes have been created specifically for the Windows Server class.

Which tab in the Attributes section is used?

A. Child

B. Added

C. All

D. Derived

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Within the CMDB class dictionary in ServiceNow, attributes can be inherited from parent classes or defined directly on a specific class. To identify attributes created specifically for the Windows Server class, administrators must use the Added tab.

The Added tab displays attributes that are unique to the selected class and not inherited from parent classes (such as Server or Computer). This is essential for understanding class-specific extensions—

like Windows-only configuration details—that were introduced to support platform requirements, discovery enhancements, or organizational needs.

The All tab shows every attribute available to the class, including inherited and added attributes, which makes it difficult to isolate class-specific additions. The Child tab focuses on attributes inherited by subclasses, not attributes introduced at the current class level. The Derived tab shows attributes calculated or derived from other data, not necessarily those created specifically for the class.

Using the Added tab supports best practices for configuration transparency, impact analysis during upgrades, and governance—especially important in Data Foundations to minimize unnecessary customization and maintain

upgrade-safe designs.

Therefore, the correct answer is B – Added.

Question: 11

An organization is changing data centers and needs to know the consequences of the planned changes.

How can Application Service Mapping be used as part of Change Management?

A. To identify which devices will go ofline first
B. To understand the business impact of CIs
C. To understand the physical location of CIs

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Application Service Mapping is a critical capability in ServiceNow for enabling business-aware Change Management. Its primary value is not in identifying physical shutdown sequences or CI locations, but in translating technical changes into business impact.

When an organization plans a data center move, multiple infrastructure components—servers, databases, network devices—may be affected. On their own, these technical CIs provide little insight into business risk. Application Service Mapping connects these CIs to Application Services and Business Services as defined by the Common Service Data Model (CSDM). This relationship allows Change Managers to see which business services, customers, and processes are impacted by the planned change.

By leveraging service maps, Change Management can answer critical questions such as:

Which customer-facing services may experience downtime?

What revenue-generating or mission-critical services are at risk?

Which stakeholders must be notified or involved in approvals?

Option A is incorrect because service mapping does not determine shutdown order; that is handled by infrastructure planning. Option C focuses on physical location data, which is typically managed through Location CIs and Discovery, not service mapping.

Therefore, the correct answer is B – To understand the business impact of CIs, which aligns directly with ITIL 4, CSDM, and Change Management best practices.

Question: 12

What ensures data volume in the CMDB is manageable?

A. Business Rules

B. Scheduled Jobs

C. Archive Policies

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Managing CMDB data volume is a key Data Foundations governance objective. Over time, CMDBs naturally accumulate retired, obsolete, or decommissioned CIs. If these records are not properly managed, they degrade CMDB performance, reduce reporting accuracy, and negatively impact discovery reconciliation and health scores.

Archive Policies are the mechanism designed to address this challenge. They define when CI records should be archived or deleted based on lifecycle state, age, or retention requirements. By automating archival and cleanup, archive policies ensure that only relevant, active CIs remain in the operational CMDB, keeping data volume manageable and performant.

Business Rules (Option A) are used to enforce logic during record creation or updates, not for longterm data volume control. Scheduled Jobs (Option B) may execute tasks, but without archive policies they have no governance logic to determine what should be removed or retained.

Archive policies work in conjunction with CMDB Data Manager to enforce lifecycle-based retention and cleanup, making them the correct and verified answer.

Therefore, Option C – Archive Policies is correct.

Question: 13

(Choose 2 options)

The Configuration Manager is preparing justification to utilize the CMDB Data Foundations Dashboard.

Which benefits align with the usage of this dashboard?

A. It automates approval processes for Change Management

B. It provides actionable insights to improve data quality and completeness

C. It helps detect and eliminate duplicate records in the CMDB

D. It enables monitoring and tracking of CMDB health over time

Answer: BD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Foundations Dashboard is designed to provide visibility, insight, and guidance into the overall health of CMDB data. Its purpose is not to automate ITSM workflows, but to enable informed decision-making and continuous improvement of configuration data.

One of its primary benefits is providing actionable insights to improve data quality and completeness (Option B). The dashboard highlights gaps in CI attributes, missing relationships, and compliance issues, enabling CMDB administrators and data owners to take targeted corrective actions using Get Well Playbooks.

Another core benefit is enabling organizations to monitor and track CMDB health over time (Option D). The dashboard presents trends across health dimensions—completeness, correctness, and compliance—allowing teams to measure progress, justify investments, and demonstrate maturity improvements aligned to CSDM adoption stages.

Option A is incorrect because Change Management approvals are handled by workflow and policy engines, not the Data Foundations Dashboard. Option C is also incorrect because duplicate detection and remediation are handled through de-duplication tools and the Duplicate CI Remediator, not directly by the dashboard itself.

Therefore, the correct answers are B and D, which accurately reflect the strategic and operational value of the CMDB Data Foundations Dashboard.

Question: 14

A Configuration Manager needs to leverage a policy type to automate the creation and assignment of tasks to validate the existence of CIs.

Which policy type should be used to accomplish this goal?

A. Certification

B. Delete

C. Retire

D. Attestation

Answer: D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, validating whether Configuration Items (CIs) still exist is a core CMDB governance activity. Over time, environments change rapidly—servers are decommissioned, cloud resources are torn down, and applications are replaced. If existence validation is not enforced, the CMDB quickly fills with obsolete or “ghost” CIs.

Attestation policies are specifically designed to address this need. An attestation policy automatically generates and

assigns tasks to responsible users or groups, asking them to confirm that a CI still exists and is still relevant. This process focuses on acknowledgment rather than deep data validation, making it lightweight and scalable across large CMDBs.

Certification policies (Option A) are used when specific attributes must be validated, such as lifecycle status, support group, or environment. While important for data correctness, certification is not intended solely to confirm CI existence. Delete (Option B) and Retire (Option C) policies are lifecycle actions that remove or transition records, but they do not validate existence before taking action.

Attestation integrates cleanly with CMDB Workspace, assigns tasks automatically, and supports auditability—ensuring accountability for CI ownership. This makes it the correct and Data Foundations–aligned policy type for validating CI existence.

Therefore, Option D – Attestation is the correct answer.

Question: 15

(Choose 2 options)

Configuration Management needs to ensure data quality for all CIs in the CMDB.

What areas of data quality for CIs are included in the CMDB Health Dashboard?

A. Downgraded CIs

B. Upgraded CIs

C. Missing CIs

D. Stale CIs

E. Duplicate CIs

Answer: DE

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Health Dashboard is a central component of CMDB Data Foundations insight and governance. It measures and tracks data quality using well-defined health indicators that focus on the accuracy, relevance, and usability of CI data.

Two key data quality areas included in the dashboard are Stale CIs and Duplicate CIs.

Stale CIs (Option D) refer to configuration items that have not been updated within a defined time window. These records are risky because they may no longer reflect the current state of the environment, leading to inaccurate impact

analysis, poor change decisions, and misrouted incidents. Monitoring staleness helps organizations identify where discovery, integrations, or ownership processes are failing.

Duplicate CIs (Option E) occur when the same real-world asset or service is represented by multiple records. Duplicates undermine trust in the CMDB, distort reporting, and break service mappings. The CMDB Health Dashboard highlights duplicate trends and integrates with de-duplication and remediation workflows to address them.

Options A (Downgraded CIs), B (Upgraded CIs), and C (Missing CIs) are not standard CMDB Health Dashboard quality dimensions. While "missing" data may be inferred through completeness checks, Missing CIs as a category is not directly tracked.

Therefore, the correct answers are D – Stale CIs and E – Duplicate CIs, which are core CMDB Health indicators used to maintain high-quality configuration data.

Question: 16

A CMDB Manager wants to improve data quality using the CMDB Health Dashboard.

What needs to happen to generate CMDB health scores?

A. The scheduled jobs for the CMDB Health Dashboard must be activated

B. Nothing, CMDB health scores are calculated by default

C. The plugin, CMDB health calculation, needs to be installed

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, the CMDB Health Dashboard does not calculate health scores in real time by default. Instead, health scores are generated and refreshed by scheduled calculation jobs that evaluate CI data against defined health rules across the dimensions of completeness, correctness, and compliance.

To generate and maintain CMDB health scores, the scheduled jobs for CMDB Health must be active. These jobs periodically scan the CMDB, apply health rules (for example, required attributes populated, lifecycle status compliance, certification results), and calculate scores that are displayed on the dashboard and scorecards. Without these scheduled jobs running, the dashboard cannot produce current or meaningful health metrics.

Option B is incorrect because CMDB health scoring is not automatic or real-time; it depends on scheduled processing. Option C is also incorrect because CMDB Health is part of the core CMDB/Data Foundations capability in ServiceNow and does not require a separate "CMDB health calculation" plugin to be installed in modern implementations.

Activating and maintaining these scheduled jobs ensures that health scores remain accurate, trendable over time, and useful for governance decisions. This is a foundational requirement for using the CMDB Health Dashboard as a data quality improvement tool.

Therefore, the correct answer is A – The scheduled jobs for the CMDB Health Dashboard must be activated.

Question: 17

In a company, there is a need to understand the CSDM maturity level required. Different stakeholders listed several use cases they expect over time.

Which use case requires information objects?

A. The Asset Management team wants to understand asset lifecycle compliance in a Business

Application context

B. The Event Operations team wants to automate their events into incidents for operational actions

C. The Customer Service team wants to onboard proactive case management

D. The SecOps team wants to understand the operational risk in the Business Application context

E. The Business Service Management team wants to understand the operational impact for their consumer parties

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Within the Common Service Data Model (CSDM), information objects are used to represent non-CI data entities that provide important business or governance context but are not configuration items themselves. These objects are especially important when extending service visibility beyond pure infrastructure and application relationships.

The use case described in Option A—understanding asset lifecycle compliance in a Business Application context—explicitly requires information objects. Asset lifecycle data (such as financial state, depreciation, warranty, and compliance milestones) is typically managed in IT Asset Management (ITAM) and must be associated to Business Applications without converting every asset-related data point into a CI. Information objects enable this linkage while maintaining clean CMDB boundaries.

Option B focuses on event-to-incident automation, which relies on CIs, technical services, and operational relationships, not information objects. Option C (proactive case management) is primarily a CSM and service offering use case. Option D (SecOps risk context) relies on application services and business application relationships, not information objects. Option E (business service impact) is addressed through service modeling and service mapping, again without requiring information objects.

Information objects are introduced as organizations mature and need to integrate governance, financial, or compliance data with service and application models—making asset lifecycle compliance the correct match.

Therefore, the correct answer is A.

Question: 18

A CMDB Administrator has installed a Service Graph Connector and customized a script transform.

What will happen on subsequent upgrades if the default definition of the script transform is updated?

