CIS Data Foundations study guide

Start here

Domain	Weight	Covers
Govern the CMDB	**35%**	Health Dashboard scorecards, duplicates, reclassification, remediation, life cycle, Data Manager
Gain Insight	20%	Query Builder, Unified Map, Intelligent Search, Data Foundations Dashboards, value to other products
Ingest Data	19%	Discovery, Service Mapping, ACC, Service Graph Connectors, IntegrationHub ETL, Import Sets, group data
Configure the CMDB	15%	CI Class Manager, IRE, identification and reconciliation rules, CMDB 360
CSDM Fundamentals	11%	Domains, service types, Foundation/Crawl/Walk/Run/Fly, key tables

CMDB basics

The CMDB is a single system of record for CIs, services, and configured assets, and the hub that Service Management, Asset Management, and Operations Management depend on. A CI can be logical (an application), physical (a server), or an operational construct (a cluster).

The three pillars and the four deployment steps

> **Ingest** gets data in. **Govern** keeps it trustworthy. **Insight** gets value out. It's a continuous cycle, not a one-time setup.

> The steps to a successful CMDB, in order: **1** set your direction, **2** build a team and governance model, **3** design your service data model, **4** operationalize and maintain. You don't need a comprehensive CMDB right away.

Two critical success factors: stakeholder **alignment** on goals, and a **strategic connection** to both organizational strategy and IT's tactical goals. Plan and Govern are the practices that minimize technical debt, and clear governance plus continuous health monitoring is what makes a CMDB operational.

Teams, governance, and data model design

- Form a **Configuration Control Board (CCB)**: the team that sets direction, decides what the CMDB needs to support strategic initiatives, and decides when it's ready to use. Get executive buy-in early.
- The **CMDB team doesn't own the data** in the CMDB. Data owners must be part of governance; the CMDB team can usually only run basic checks like duplicates and missing attributes.
- Put in only the data your use cases need. Start with out-of-the-box classes (for example, Firewall, Server, and Application for a security use case), and add classes as new use cases come up.
- For each class you ingest: review its base attributes, identify required and missing attributes, and identify the authoritative source for each attribute. Custom classes and attributes create technical debt, and a CMDB health check looks for them.
- **Discoverable** attributes (host name, IP address, operating system) come from discovery tools. **Non-discoverable** ones (groups, owners, location) need manual entry or another source; use Data Manager certification and Health Dashboard recommended fields to find gaps.

Key terms

- **Configuration Item [`cmdb_ci`]** is the core CI table that every CI class extends, and it references foundation tables like `sys_user` and `cmn_location`. (It extends the Base Configuration Item [`cmdb`] table.)
- A **class** is a table of CIs that share attributes. A table can be both a parent and a child (Server extends Computer; Windows Server extends Server). View the structure in CI Class Manager or the schema map.
- **Logical CIs** take up no physical space and depend on physical components, for example service offerings, services, and business applications. They're entered manually.
- An asset becomes a CI when it's deployed and stays one until it's returned to stock or disposed of. Configuration Management focuses on operational use; IT Asset Management focuses on financial and contractual life cycle.

Asset vs. CI

Asset only (financial)	Both	CI only (operational)
Contracts, consumables, software licenses, buildings	PCs, printers, routers, switches, physical servers, on-prem data centers	Applications, databases, virtual servers, websites, Dynamic CI Groups, cloud environments

An asset is often a CI, but a CI isn't necessarily an asset. A business service is a CI with no financial value of its own.

The three "application" records

	Business Application	Application Service	Application
What it is	Inventory/portfolio record ("we have a Payroll System")	A running deployment of that app (per environment, region, or line of business)	Installed software or running process on a host
Used in Incident/Problem/Change?	**Never**	**Yes**, the usual pick when a caller reports an app issue	Can be the root cause
How it's created	Manually	Manually or via Service Mapping	Discovered. Don't populate manually.
Version specific?	No	Yes (unique instance)	Yes

Foundational data

Foundational data lives in common tables **outside** the CMDB and is referenced **from** CI fields: `cmn_location`, `cmn_building`, `cmn_cost_center`, `core_company`, `cmn_department`, `sys_user`, `sys_user_group`. It usually arrives from LDAP, Active Directory, or HR sources.