A. The upgrade stops and reports an error

B. A skipped change is created and no change is made to the script transform definition

C. The Service Graph Connector upgrade refuses to start

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, Service Graph Connectors deliver data ingestion patterns using protected, upgradesafe artifacts, including script transforms. When a customer customizes a script transform provided by a Service Graph Connector, ServiceNow follows standard update set and upgrade behavior to protect customer customizations.

During a subsequent upgrade, if the out-of-box (default) script transform definition changes, ServiceNow does not overwrite the customized version. Instead, the platform records a skipped change, indicating that an update was available but intentionally not applied due to a local customization. This ensures customer-specific logic is preserved while still maintaining transparency about what changed in the newer release.

Option A is incorrect because upgrades do not halt due to customized transforms. Option C is also incorrect because Service Graph Connector upgrades proceed normally; they do not refuse to start because of customizations.

This behavior aligns with Data Foundations best practices: avoid modifying OOTB content when possible, but when customization is necessary, ensure it is protected during upgrades. Administrators should review skipped changes after upgrades to decide whether to manually adopt new OOTB logic.

Therefore, the correct answer is B – A skipped change is created and no change is made to the script transform definition.

Question: 19

The CMDB Configuration Management team has successfully developed a healthy and trusted CMDB. They have integrated discovered infrastructure data, accurately referenced non-discoverable data (such as change and support group information), and made the CMDB service-aware using Service Mapping.

Which field on an Incident form is automatically populated after a CI is selected that references an appropriate support group?

A. Managed by Group

B. Approval Group

C. Assignment Group

D. Change Group

E. Support Group

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In a mature CMDB implementation within ServiceNow, CI operational attributes are leveraged to automate ITSM workflows. One of the most important outcomes of accurate Configuration Management is automatic incident routing.

When a CI is selected on an Incident record, ServiceNow evaluates the CI's Support Group attribute. If populated correctly, the platform automatically copies this value into the Assignment Group field on the Incident. This ensures incidents are routed to the correct resolver group without manual triage, reducing mean time to resolution (MTTR).

The Support Group is a CI attribute, not an incident field that drives workflow directly. The Assignment Group is the operational field used by Incident Management to assign ownership. Managed by Group, Approval Group, and Change Group are used in other governance and lifecycle contexts and are not auto-populated during incident creation.

This behavior is a direct result of Data Foundations best practices: maintaining accurate CI-to- support-group relationships to enable automation and consistency across ITSM processes.

Therefore, the correct answer is C – Assignment Group.

Question: 20

(Choose 2 options)

Configuration Management requires an accurate inventory of devices to be reflected in the CMDB.

Which are common use cases for using Agent Client Collector (ACC)?

A. Servers in the data center

B. Network devices in the DMZ

C. Devices in secure environments

D. Devices that intermittently connect to the network

Answer: CD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The Agent Client Collector (ACC) in ServiceNow is designed to collect inventory data from endpoints that are not consistently reachable by traditional Discovery methods. ACC is especially valuable where credential-based, network-based discovery is impractical or impossible.

Devices in secure environments (Option C), such as isolated networks, restricted zones, or highly regulated environments, often block inbound discovery traffic. ACC runs locally on the device and securely sends inventory data outward, making it ideal for these scenarios.

Devices that intermittently connect to the network (Option D), such as laptops, remote endpoints, or roaming devices, are another core use case. Traditional Discovery requires the device to be reachable during scheduled scans, which is unreliable for mobile or off-network assets. ACC ensures inventory data is collected whenever the device is online.

Option A (data center servers) is better served by agentless Discovery, which provides deeper infrastructure and relationship data. Option B (network devices in the DMZ) are typically discovered using SNMP and network discovery, not ACC.

ACC complements Discovery as part of a layered ingestion strategy, ensuring accurate inventory coverage across diverse environments.

Therefore, the correct answers are C – Devices in secure environments and D – Devices that intermittently connect to the network.

Question: 21

(Choose 2 options)

A CMDB Administrator is evaluating whether to monitor the metrics provided on the CMDB Data Foundations Dashboard.

Which benefits support the decision to continually monitor the results on this dashboard?

A. Provides a list of all CIs that failed health audits

B. Provides metrics on active CIs updated in the last 90 days

C. Provides metrics for CIs processed by the Identification and Reconciliation Engine (IRE)

D. Reports on all orphan CIs in the CMDB

Answer: BC

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Foundations Dashboard in ServiceNow is intended to provide ongoing, trend-based visibility into how well the CMDB ingestion and maintenance processes are functioning—not just point-in-time issue lists. This is why continual monitoring of its metrics is valuable.

Option B is correct because tracking active CIs updated in the last 90 days provides a strong indicator of data freshness and operational relevance. A healthy CMDB should reflect recent updates from Discovery, integrations, and governed manual processes. Monitoring this metric over time helps organizations detect stagnation, discovery failures, or integration issues early.

Option C is also correct because metrics for CIs processed by the Identification and Reconciliation Engine (IRE) directly indicate the effectiveness and adoption of governed ingestion practices. Consistent IRE processing confirms that integrations are not bypassing identification rules, reducing duplicates and improving trust in CMDB data. Trending this metric helps validate Data Foundations maturity.

Option A is incorrect because the dashboard is not designed to provide exhaustive audit failure lists; those are handled through certification and remediation workflows. Option D is also incorrect because orphan CIs are a specific health condition surfaced via health rules and remediation tools, not a core benefit metric for continual dashboard monitoring.

Therefore, the correct answers are B and C.

Question: 22

(Choose 2 options)

A CMDB Administrator needs to create a new CI class for an Internet of Things (IoT) Sensor in ServiceNow.

What are the recommended practices for this activity?

A. Delete an unused class and replace it with the new one

B. Install or update the CMDB CI Class Models Store application and verify the class does not already exist

C. Add a new class under an appropriate parent class

D. Modify an existing class

Answer: B C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Creating new CI classes is a high-impact configuration activity and must follow strict Data Foundations and CSDM-aligned best practices to avoid long-term technical debt and upgrade risk.

Option B is a recommended first step. Before creating any new CI class, administrators should install or update the CMDB CI Class Models Store application and verify whether an appropriate class already exists. ServiceNow frequently delivers new CI classes through updates and class model packages, and duplicating an existing or planned class can lead to fragmentation and governance issues.

Option C is also correct. When a new class is truly required, it should be added under an appropriate parent class to inherit attributes, behaviors, and discovery patterns. For an IoT Sensor, this might be under a hardware or device-related parent class, ensuring consistency and minimizing customization.

Option A is incorrect and dangerous—deleting unused classes can break dependencies and historical data. Option D is also discouraged; modifying existing classes to repurpose them violates upgradesafe design principles and can negatively impact discovery, integrations, and reporting.

By verifying existing models first and extending the class hierarchy correctly, organizations maintain a clean, scalable, and upgrade-safe CMDB.

Therefore, the correct answers are B and C.

Question: 23

A new custom class is needed to reflect a new application being managed in the CMDB.

Which roles are minimally needed to add this custom CI class?

A. sn_cmdb_admin and personalize_dictionary

B. sn_cmdb_admin and personalize_form

C. sn_cmdb_admin and personalize_dictionary

D. itil_admin and personalize_form

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Creating a custom CI class in the CMDB is a dictionary-level configuration activity and must be performed with the correct minimum privileges to ensure governance and upgrade safety in ServiceNow.

The sn_cmdb_admin role is required because it grants administrative access to CMDB structures, including CI class hierarchy management. This role ensures that changes align with CMDB governance controls and Data Foundations practices.

The personalize_dictionary role is also required because adding a new CI class involves creating or extending dictionary entries (tables, attributes, inheritance). Without dictionary-level access, a user cannot define new classes or attributes in the CMDB schema.

Option A is incorrect because sn_cmdb_admin alone is insufficient without dictionary privileges. Option B and D focus on form personalization, which affects UI layout only and does not allow creation of new CI classes. Additionally, itil_admin is not the correct role for schema-level CMDB changes.

Therefore, the minimal and correct role combination is sn_cmdb_admin and personalize_dictionary, making Option C the verified answer.

Question: 24

When integrating data into the CMDB using Import Sets and Transform Maps, which type of script is added to ensure the data is processed through the Identification and Reconciliation Engine (IRE)?

A. onBefore

B. onComplete

C. onAfter

D. onStart

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

When using Import Sets and Transform Maps to ingest data into the CMDB, it is critical that records are processed through the Identification and Reconciliation Engine (IRE) to prevent duplicates and enforce source precedence. In ServiceNow, this is achieved by invoking the IRE after the transform logic has completed.

The onAfter transform script is the correct place to call the IRE API. At this stage, the transformed

data has already been mapped and prepared, allowing the IRE to correctly identify whether a CI already exists and reconcile updates according to defined rules.

The onBefore and onStart scripts execute too early—before data mapping is complete—making them unsuitable for IRE processing. The onComplete script runs after the entire import job finishes and is not intended for per-record CI identification and reconciliation.

Because Import Sets can bypass IRE if not configured correctly, using an onAfter script is a critical Data Foundations safeguard when this ingestion method is chosen.

Therefore, the correct answer is C – onAfter.

Question: 25

A customer wants recently imported server records to be automatically reclassified into more specific CMDB classes after being discovered by ServiceNow Discovery.

During the discovery process, if existing Server records are reclassified into the Linux Server and Windows Server classes, which reclassification operation occurs?

A. Class Switch

B. Class Downgrade

C. Class Upgrade

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In the CMDB class hierarchy, Server is a generic parent class, while Linux Server and Windows Server are more specific child classes. When ServiceNow Discovery detects sufficient evidence (such as OS signatures) to move a CI from a generic class to a more specialized one, this action is called a Class Upgrade.

A Class Upgrade occurs when a CI is reclassified down the hierarchy into a more specific subclass, enriching the record with additional attributes, behaviors, and discovery patterns appropriate to that class. This is a standard and expected behavior in mature CMDB implementations and aligns with Data Foundations best practices.

A Class Switch would imply lateral movement between unrelated classes, which is not what happens here. A Class Downgrade would move a CI from a specific class back to a more generic one, typically when discovery confidence is reduced—not the case in this scenario.

By performing class upgrades automatically, Discovery improves CMDB accuracy, reporting precision, and service mapping quality without manual intervention.

Therefore, the correct answer is C – Class Upgrade.

Question: 26

A CMDB Configuration Manager is reviewing the metrics on the CMDB Health Dashboard – Correctness Scorecard for the Server class, which consists of 60,000 servers in the CMDB.

For the Duplicate metric, it shows Healthy CIs / Evaluated = 59,000 / 60,000

For the Orphan metric, it shows Healthy CIs / Evaluated = 45,000 / 50,000

Which configuration explains the difference in the scope of Server CIs evaluated (60,000 vs 50,000) between the two metrics?

A. The Orphan metric has a CMDB Group configured for the Server class

B. The Orphan metric has a Health Inclusion rule configured for the Server class

C. The Duplicate metric has a Health Inclusion rule configured for the Server class

D. The Duplicate metric has a CMDB Group configured for the Server class

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, each CMDB Health metric can independently define which CIs are in scope for evaluation. This scoping is controlled primarily through Health Inclusion Rules, not CMDB Groups.

In this scenario, the Duplicate metric evaluates all 60,000 Server CIs, indicating no inclusion rule is restricting its scope. In contrast, the Orphan metric evaluates only 50,000 Server CIs, which means 10,000 servers are intentionally excluded from that metric's evaluation.

This difference is explained by a Health Inclusion rule configured specifically for the Orphan metric on the Server class. Health Inclusion rules allow administrators to define conditions—such as lifecycle state, environment, discovery source, or operational status—that determine whether a CI should be included in a specific health calculation. For example, retired servers or servers in build states may be excluded from orphan checks.

CMDB Groups are not used by the CMDB Health Engine to determine metric scope; they are used for reporting, assignment, and operational grouping. Therefore, Options A and D are incorrect. Option C is also incorrect because the Duplicate metric clearly evaluates the full population of 60,000 servers.

Thus, the scope difference is correctly explained by the Orphan metric having a Health Inclusion rule configured, making Option B the verified answer.

Question: 27

(Choose 2 options)

Which are values of CMDB?

A. Strengthening operational resiliency

B. Streamlining Incident and Change Management

C. Automating maintenance for CI relationships

D. Collecting and managing financial data

Answer: AB

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB is a foundational capability that enables organizations to operate IT services with confidence, resilience, and efficiency. Its value lies not in automation for its own sake or financial data management, but in how it supports service-aware operations and decision-making.

Strengthening operational resiliency (Option A) is a core value of the CMDB. By maintaining accurate configuration data and relationships, organizations can better understand dependencies, assess risk, and recover more quickly from incidents or outages. A trusted CMDB enables proactive problem management and informed change planning, directly contributing to resiliency.

Streamlining Incident and Change Management (Option B) is another primary value. Accurate CI data allows incidents to be routed automatically to the correct support groups, enables faster root-cause analysis, and supports risk-based change assessments. This reduces manual effort, improves response times, and lowers operational risk.

Option C is incorrect because automating CI relationship maintenance is a capability enabled by tools like Discovery and Service Mapping—not a value in itself. Option D is also incorrect because financial data management is the domain of IT Asset Management (ITAM) and Financial Management, not the CMDB's core value proposition.

In summary, the CMDB delivers value by improving operational resilience and optimizing ITSM processes, making Options A and B the correct answers.

Question: 28

A development team is working on a project where an application will be deployed to many servers. There are several security requirements that must be checked to adhere to lawful regulatory compliance because the application will be holding customer personal data (PII and PCI).

Where in the CSDM should the development team store the information that will be used to satisfy audits?

A. Technology Management Service Offerings (Technical Service Offerings) and Groups

B. Business Applications and Information Objects

C. Customer Service Offerings and Databases

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Within the Common Service Data Model (CSDM), regulatory, security, and compliance-related information—especially for PII and PCI—must be modeled at the business and information level, not at the infrastructure or service offering level.

The correct location for this data is Business Applications combined with Information Objects.

Business Applications represent the logical applications that support business capabilities and processes. Since compliance obligations (such as GDPR, PCI-DSS, or HIPAA) are assessed based on how the business uses data—not how many servers host the application—this is the correct anchor point for audit-relevant context.

Information Objects are explicitly designed to capture what data is processed, stored, or transmitted by an application, including data classifications such as PII, PCI, PHI, or confidential business data. They allow organizations to document regulatory scope, retention rules, encryption requirements, and audit controls without overloading CI records or polluting infrastructure classes.

Option A is incorrect because Technical Service Offerings and Groups focus on operational support and service delivery, not regulatory data context. Option C is also incorrect because Customer Service Offerings describe how services are consumed, while databases are technical components; neither is the authoritative place for compliance definitions.

Therefore, Business Applications and Information Objects are the correct CSDM constructs to support audits and regulatory compliance, making Option B the correct answer.

Question: 29

A Platform Data Owner wants to improve data quality with reconciliation rules across five discovery sources. The Data Owner knows the best option is to include CMDB 360 / Multisource CMDB to manage and monitor discovery sources. The company currently does not have the ITOM Discovery license required for CMDB 360 / Multisource CMDB.
What can the Data Owner do in this case?

A. ITOM Discovery must be purchased to take advantage of multisource IRE rules

B. The IRE reconciliation rules can use discovery sources regardless of CMDB 360 being enabled

C. CMDB 360 / Multisource is a platform product that can be used immediately

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):
The Identification and Reconciliation Engine (IRE) is a core platform capability in ServiceNow and does not require CMDB 360 / Multisource CMDB to function. Even without the ITOM Discovery license, organizations can still define and use IRE reconciliation rules across multiple data sources.

IRE rules are source-aware and can evaluate attributes based on source precedence, regardless of whether CMDB 360 is enabled. CMDB 360 enhances visibility, governance, and monitoring of multiple sources, but it is not a

prerequisite for reconciliation logic itself.

Option A is incorrect because purchasing ITOM Discovery is not mandatory to use multisource reconciliation. Option C is also incorrect because CMDB 360 / Multisource CMDB is a licensed addon, not a universally available platform feature.

Therefore, the Data Owner can proceed by configuring IRE reconciliation rules directly, making Option B the correct answer.

Question: 30

The CMDB Configuration Manager is using the CI Class Manager to define group ownership of CI classes and needs to leverage the ownership value specified in the CI Class Manager.

When creating a CMDB Data Manager policy, which group reference field should be set?

A. Approval Group

B. Managed By Group

C. Support Group

D. Change Group

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow CMDB governance, the CI Class Manager allows administrators to assign ownership and accountability at the class level. This ownership is used by governance tools—especially CMDB Data Manager—to automatically determine who receives tasks and actions related to data lifecycle management.

The group reference field that aligns with class ownership and is consumed by CMDB Data Manager policies is the Managed By Group. This field represents the team responsible for the technical stewardship and lifecycle management of CIs within that class.

When CMDB Data Manager executes policies such as retention, archival, or cleanup, it uses the Managed By Group to assign tasks to the appropriate data owners. This ensures governance actions are routed to the correct accountable team without manual intervention.

Approval Group, Support Group, and Change Group serve different purposes. Approval Group is used for workflow approvals, Support Group is used for operational ticket routing, and Change Group supports Change Management governance. None of these reflect data ownership at the class level.

Therefore, to leverage CI class ownership defined in the CI Class Manager within CMDB Data Manager policies, the correct field is Managed By Group, making Option B the verified answer.

Question: 31

(Choose 2 options)

A Configuration Management Process Owner is preparing solution options for presentation to technical governance for ingesting custom CIs into the CMDB. The solution must align with best practices, minimize future technical debt, and ensure upgrade compliance.

Which solutions accomplish this?

A. Repurposing a base CI class and renaming attributes as required

B. Installing or upgrading the CMDB CI Class Models Store application to find a suitable existing CI class accommodating any new attributes

C. Extending the existing Asset class table to custom CI class attributes

D. Extending an existing CI class table to accommodate the custom CI class attributes

Answer: BD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, ingesting custom CIs must be done with a strong focus on upgrade safety, governance, and long-term maintainability. Data Foundations guidance explicitly discourages repurposing or overloading base classes, as this creates technical debt and upgrade risk.

Option B is a best practice because the CMDB CI Class Models Store delivers ServiceNow-supported CI classes that align with platform evolution. Before creating or extending classes, administrators should verify whether a suitable class already exists or has been introduced in newer releases. This avoids duplication and ensures future compatibility.

Option D is also correct. When no suitable class exists, extending an existing CI class (under the appropriate parent) to add required attributes preserves inheritance, discovery behavior, reporting, and upgrade compatibility. This approach is preferred over creating entirely new, disconnected schemas.

Option A is incorrect because repurposing base classes and renaming attributes breaks standard semantics, causes confusion, and complicates upgrades. Option C is incorrect because extending Asset tables to represent CIs conflates ITAM and CMDB concerns; assets and CIs serve different purposes and lifecycles.

Therefore, the solutions that minimize technical debt and ensure upgrade compliance are B and D.

Question: 32

(Choose 2 options)

A CMDB Administrator wants to run the "Services Have Owners Identified" Get Well Playbook to remediate issues shown in the CMDB Data Foundations Dashboard.

Which remediation plays would be used?

A. Fix Data

B. Analyze Data

C. Report Data

D. Govern Data

Answer: AD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Foundations Dashboard is paired with Get Well Playbooks that guide administrators through structured remediation. The "Services Have Owners Identified" playbook focuses on closing ownership gaps for services, which is a governance and data correction activity.

Fix Data (Option A) is used to correct missing or incorrect values, such as populating owner fields, assigning responsible groups, or updating relationships. In this playbook, Fix Data actions are required to actually remediate the issue by assigning owners to services.

Govern Data (Option D) is also required because ownership is not a one-time correction—it must be enforced and sustained. Govern Data establishes policies, ownership accountability, and controls (such as certifications or attestations) to ensure services continue to have owners over time and do not regress.

Analyze Data (Option B) is used to understand patterns and root causes, but it does not remediate the issue. Report Data (Option C) provides visibility and communication, not corrective action.

Therefore, the remediation plays that apply to the Services Have Owners Identified playbook are Fix Data and Govern Data, making Options A and D correct.

Question: 33

A Platform Owner is collaborating with stakeholders in the manufacturing industry to align their CIs with the CSDM 5 framework. They need to map production line monitoring systems to the appropriate CSDM domain.

Which CSDM 5 domain does the Platform Owner use?

A. Build and Integration (Build)

B. Foundation

C. Service Consumption (Sell/Consume)

D. Design and Planning (Design)

E. Service Delivery (Manage Technical)

Answer: E

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In CSDM 5, production line monitoring systems (such as SCADA, MES, IoT telemetry platforms, and operational monitoring tools) are technical systems responsible for operating, monitoring, and supporting services, not for designing them or consuming them. These systems directly align with

the Service Delivery (Manage Technical) domain.

The Service Delivery domain is used to model how services are technically delivered and operated, including the infrastructure, platforms, and operational technologies that ensure availability, performance, and reliability. In a manufacturing context, production line monitoring systems continuously observe equipment health, throughput, alerts, and operational metrics—making them part of the technical service delivery layer.

Option A (Build and Integration) applies to CI/CD pipelines and system construction activities. Option B (Foundation) focuses on base CIs such as locations, people, and organizations. Option C (Service Consumption) models how customers or consumers use services, which is not applicable here.

Option D (Design and Planning) is used for service architecture and planning artifacts, not live operational systems.

Therefore, production line monitoring systems correctly belong in Service Delivery (Manage Technical), making Option E the correct answer.

Question: 34

A CMDB Administrator has been tasked with gathering information for a presentation to leadership. The Administrator needs to provide Duplicate CI, Orphan CI, and Stale CI metrics.

Which scorecard provides this information on the CMDB Health Dashboard?

A. Compliance

B. Correctness

C. Completeness

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

On the CMDB Health Dashboard in ServiceNow, health metrics are grouped into three primary scorecards: Completeness, Correctness, and Compliance. Each scorecard focuses on a distinct aspect of data quality.

Duplicate CIs, Orphan CIs, and Stale CIs are all indicators of data accuracy and reliability, which fall under the Correctness scorecard.

Duplicate CIs indicate multiple records representing the same real-world item.

Orphan CIs are missing required relationships.

Stale CIs have not been updated within an expected timeframe.

All three conditions reflect whether the CMDB data is correct and trustworthy, not whether it is complete or policy-compliant.

The Completeness scorecard focuses on missing required attributes or relationships. The Compliance scorecard evaluates adherence to policies such as certifications, lifecycle rules, or patch compliance.

Since leadership reporting typically focuses on trust and accuracy of CMDB data, the Correctness scorecard is the authoritative source for these metrics.

Therefore, the correct answer is B – Correctness.

Question: 35

Which type of CMDB Data Manager policy creates tasks that allow the assigned individual to update fields on the CI record?

A. Audit

B. Certification

C. Attestation

D. Compliance

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In CMDB governance, different CMDB Data Manager policy types serve different validation and enforcement purposes.

When the objective is to allow an assigned individual to review and update specific fields on a CI record, the correct policy type is Certification.

A Certification policy creates actionable tasks that require the assignee to validate and, if necessary, correct specific CI attributes, such as lifecycle status, support group, environment, or ownership. During certification, the user can directly update CI fields to bring the record into compliance with defined standards.

Attestation (Option C) only asks the user to confirm that a CI still exists or is still valid; it does not require or enable attribute-level updates. Audit (Option A) is used for reporting and evidence collection, not remediation. Compliance (Option D) measures adherence to rules but does not itself generate editable remediation tasks.

Certification is therefore the primary mechanism used when human validation and correction of CI data is required—making it a cornerstone of CMDB data quality management.

Hence, the correct answer is B – Certification.

Question: 36

A healthcare provider faces a critical incident affecting its patient management system. The provider needs to identify the users impacted to mitigate disruption effectively.

Which CSDM-related data should they leverage?

A. Incident history of similar CIs

B. Service Offerings by Department or Location

C. Service environment attribute

D. Affected CI [task_ci] related list

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In a healthcare environment, identifying who is impacted during a critical incident is essential to patient safety and continuity of care. Within the Common Service Data Model (CSDM), the most effective way to determine impacted users is through Service Offerings, particularly when they are defined by department or location.

Service Offerings represent how a service is consumed by specific user groups. In this case, a patient management system may have different offerings for departments such as Emergency, Inpatient Care, or Outpatient Services, or be scoped by hospital location. These offerings explicitly define consumer context, allowing incident responders to immediately identify which clinicians, staff, or facilities are affected.

Option D (Affected CI related list) identifies technical impact but does not translate that impact into user or consumer context. Option A provides historical insight but does not identify current impacted users. Option C (service environment) helps differentiate production vs non-production but does not identify who is impacted.

By leveraging Service Offerings by Department or Location, the provider can quickly notify the right users, prioritize response based on clinical impact, and coordinate mitigation effectively—aligning with CSDM and ITIL best practices.

Therefore, the correct answer is B – Service Offerings by Department or Location.

Question: 37

A Service Owner is using Unified Map to understand the composition of a service but wants to filter irrelevant information.

Which options are available to the Service Owner from the filter panel? (Choose 2 options)

A. CI type

B. Discovery source

C. Managed by group

D. Business criticality

Answer: A D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The Unified Map in ServiceNow provides a consolidated view of services and their underlying components, integrating Discovery and Service Mapping data. To make this view actionable, Service Owners can apply filters to focus on relevant elements and reduce visual noise.

Filtering by CI type (Option A) is a core capability. It allows the Service Owner to show or hide categories such as servers, databases, load balancers, or applications—making it easier to analyze specific layers of the service.

Filtering by Business Criticality (Option D) is also available and highly valuable. This enables Service Owners to prioritize views around high-impact components, ensuring attention is focused on CIs that pose the greatest risk to service delivery.

Option B (Discovery source) is not typically exposed as a Unified Map filter because the map focuses on operational and service context, not ingestion provenance. Option C (Managed by group) is a governance attribute and is not a standard visual filter within the Unified Map.

Thus, the correct filter options are A – CI type and D – Business criticality.

Question: 38

(Choose 2 options)

Which ServiceNow solutions create automatic relationships?

A. IntegrationHub ETL

B. Service Mapping

C. Discovery

D. Workflow Studio

Answer: BC

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Automatic relationship creation is fundamental to maintaining a service-aware and trustworthy CMDB. In ServiceNow, this capability is primarily delivered by Discovery and Service Mapping.

Discovery (Option C) automatically identifies infrastructure components—such as servers, network devices, and storage—and creates technical relationships between them (for example, “runs on,” “connected to,” or “depends on”). These relationships form the backbone of infrastructure dependency mapping.

Service Mapping (Option B) builds on Discovery by creating application- and service-level relationships. It maps how application components interact across servers, databases, and middleware, resulting in accurate Application Service models aligned with CSDM. These relationships are created and maintained automatically as the environment changes.

Option A (IntegrationHub ETL) focuses on data ingestion and transformation; it does not inherently create or maintain relationships unless explicitly scripted. Option D (Workflow Studio) orchestrates processes and automations but does not discover or infer CI relationships.

Therefore, the ServiceNow solutions that create automatic relationships are Service Mapping and Discovery, making Options B and C correct.

Question: 39

(Choose 2 options)

A CMDB Administrator wants to create a CMDB query to find all databases located in Seattle that are connected to application services. They also want to include incidents related to those databases.

Which actions should be taken to build this query?

A. Add to the canvas the Incident table from the Non-CMDB Tables list

B. Add property columns to the Application Service node

C. Add a filter to the Database node for Location = Seattle

D. Set the relationship level to up to 2nd-level relationships

Answer: AC

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

When building advanced CMDB queries using CMDB Query Builder in ServiceNow, the correct approach is to model CI scope, relationships, and task context explicitly on the canvas.

To limit results to databases in a specific location, the administrator must filter the Database CI node by the Location attribute. Therefore, Option C is required to scope the query to Database CIs where Location = Seattle.

To include Incidents related to those databases, the Incident table must be added from the Non- CMDB Tables list and linked through the task_ci relationship. This is exactly what Option A provides. CMDB Query Builder separates CMDB tables (CIs) from task and transactional tables, so incidents must be explicitly added from the Non-CMDB section.

Option B is incorrect because property columns on Application Services do not scope databases or incidents. Option D is unnecessary because relationship depth alone does not include non-CMDB task data and does not filter by location.

Thus, the correct actions are A (add Incident table) and C (filter Database by location).

Question: 40

A CMDB Data Manager needs to access the ServiceNow platform to create, publish, and manage policies that automate and govern CI lifecycle operations, ensuring the CMDB remains healthy and efficient.

Where can the Data Manager do this?

A. CMDB Workspace – CMDB 360 tab

B. Service Operations Workspace

C. CI Class Manager

D. CMDB Workspace – Management tab

Answer: D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The CMDB Data Manager performs governance activities such as creating, publishing, and managing lifecycle policies (archival, certification, attestation, cleanup) to ensure long-term CMDB health.
These activities are executed within the CMDB Workspace, specifically under the Management tab.

In ServiceNow, the CMDB Workspace – Management tab is the centralized location for CMDB governance operations. From here, Data Managers can define policy logic, assign ownership, schedule execution, monitor outcomes, and manage remediation tasks generated by those policies.

Option A (CMDB 360 tab) focuses on visibility and analysis of CI data and relationships, not policy authoring. Option B (Service Operations Workspace) is used for operational response and service monitoring, not CMDB governance. Option C (CI Class Manager) is used to define class hierarchy and ownership, but it does not manage lifecycle policies.

Therefore, the correct location for CMDB Data Manager policy management is CMDB Workspace – Management tab, making Option D correct.

Question: 41

A CMDB Administrator wants only the CIs of Principal Classes to appear in CI reference fields, such as the CI reference field on an Incident form.

Where does the CMDB Administrator designate Principal Classes?

A. CMDB Workspace
B. CI Class Manager
C. System Properties
D. CMDB Data Manager

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Principal Classes are a key CMDB configuration concept used to control which CI classes are selectable in reference fields across ITSM processes. This helps reduce noise, prevent incorrect CI selection, and improve data quality.

In ServiceNow, Principal Classes are designated within the CI Class Manager. This tool allows CMDB administrators to manage the CI class hierarchy, define ownership, and explicitly mark classes as principal. Once a class is marked as principal, its CIs become available in CI reference fields such as those on Incident, Change, and Problem forms.

Option A (CMDB Workspace) provides operational and analytical views but does not control schemalevel class behavior. Option C (System Properties) does not manage class designation. Option D (CMDB Data Manager) governs lifecycle and data quality policies, not reference field behavior.

By configuring Principal Classes in the CI Class Manager, organizations ensure that only relevant, high-value CI classes are exposed to end users, aligning with Data Foundations best practices.

Therefore, the correct answer is B – CI Class Manager.

Question: 42

A Configuration Management Governance team is transitioning from utilizing legacy CMDB status fields to CSDM lifecycle status fields.

Which table can be modified?

A. Life Cycle Stages

B. Life Cycle Mapping

C. Life Cycle Stage Status

D. Life Cycle Controls

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

When organizations transition from legacy CMDB status fields (such as custom install status or operational status values) to CSDM-aligned lifecycle status fields, the goal is to map old values to standardized lifecycle stages without disrupting existing processes. In ServiceNow, this is achieved through the Life Cycle Mapping table.

The Life Cycle Mapping table is specifically designed to translate legacy or custom status values into CSDM lifecycle stages and statuses. This allows organizations to preserve historical data and integrations while progressively adopting CSDM standards. By modifying this table, administrators can define how existing status values correspond to CSDM lifecycle stages such as Plan, Build, Deploy, Operate, and Retire.

The Life Cycle Stages table (Option A) defines the standard stages themselves and should not be

modified, as these are core to CSDM governance. Life Cycle Stage Status (Option C) defines valid statuses within a stage and is also part of the standardized model. Life Cycle Controls (Option D) enforce governance rules but do not perform value translation.

Therefore, to safely transition from legacy status fields to CSDM lifecycle statuses, the correct and supported approach is to modify the Life Cycle Mapping table, making Option B the correct answer.

Question: 43

(Choose 2 options)

A Change Manager wants to gain value from CSDM.

How will the Change Management process benefit from CSDM?

A. Identify blackout windows

B. Determine the root cause of the change issue

C. Route the change dynamically

D. Understand the impact of the change on services

Answer: AD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

CSDM significantly enhances Change Management by providing service-aware context, enabling better planning, risk assessment, and stakeholder communication.

One key benefit is the ability to identify blackout windows (Option A). Through CSDM-aligned Business Services, Service Offerings, and service calendars, Change Managers can clearly see when services are unavailable for change due to business constraints, regulatory requirements, or peak usage periods. This helps prevent changes from being scheduled during high-risk windows.

Another critical benefit is the ability to understand the impact of the change on services (Option D). CSDM establishes clear relationships between infrastructure CIs, Application Services, and Business Services. When a change is proposed, these relationships enable accurate impact analysis, allowing Change Managers to assess risk based on business criticality rather than just technical scope.

Option B (root cause determination) is primarily a Problem Management function. Option C (dynamic routing of changes) is driven by workflow and approval logic, not directly by CSDM.

Therefore, the correct answers are A – Identify blackout windows and D – Understand the impact of the change on services.

Question: 44

A CMDB Administrator wants to remove all Linux Servers in the organization that have not been updated in six months.

Which recommended action should the Administrator take in Data Foundations?

A. Create a business rule

B. Create an archive policy

C. Create a scheduled job

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

Removing obsolete or inactive CIs from the CMDB must be handled carefully to avoid data loss, audit issues, and unintended operational impact. In ServiceNow, the recommended and governed approach is to use an archive policy.

Archive policies are designed to manage CI lifecycle cleanup based on defined conditions such as class, last updated date, lifecycle status, or operational state. In this scenario, the condition would target Linux Server CIs that have not been updated in six months. Archive policies can either archive or permanently delete records in a controlled, auditable manner, ensuring compliance with data retention and governance standards.

Creating a business rule (Option A) is strongly discouraged for bulk CMDB cleanup because it introduces technical debt, upgrade risk, and unpredictable side effects. A scheduled job (Option C) may automate execution but lacks governance logic and lifecycle awareness on its own.

Archive policies integrate with CMDB Data Manager, provide visibility into actions taken, and support approval and rollback where appropriate. This aligns fully with Data Foundations best practices for maintaining a lean, accurate, and trusted CMDB.

Therefore, the correct and recommended action is B – Create an archive policy.

Question: 45

A Configuration Manager is planning the implementation of the CMDB.

Which is the prescribed CSDM rollout order?

A. Initial, Developing, Defined, Managed

B. Architecture, Business, Technical, Governance

C. Initiate, Plan, Execute, Deliver, Close

D. Crawl, Walk, Run, Fly

Answer: D

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The Common Service Data Model (CSDM) prescribes an incremental, maturity-based rollout approach to reduce risk and ensure sustainable adoption. The recommended order is Crawl, Walk, Run, Fly, which aligns implementation effort with increasing organizational capability and value realization.

Crawl focuses on foundational data hygiene: core CI classes, identification rules, reconciliation, basic Discovery ingestion, and CMDB Health basics.

Walk introduces service context, including Business Services, Application Services, and relationships that enable impact analysis for Incident and Change.

Run expands into operational excellence with Service Mapping, service offerings, advanced governance, and process automation.

Fly represents optimization and scale, leveraging analytics, AI/ML, proactive operations, and cross-domain integration (e.g., SecOps, APM, CSM).

This progression ensures teams do not over-model early or introduce complexity before data quality and governance are established. The other options describe generic project lifecycles or organizational categorizations, not the CSDM-recommended adoption path.

Therefore, the correct answer is D – Crawl, Walk, Run, Fly.

Question: 46

According to the Common Service Data Model (CSDM), a server team is requesting a catalog item be created for infrastructure requests.

Which role is involved in initiating the request and defining requirements?

A. Application Service Owners

B. Technology Service Owners

C. Enterprise Architect

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In CSDM, Technology Services (and their Service Offerings) represent how technical capabilities are delivered and consumed by internal teams. When a server team requests a catalog item for infrastructure services (e.g., VM provisioning, storage, OS builds), the role responsible for initiating the request and defining requirements is the Technology Service Owner.

Technology Service Owners understand the operational capabilities, constraints, SLAs, and fulfillment workflows required to deliver infrastructure services. They define catalog requirements such as options, approvals, fulfillment tasks, and guardrails—ensuring the request aligns with standardization, security, and operational readiness.

Application Service Owners focus on how applications are delivered and supported, not on defining infrastructure catalog items. Enterprise Architects provide standards and guidance but do not initiate or define catalog request requirements.

Thus, the correct role is B – Technology Service Owners.

Question: 47

A CMDB Administrator is reviewing the CMDB and notices that many Hardware CIs are missing serial numbers. The Administrator is concerned this may cause duplicate CIs and wants to resolve the issue quickly.

What structured guidelines provided by ServiceNow are available to troubleshoot and resolve the issue?

A. CMDB Data Foundations Dashboard Playbooks

B. CSDM Data Foundations Dashboard Playbooks

C. CMDB Health Dashboard Playbooks

D. CSDM Now Create Playbooks

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

When data quality issues such as missing serial numbers threaten CMDB integrity and increase the risk of duplicates, ServiceNow provides prescriptive, step-by-step remediation guidance through the CMDB Data Foundations Dashboard Playbooks.

These playbooks are specifically designed to help administrators identify root causes, assess ingestion and governance gaps, and apply corrective actions using structured remediation plays (Analyze Data, Fix Data, Govern Data). For missing serial numbers, the playbooks guide teams to review Discovery patterns, identification rules, reconciliation sources, and governance controls to ensure authoritative data capture and prevention of future issues.

The CMDB Health Dashboard Playbooks focus on health scoring and metrics, not guided remediation. CSDM Data Foundations Dashboard Playbooks is not a distinct product naming; the correct construct is CMDB Data Foundations. Now Create Playbooks provide implementation project guidance, not operational troubleshooting for live data issues.

Therefore, the correct answer is A – CMDB Data Foundations Dashboard Playbooks, which are purpose-built to quickly

troubleshoot and remediate CMDB data quality problems while aligning with best practices in ServiceNow.

Question: 48

A CMDB Administrator wants to improve data quality related to the CSDM.

Which action should the Administrator take to meet this goal?

A. Use the CSDM Data Foundations Dashboard

B. Start the ServiceNow Health Scan

C. Use the default configured CMDB Health Dashboard

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

To specifically improve data quality related to CSDM, the most effective and prescribed action is to use the CSDM Data Foundations Dashboard. In ServiceNow, this dashboard is purpose-built to assess and improve CSDM alignment, not just general CMDB hygiene.

The CSDM Data Foundations Dashboard focuses on service modeling readiness, highlighting gaps such as missing service ownership, incomplete relationships between Business Applications and

Application Services, unmanaged services, and misaligned lifecycle states. It provides Get Well Playbooks that guide administrators through structured remediation using Analyze Data, Fix Data, and Govern Data plays—directly tied to CSDM outcomes.

Option C (default CMDB Health Dashboard) is valuable, but it measures generic CMDB data quality dimensions (completeness, correctness, compliance) and does not specifically evaluate CSDM constructs or service modeling maturity. Option B (ServiceNow Health Scan) provides platform-wide configuration and performance recommendations, but it is not focused on CMDB or CSDM data quality.

Therefore, to improve CSDM-specific data quality, the administrator should use the CSDM Data Foundations Dashboard, making Option A the correct answer.

Question: 49

A Configuration Manager wants to use the Unified Map.

Where would it be accessed?

A. CI Class Manager

B. CMDB Data Manager

C. CMDB Workspace

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The Unified Map is a visualization capability used to understand service composition, dependencies, and relationships across CIs, Application Services, and infrastructure. In ServiceNow, the Unified Map is accessed through the CMDB Workspace.

CMDB Workspace serves as the central experience for CMDB operations, analytics, and visualization. From within the workspace, users can launch the Unified Map to explore how services are constructed, identify dependencies, and analyze impact—leveraging data from Discovery and Service Mapping.

Option A (CI Class Manager) is used for class hierarchy, ownership, and principal class configuration, not visualization. Option B (CMDB Data Manager) is focused on governance and lifecycle policies, such as archival, certification, and attestation, not service mapping views.

Because Unified Map is an operational and analytical visualization tool, it is correctly accessed via CMDB Workspace.

Therefore, the correct answer is C – CMDB Workspace.

Question: 50

A Data Center Manager is working with the CMDB CI Class Manager to define the relationship between Application Servers and the Applications they host. The company has multiple Application Servers that host one or more Applications.

Which describes the relationship between the Application Server table and the Application table?

A. Many-to-many

B. Many-to-one

C. One-to-one

D. One-to-many

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In CMDB modeling, accurately defining relationships is critical for impact analysis, service mapping, and Change Management. In this scenario, Application Servers can host multiple Applications, and Applications can also run across multiple Application Servers (for example, in clustered, load- balanced, or distributed architectures).

This architectural reality defines a many-to-many relationship between the Application Server table and the Application table.

In ServiceNow, many-to-many relationships are common for application hosting models, especially in modern environments that use horizontal scaling, redundancy, or containerized workloads. Modeling this correctly ensures that incidents, changes, and outages affecting a single server can be accurately traced to all impacted applications—and vice versa.

A one-to-many or many-to-one relationship would incorrectly assume exclusivity in one direction, which does not reflect real-world application deployment patterns. A one-to-one relationship would be even more restrictive and inaccurate.

Therefore, the correct relationship type is A – Many-to-many, which aligns with CMDB best practices and CSDM service modeling principles.

Question: 51

A Configuration Manager has configured multiple data sources that are all authorized to update the same class and the same set of class attributes in the CMDB.

What can the Configuration Manager do to control which data source should be the authoritative source of truth for a specific class or set of class attributes?

A. Assign a priority to each data source in reconciliation rules

B. Manually run the data source updates in the correct order

C. Configure data refresh rules with a specific time period

D. Assign a run order to each data source in the identification rules

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, controlling source precedence when multiple authorized data sources update the same CI attributes is a core responsibility of Identification and Reconciliation Engine (IRE) governance.

The correct and supported method is to assign priority to each data source in reconciliation rules. Reconciliation rules determine which source wins when multiple sources attempt to update the same attribute on a CI. By defining source precedence, the Configuration Manager ensures that the most authoritative system of record (for example, Discovery over manual imports, or HR systems over spreadsheets) consistently controls specific attributes or classes.

Option B is incorrect because manually sequencing data source runs is unreliable, does not scale, and violates Data

Foundations best practices. Option C only controls how often data is refreshed, not which source is authoritative. Option D is incorrect because identification rules are used to uniquely identify CIs—not to control attribute-level precedence.

Using reconciliation rules provides deterministic, auditable, and automated control, which is essential for maintaining CMDB trust and avoiding data conflicts.

Therefore, the correct answer is A – Assign a priority to each data source in reconciliation rules.

Question: 52

With CMDB 360 / Multisource CMDB, Dynamic Reconciliation Rules are enabled. Based on management requirements, a CMDB Administrator needs to configure multiple Dynamic Reconciliation Rules.

Which are available Dynamic Rule Types within the Create Reconciliation Rule wizard? (Choose 2 options)

A. Smallest Value

B. Most Reported

C. Last Updated

D. Last Created

Answer: BC

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

CMDB 360 / Multisource CMDB extends the standard IRE by enabling Dynamic Reconciliation Rules, which determine attribute values dynamically based on incoming data patterns rather than fixed source priority.

Within the Create Reconciliation Rule wizard, two supported dynamic rule types are:

Most Reported (Option B): selects the attribute value that is reported most frequently across all sources. This is useful when multiple sources contribute data and consensus is a strong indicator of correctness.

Last Updated (Option C): selects the most recently updated value, which is useful for rapidly changing attributes such as IP address or operational state.

Option A (Smallest Value) and Option D (Last Created) are not supported dynamic reconciliation rule types in ServiceNow.

Dynamic reconciliation rules are particularly valuable in complex, multisource environments where rigid source

precedence is insufficient and data confidence must be inferred.

Therefore, the correct answers are B – Most Reported and C – Last Updated.

Question: 53

A Change Manager aims to streamline ITSM processes by automatically populating fields on the Change form when a CI is selected. The Configuration Management team ensures that the Change Group field is populated for all managed CIs.

As a result, which base system field on the Change form will be automatically populated after selecting a CI?

A. Change Group

B. Assignment Group

C. Managed by Group

D. Approval Group

Answer: A

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In a mature Configuration Management implementation within ServiceNow, CI operational attributes are leveraged to automate Change Management workflows and reduce manual effort.

When a CI is selected on a Change record, ServiceNow evaluates the CI's Change Group attribute. If this field is populated on the CI, the platform automatically copies its value into the Change Group field on the Change form. This ensures that change ownership and governance are immediately aligned with the responsible technical team.

The Change Group is distinct from the Assignment Group, which is used primarily in Incident and Task routing. Managed by Group represents lifecycle ownership and is used by CMDB governance tools, while Approval Group controls approval workflows but is not auto-populated from CI selection.

This behavior demonstrates the value of accurate CI attributes: once populated consistently, they enable automatic field population, reduced manual errors, and faster processing across ITSM workflows.

Therefore, the correct answer is A – Change Group.

Question: 54

The Incident Process Owner asks which classes of CSDM are used on the Incident form. Which classes are

appropriate?

A. Application Service

B. Business Application

C. Service Offering

D. Service Portfolio

Answer: A, C

Explanation:

In the Common Service Data Model (CSDM), the Incident form is designed to capture operational impact and enable effective incident routing, prioritization, and communication. To achieve this, CSDM prescribes using classes that represent how services are delivered and consumed, not how they are planned or governed.

Application Service (Option A) is an appropriate class on the Incident form because it represents the technical service that is running in production and is directly impacted during an incident.
Application Services are service-mapped, relate to underlying infrastructure, and support impact analysis, root cause investigation, and automated assignment. This makes them ideal for associating incidents with technical outages or degradations.

Service Offering (Option C) is also appropriate because it represents how a service is consumed by users or business units. Service Offerings allow Incident Management to understand who is affected, enable targeted communications, and support SLA/OLA alignment. For example, an email service offering for a specific department clearly identifies the impacted consumer group.

Business Application (Option B) is not recommended on the Incident form. Business Applications are logical representations used for portfolio, ownership, and governance purposes, not day-to-day operational incident handling. Using them directly on incidents can reduce precision and automation.

Service Portfolio (Option D) is a strategic construct used for service lifecycle management and is never associated with operational incidents.

Therefore, according to CSDM best practices, the correct classes used on the Incident form are Application Service and Service Offering, making Options A and C the correct answers.

Question: 55

An organization is updating the CMDB to include new asset types like loT devices Relevant CI classes need to be added and outdated ones need to be removed from the Principal Class filler to ensure accurate display in ITSM processes.

Which roles are needed to add or remove classes?

A. personalize.dictionary

B. sn_cmdb_admin

C. sn_csdm_admin

D. cmdb_query_ builder

Answer: A, B

Explanation:

Managing CI classes and Principal Class designation is a schema-level CMDB activity that directly impacts how CIs appear in ITSM processes such as Incident, Change, and Problem. In ServiceNow, these activities require specific administrative privileges to ensure governance, security, and upgrade safety.

The sn_cmdb_admin role is required because it provides administrative access to CMDB structures, including CI class hierarchy management, Principal Class configuration, and overall CMDB governance. Without this role, users cannot add, remove, or govern CI classes effectively.

The personalize_dictionary role is also required because adding or removing CI classes involves dictionary-level changes. CI classes are implemented as tables that extend the CMDB schema, and modifying the Principal Class filter relies on dictionary metadata. This role grants permission to create, modify, or remove class definitions safely.

The sn_csdm_admin role focuses on managing CSDM constructs (domains, services, lifecycle alignment) but does not grant dictionary or schema modification rights. The cmdb_query_builder role is used only for querying and reporting and does not allow structural changes.

Therefore, the two required roles are personalize.dictionary and sn_cmdb_admin, making Options A and B correct.

Question: 56

A ServiceNow Administrator wants to implement Sen/ice Graph Connectors to provide integrations to many third-party solutions that the company wants integrated into the CMDB

Which categories of connectors are available to the Administrator"?

A. Cloud

B. Observability

C. DevOps

D. Workflow Automation

Answer: A, B

Explanation:

Service Graph Connectors are a key Data Foundations ingestion capability in ServiceNow. They provide out-of-the-box, upgrade-safe integrations that ingest data into the CMDB using the Identification and Reconciliation Engine (IRE),

ensuring data quality and source governance.

Two primary categories of Service Graph Connectors are:

Cloud (Option A): These connectors integrate with major cloud providers and platforms (such as AWS, Azure, and GCP) to ingest infrastructure, platform, and service data into the CMDB. They are essential for managing hybrid and multi-cloud environments and maintaining accurate cloud CI relationships.

Observability (Option B): These connectors integrate with monitoring and observability tools (such as APM, infrastructure monitoring, and telemetry platforms). They provide near-real-time operational data that enriches CI records and supports incident correlation, impact analysis, and service health insights.

Option C (DevOps) is incorrect because DevOps integrations are typically handled through CI/CD tools and workflow integrations rather than Service Graph Connectors. Option D (Workflow Automation) is unrelated; Service Graph Connectors focus on data ingestion, not process orchestration.

Therefore, the correct connector categories are Cloud and Observability, making Options A and B the correct answers.

Question: 57

A CMDB Administrator is using the Duplicate CI Remediator to address a de-duplication task. On the first tab of the wizard, the Main CI is selected.

Which attributes are used to identify the Main CI? (Choose two)

A. Oldest Created

B. Most Related Items

C. Newest Created

D. Least Related Items

Answer: A B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, the Duplicate CI Remediator is a governed tool designed to safely resolve duplicate Configuration Items while preserving the most valuable and authoritative record. The first step in the remediation wizard is identifying the Main CI, which will be retained after remediation.

ServiceNow uses two primary attributes to help determine the best candidate for the Main CI:

Oldest Created (Option A)

The oldest CI is often preferred because it typically has a longer operational history, may be referenced by historical incidents, changes, problems, or audits, and is more likely to be embedded in downstream processes and reports.

Retaining the oldest CI helps avoid breaking historical references.

Most Related Items (Option B)

A CI with the most relationships (for example, links to applications, services, incidents, or other CIs) is generally the most valuable from a business and operational context perspective. Preserving these relationships is critical for impact analysis, Change Management, and CSDM-aligned service modeling.

Options C (Newest Created) and D (Least Related Items) are not used as selection criteria because newer or weakly-related CIs typically contain less historical and relational value and are better candidates for removal or merging.

By prioritizing Oldest Created and Most Related Items, the Duplicate CI Remediator aligns with CMDB Data Foundations best practices, ensuring minimal data loss, preserved business context, and safer de-duplication outcomes.

Therefore, the correct answers are A and B.

Question: 58

The following identification rule for a Hardware CI class has been defined

Hardware Rule: Identifier Entries Table	Criterion attributes	Priority
Serial Number	serial_number, serial_number_type	100
Hardware	serial_number	200
Hardware	name	300
Network Adapter	mac_address, name	400

Two new CI records are imported into the Hardware class of the CMDB:

CI1: The name of this CI record matches the name of an existing CI record in the CMDB.

CI2: The IP address of this CI record matches the IP address of an existing CI record in the CMDB.

Which is correct based on the identification rule and the imported CI records?

A. CI1 will be inserted as a new record and CI2 will be updated with the matching record
B. CI1 and CI2 both will be inserted as new records
C. CI1 will be updated with the matching record and CI2 will be inserted as a new record
D. CI1 and CI2 both will be updated with matching records

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

This question tests understanding of how the Identification and Reconciliation Engine (IRE) evaluates incoming CI data against Identification Rules and their priority order in ServiceNow.

From the identification rule shown:

Serial number (+ type) → Priority 100

Serial number → Priority 200

Name (Hardware) → Priority 300

MAC address + name (Network Adapter) → Priority 400

For a CI to be identified and matched, the incoming record must satisfy one complete identifier entry exactly as defined for that class.

CI1 (Name match only)

Although the name matches an existing Hardware CI, name alone is a low-priority identifier (300) and is not sufficient to uniquely identify a Hardware CI unless no higher-priority identifiers exist and the identifier entry criteria are fully satisfied. In practice, Hardware identification relies on serial number–based identifiers, not name-only matching, to avoid false positives. Therefore, CI1 cannot be confidently matched and is inserted as a new record.

CI2 (IP address match)

IP address is not part of any Hardware identification rule shown. IP address is typically used for discovery correlation or network relationships, not as a primary Hardware identifier. Since no identifier entry includes IP address, CI2 does not match any valid identification rule and is also inserted as a new record.

Because neither CI satisfies a valid identifier entry, both records are inserted as new CIs.

Question: 59

DRAG DROP

A CMDB Administrator needs to set up CMDB 360/Multisource CMDB Drag and drop the system property to the functionality Some options may not apply.

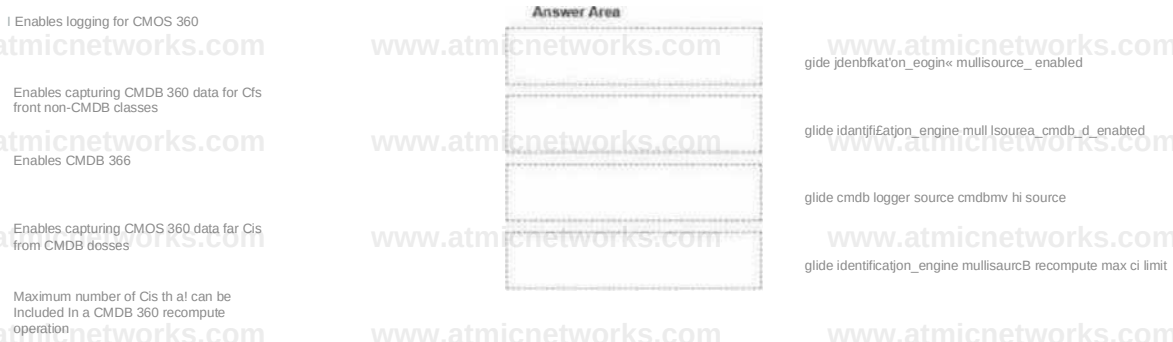

Answer:

Functionality	System Property
Enables CMDB 360	glide. ident ification_engine .multisource_enabled
Enables capturing CMDB 360 data for Cis from CMDB classes	glide.identification_engine.multisource_cmdb_ci_enable d
Enables logging for CMDB 360	glide.cmdb.logger.source.cmdb_multisource
Maximum number of Cisthat can be included in a CMDB 360 recompute operation	glide.identification_engine,multisource_recompute,max_c i limit

In ServiceNow, CMDB 360 / Multisource CMDB extends the Identification and Reconciliation Engine (IRE) to support multi-source visibility, reconciliation, and confidence scoring. Configuration is controlled through system properties, each enabling a specific capability.

glide.identification_engine.multisource_enabled is the master switch that activates CMDB 360 functionality.

glide.identification_engine.multisource_cmdb_ci_enabled allows CMDB 360 to collect and evaluate data specifically from CMDB CI classes.

glide.cmdb.logger.source.cmdb_multisource enables detailed logging for troubleshooting multisource ingestion and reconciliation.

glide.identification_engine.multisource_recompute.max_ci_limit protects performance by limiting how many CIs can be processed during a recompute operation.

Capturing CMDB 360 data from non-CMDB classes is not supported, which is why no property applies to that option.

This mapping aligns with Data Foundations and CMDB governance best practices, ensuring CMDB 360 is enabled safely, transparently, and at scale.

Question: 60

DRAG DROP

ACMDB Administrator seeks to understand the available tools for preventing addressing, and remediating duplicate CIs

Drag and drop each feature with the corresponding outcome Some options may not apply.

Answer:

Explanation:

Feature	Corresponding Outcome
Certification Tasks	Can be assigned to groups for resolving duplicate Cis
CMDB Health Dashboard-Correctness Scorecard	Offers insight into duplicate Cis Within the CMDB
De-duplication Templates	Offers a solution to resolve de-duplication tasks **in** bulk
Duplicate CI Remediator	Provides a wizard to resolve de-duplication tasks individually

In ServiceNow, duplicate CI management spans detection, insight, and remediation:

The CMDB Health Dashboard – Correctness surfaces where duplicates exist and their impact, providing visibility and prioritization.

Certification Tasks support governed resolution, allowing assignment to individuals or groups to validate and correct data.

De-duplication Templates enable bulk remediation, ideal when standardized merge rules can be applied across many duplicates.

The Duplicate CI Remediator offers a guided, record-by-record wizard, preserving relationships and history for safer individual resolutions.

De-duplication Tasks are the work items generated by these processes, but they do not themselves represent an outcome—hence not applicable.

This mapping aligns with CMDB Data Foundations best practices, ensuring duplicates are identified, prioritized, and remediated efficiently and safely.

Question: 61

The CMDB Configuration Management team wants to manage de-duplication tasks generated from data ingested into the CMDB via the Identification and Reconciliation Engine (IRE).

In which area of the CMDB Workspace can they locate these de-duplication tasks?

A. Important actions tile under the Home tab
B. CMDB feature adoption tile under the Insights tab
C. Total status under the My Work tab

Answer: C

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

In ServiceNow, de-duplication tasks generated by the Identification and Reconciliation Engine (IRE) are operational governance tasks that require action by CMDB administrators or data owners. These tasks are surfaced in a role-based and actionable manner to ensure they are addressed promptly.

Within the CMDB Workspace, such tasks are accessed via the My Work tab, specifically under the Total status section. This area consolidates all actionable CMDB-related work items assigned to the user or their groups, including duplicate CI remediation tasks, data certification tasks, attestation requests, and lifecycle governance actions generated by CMDB Data Manager and IRE processes.

Option A is incorrect because the Home tab focuses on high-level navigation and featured actions, not task execution. Option B is also incorrect because the Insights tab and feature adoption tiles provide visibility and analytics, not task management.

By centralizing de-duplication tasks in My Work, ServiceNow ensures CMDB governance is embedded into daily operations, improving responsiveness and data quality while maintaining accountability.

Therefore, the correct answer is C – Total status under the My Work tab.

Question: 62

(Choose two)

A CMDB Administrator is leveraging CI data as part of an Integrated Risk Management (IRM) implementation and the Entity Scoping process. The Administrator wants to leverage the CSDM Data Foundations Dashboard playbooks under the Run tab.

Which CSDM relationships are leveraged using the CSDM playbooks?

A. Business Applications that have their relationships to Information Objects
B. Locations that have established parent records
C. Business Applications that have relationships to Application Services
D. Logical CIs that have relationships with Information Objects

Answer: AD

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

The Run tab in the CSDM Data Foundations Dashboard focuses on enabling operationalized, risk- aware use cases, including Integrated Risk Management (IRM) and Entity Scoping. These use cases require organizations to understand what data is processed, where it resides, and which technical components are involved, rather than only service impact for ITSM.

Information Objects play a central role at this stage.

Option A is correct because Business Applications related to Information Objects allow organizations to identify what types of data (PII, PCI, PHI, regulated data) are processed by each business application. This relationship is essential for risk classification, regulatory compliance, and audit scoping in IRM. Without it, risk assessments lack data sensitivity context.

Option D is also correct because Logical CIs (such as databases, schemas, or data stores) related to Information Objects establish where sensitive data is stored or processed at a technical level. This enables IRM to trace risk from business context down to technical exposure, supporting control testing, issue management, and remediation prioritization.

Option B (Location hierarchy) supports foundational data quality but does not directly enable risk or entity scoping.

Option C (Business Applications to Application Services) is critical for service impact and Change/Incident Management, but it is more aligned to service operations rather than risk and data-centric scoping, which is the focus of the Run playbooks for IRM.

Therefore, the correct answers are A and D, as they directly support IRM entity scoping, regulatory analysis, and risk visibility through CSDM-aligned data modeling.

Question: 63

The Apache Web Server Identification Rule is configured with the following criterion attributes:

Class

Configuration file

Version

kfentifer Entry
yys.cliM.n*™! tonfii-fite version

Yesterday, an Apache Web Server CI was discovered as part of Service Mapping.

Today, the application owner upgraded Apache to a different version and reran discovery of the service.

What will happen in the CMDB?

A. The existing Apache Web Server CI will be reconciled and its version will be updated.

B. A new Apache Web Server CI is created.

C. The Apache Web Server CI will be reclassified as a Web Server CI.

D. A duplication error will occur.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

This scenario hinges on how the Identification and Reconciliation Engine (IRE) in ServiceNow evaluates identification rules.

The identification rule for the Apache Web Server class includes Version as part of the identifier criteria, along with Class and Configuration file. Identification rules must match all criterion attributes exactly for an incoming record to be identified as an existing CI and reconciled.

Yesterday's discovery created an Apache Web Server CI with a specific version value. When Apache is upgraded and discovery is rerun, the Version attribute changes. Because Version is part of the identifier, the incoming record no longer matches the existing CI's identifier entry. As a result, the IRE cannot identify the existing CI as the same configuration item.

When identification fails and no matching identifier entry is found, the IRE proceeds to insert a new CI record rather than updating the existing one. This behavior is intentional and protects the CMDB from incorrectly reconciling records that no longer meet identification criteria.

Option A is incorrect because reconciliation only occurs after successful identification.

Option C is incorrect because reclassification is unrelated to identification criteria.

Option D is incorrect because this is not an error condition; it is expected IRE behavior.

This example highlights a best practice caution: volatile attributes (such as Version) should generally not be used as identifier attributes, as they can cause unintended CI duplication during upgrades.

Question: 64

DRAG DROP

Some steps need to be taken to transition from using legacy CMDB status attributes to CSDM lifecycle objects.
Drag and drop the objects / attributes to the correct descriptions.

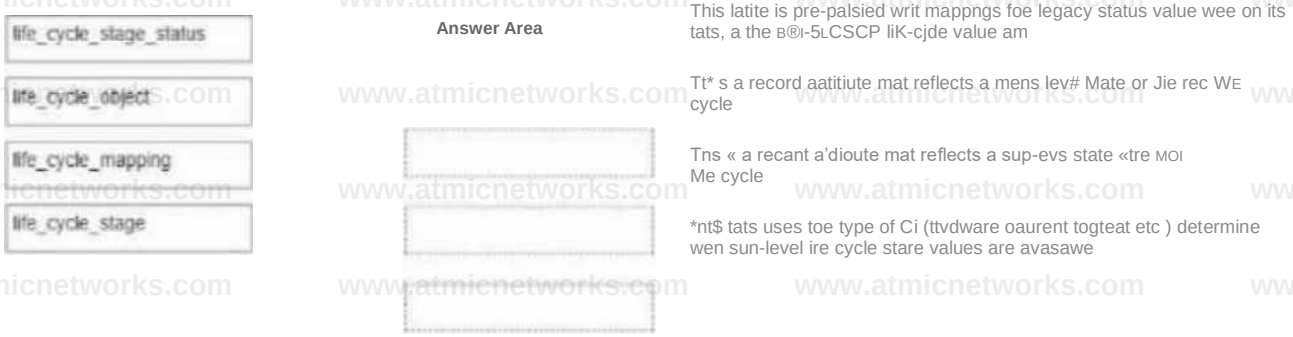

Answer:

Explanation:

Object / Attribute	Correct Description
life_cycle_mapping	This table is pre-populated with mappings for legacy status values; based on this table, legacy values are mapped to the best-fit CSDM lifecycle value pair.
life_cycie_object	This table uses the type of Cl (hardware, document, logical, etc.) to determine which sub-level lifecycle state values are available.
life_cycle_stage	This is a record attribute that reflects a meta-level state of the record lifecycle.
I if e_cyc 1 e_sta ge_s tatu s	This is a record attribute that reflects a sub-level state of the record lifecycle.

When transitioning to CSDM-aligned lifecycle management, ServiceNow introduces a structured lifecycle framework that separates high-level intent from detailed operational state.

Life Cycle Stage represents the macro phase of a CI's existence (for example: Plan, Build, Deploy, Operate, Retire).

Life Cycle Stage Status refines that phase with granular, sub-level states that vary by CI type.

Life Cycle Object determines which lifecycle stages and statuses apply, based on what kind of CI is being managed (hardware, application, document, etc.).

Life Cycle Mapping ensures a safe transition from legacy status fields by translating existing values into standardized CSDM lifecycle values—avoiding data loss or process disruption.

Together, these objects enable consistent lifecycle governance, improve reporting accuracy, and support automation across ITSM, ITOM, and IRM use cases—while remaining backward compatible with legacy CMDB data.

Question: 65

DRAG DROP

A manufacturing organization has implemented Incident Management in ServiceNow and wants to integrate additional products to enhance its functionality.

Drag each ServiceNow product to the value it brings in supporting Incident Management.

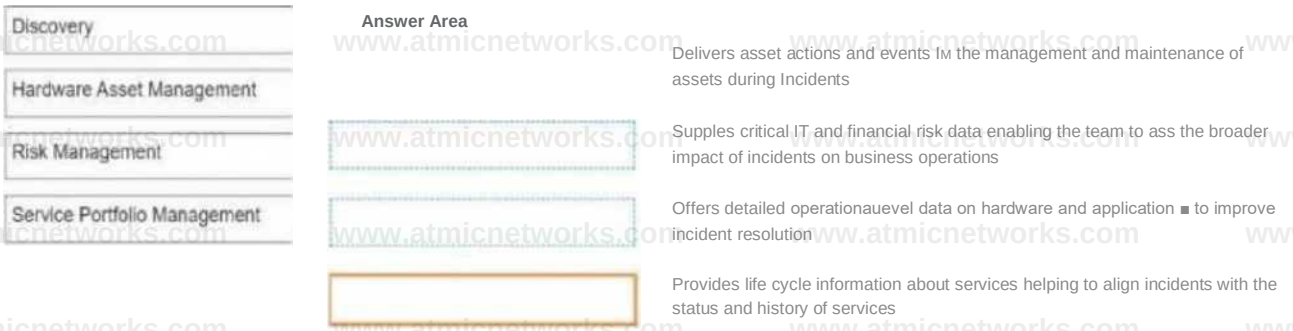

Answer:

Explanation:

Service Now Product	Value to Incident Management
Discovery	Offers detailed operational-level data on hardware and application Cis to improve incident resolution
Hardware Asset Management	Delivers asset actions and events for the management and maintenance of assets during incidents
Risk Management	Supplies critical IT and financial risk data, enabling the team to assess the broader impact of incidents on business operations
Service Portfolio Management	Provides lifecycle information about services, helping to align incidents with the status and history of services

Enhancing Incident Management in ServiceNow often involves integrating complementary products that enrich context, prioritization, and decision-making.

Discovery strengthens incident resolution by automatically populating accurate CI data and relationships, allowing responders to quickly understand affected infrastructure and applications.

Hardware Asset Management (HAM) adds visibility into asset lifecycle, ownership, and maintenance actions, which is especially valuable when incidents involve physical devices or failures.

Risk Management (part of IRM) provides insight into business and financial risk exposure, helping teams prioritize incidents based on potential regulatory, safety, or financial impact.

Service Portfolio Management (SPM) connects incidents to the service lifecycle, enabling better understanding of whether an incident affects a live, retiring, or planned service—and improving communication with stakeholders.

Together, these integrations transform Incident Management from a reactive process into a contextrich, business-aligned capability.

Question: 66

DRAG DROP

Drag and drop the CMDB Health Dashboard metric to the correct description.

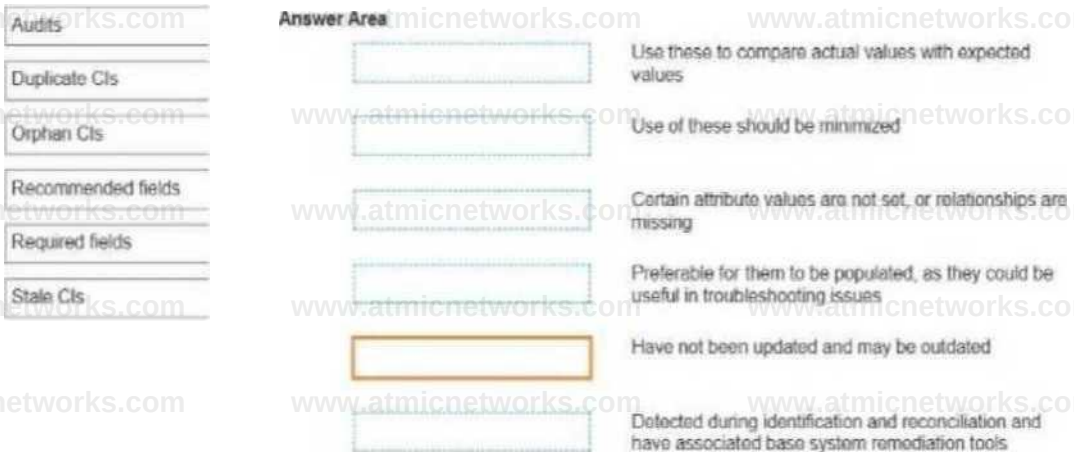

Answer:

Explanation:

CMDB Health Metric	Description
Audits	Use these to compare actual values with expected values
Duplicate Cis	Use of these should be m inimized
Orphan Cis	Certain attribute values are not set, or relationships are missing
Recommended fields	Preferable for them to be populated, as they could be useful in troubleshooting issues
Required fields	Detected during identification and reconciliation and have associated base system remediation toots
Stale Cfs	Have not been updated and may be outdated

The CMDB Health Dashboard organizes data quality into practical, action-oriented metrics:

Audits validate correctness by comparing CI values against defined expectations (rules, policies, certifications).

Duplicate CIs represent redundant records and should be reduced to improve trust and reporting accuracy.

Orphan CIs lack required relationships or key context, limiting impact analysis and service visibility.

Recommended fields are not mandatory but add diagnostic value during incidents and problem investigations.

Required fields are enforced by the platform and closely tied to Identification & Reconciliation remediation.

Stale CIs signal aging data that no longer reflects the current environment.

This mapping aligns with CMDB Data Foundations best practices and how the Health Dashboard is designed to guide prioritization and remediation.

Question: 67

DRAG DROP

A CMDB Owner starts the CSDM journey and needs to become familiar with the CSDM domains.

Drag the CMDB objects to the correct CSDM domains.

Answer Area

Desgn and Planning domain

Foundation domain

Serves Delivery domain

Seii. Consume domain

Answer:

Explanation:

CMDB Object	CSDM Domain
Business Process	Design and Planning domain
Business Application	Foundation domain
Application Service	Service Delivery domain
Business Service	Sell / Consume domain

The Common Service Data Model (CSDM) organizes CMDB data into domains that reflect how services are planned, delivered, and consumed. Correct placement of objects is essential for reporting, ownership, and downstream ITSM, ITOM, and IRM use cases.

Design and Planning domain → Business Process

Business Processes describe how the business operates and what capabilities are required. These are planning constructs used to design services and applications before they are built or delivered.

Foundation domain → Business Application

Business Applications are foundational reference objects that represent what applications exist and who owns them. They anchor ownership, lifecycle, and governance but are not directly operational.

Service Delivery domain → Application Service

Application Services represent running, operational technical services (often created via Service Mapping). They are central to Incident, Change, and Event Management, making them part of Service Delivery.

Sell / Consume domain → Business Service

Business Services represent what is offered to customers or consumers. They define value delivery, SLAs, and consumption context, which is why they belong in the Sell/Consume domain.

This mapping aligns with CSDM v5 guidance and ensures clear separation between planning constructs, foundational records, operational services, and consumer-facing services—a key principle for a scalable and governed CMDB.

Question: 68

(Choose 2 options)

The following Reconciliation Rules were configured for ServiceNow, Altiris, and SCCM for the Windows Server (cmdb_ci_win_server) class:

Discovery Source	Class	Priority
ServiceNow	Windows Server [cmdb_ci_win_server]	100
Altiris	Windows Server [cmdb_ci_win_server]	200
SCCM	Windows Server [cmdb_ci_win_server]	300

Which statements are true?

A. Data collected with a discovery source of Altiris can update records inserted by SCCM into the Windows Server table.

B. Data collected with a discovery source of ServiceNow can insert new records into the Windows Server table, but cannot update records created by Altiris or SCCM.

C. Data collected with a discovery source of SCCM can be inserted as new records in the Windows Server table.

D. Data collected with a discovery source of SCCM can update any record in the Windows Server table because it has the highest priority number.

Answer: A C

Explanation:

This question tests understanding of reconciliation source priority in the Identification and Reconciliation Engine (IRE) in ServiceNow.

In reconciliation rules, lower numeric values represent higher priority. Therefore, the priority order is:

ServiceNow (100) – highest authority

Altiris (200)

SCCM (300) – lowest authority

Why A is correct

Because Altiris (200) has higher priority than SCCM (300), data from Altiris can update records originally inserted by SCCM. This is exactly how reconciliation precedence works—higher-priority sources can overwrite lower-priority ones.

Why C is correct

SCCM, even though it has the lowest priority, is still an authorized discovery source. It can insert new records into the Windows Server table when no existing CI is identified. Priority only affects updates, not the ability to create records.

Why B is incorrect

ServiceNow (priority 100) can update records from Altiris and SCCM because it has the highest priority. The statement incorrectly claims it cannot.

Why D is incorrect

SCCM does not have the highest authority. A higher numeric value means lower priority, so SCCM cannot update records created by higher-priority sources.

Question: 69

The CMDB Administrator has set up two Dynamic Reconciliation Rules within the ServiceNow Production Instance. The 'Server' class has a Dynamic

Reconciliation Rule of largest value for the RAM field. The 'Windows Server' class has a Dynamic

Reconciliation Rule of most reported for the RAM field.

Discovery Source	RAM Mt
Twoil	4.0%
SerMceNow	40%
LANDesk	2 048
Altwis	8.192

Given the above data in Multisource CMDB, which value is added to the CMDB for RAM for a Server CI?

A. 2,048 MB
B. 8,192 MB
C. 4,096 MB

Answer: B

Explanation:

Comprehensive and Detailed Explanation (200–300 words):

This question hinges on understanding class-specific Dynamic Reconciliation Rules in CMDB 360 / Multisource CMDB within ServiceNow.

Although two different dynamic rules are configured, the rule applied depends on the CI class being evaluated:

For the Server class, the configured rule is Largest value for the RAM attribute.

For the Windows Server class, the rule Most reported would apply—but only if the CI were evaluated as a Windows Server.

The question explicitly asks for the resulting RAM value for a Server CI, not a Windows Server CI.

Therefore, the Largest value rule governs the outcome.

Looking at the multisource values:

2,048 MB (LANDesk)

4,096 MB (Tivoli)

4,096 MB (ServiceNow)

8,192 MB (Altiris)

Under the Largest value dynamic reconciliation rule, the IRE selects the maximum numeric value, regardless of how frequently it is reported or which source provided it.

The Most reported logic (which would result in 4,096 MB) does not apply here because that rule is configured for a different class (Windows Server).

This scenario illustrates an important CMDB 360 principle: Dynamic reconciliation is evaluated per class, and child-class rules do not override parent-class rules unless the CI is actually classified under that child class.