Management tools (all free)

CMDB Workspace (from the Store, spans all three pillars), CI Class Manager, CMDB Health Dashboard, CMDB Data Manager, Duplicate CI Remediator, De-duplication Templates, Unified Map, Query Builder, and Intelligent Search are all no-extra-license tools. The CMDB and CSDM Data Foundations Dashboards are a free Store app.

1. Govern the CMDB

35% of the exam

CMDB Health Dashboard: three scorecards

Scorecard	Metrics	Question it answers
Completeness	Required, Recommended	Are the important fields populated?
Compliance	Audit (Desired State and Scripted)	Do CIs match the expected configuration?
Correctness	Duplicate, Orphan, Stale	Is the data unique, related, and current?

- Health rolls up from class level to the overall CMDB, and can also be viewed per **CMDB group** (use groups of type Health for regional views).
- **Health Inclusion Rules** are the first filter. They define which classes (and which CIs) count, so a discovery tool that scans everything doesn't drown your score. Use the **"is a"** operator to include a class hierarchy. A rule can apply to one metric and not another, which is why two metrics can evaluate different CI counts for the same class.
- Scores only appear once the dashboard's **scheduled jobs** run. Run them as an admin. Job status is in `cmdb_health_metric_status`.
- Those scheduled jobs **aren't active by default**; activate them before expecting scores. The dashboard itself is free and opens from **CMDB Workspace > Management**.
- The dashboard shows a **Health Scorecard** (overall Completeness, Correctness, Compliance) and a **Top 10** view of the classes hurting each score, with drill-down to CIs.

- Create a regional or team view: **Configuration > CMDB Groups > New**, set **Group type = Health**, add encoded queries in the **CMDB Group Contains Encoded Queries** related list, check **Show All CI**, then run the Completeness, Compliance, and Correctness jobs.

Completeness: Required vs. Recommended

	Required	Recommended
Enforced?	Yes, equivalent to mandatory in the dictionary	No, visibility only
Risk	Can make **Discovery and imports fail**	None
When	Sparingly, once data is reliable	**Preferred during initial implementation**
Configured in	Dictionary or **dictionary override** (CI Class Manager > Attributes)	CI Class Manager > Health > Recommended Fields

To require a field only on part of the hierarchy, use a **dictionary override**. For example, Serial Number lives on `cmdb_ci` and isn't required there, so override it on Hardware to make it mandatory for Hardware and its children. If a mandatory field makes Discovery or other IRE inserts fail, set `glide.required.attribute.enabled` to false (Configuration > CMDB Properties > Identification/Reconciliation Properties).

The base system ships **no** predefined required or recommended CI fields. Non-discoverable fields like Support Group, Change Group, and Managed By Group belong in Recommended, not Required.

Compliance: audits

1. **Certification Filter**: the scope (for example, all UNIX servers in one data center).
2. **Certification Template**: the conditions to pass. It references the filter.
3. **Audit**: runs the comparison on a schedule and can create tasks. Its type must be **Desired State**.

- Set up all three from **CI Class Manager > class > Health > Compliance**, or from **Compliance > Filters**, **Templates**, and **Audits**.
- Only **Desired State** and **Scripted** audits count on the Compliance scorecard. Audit types such as Compliance and Compliance Architecture are for other products.
- **Scripted audits** define the pass criteria in a script when template conditions are too limited. They need the `certification_admin` role, can't be created in CI Class Manager

(use **Compliance > Scripted Audits > Audits**), and the base system includes a sample, "Data Center Zone Usage."

- After an audit runs, activate and run the **CMDB Health Dashboard - Compliance Score Calculation** job (Configuration > Health Preferences > Scheduled Jobs). Run health jobs as an admin, or results can be inaccurate.
- Don't confuse the two: a **Desired State audit** checks automatically that a field is filled in (ServerXYZ has an owner). A **Data Manager certification policy** has a person confirm the value is right (Joe Smith really owns ServerXYZ).

Correctness: Duplicate, Orphan, Stale

Metric	Definition	Key facts
Duplicate	More than one record for one real item	Only **independent** CIs are evaluated. The oldest (by creation date) stays; newer ones get **Duplicate of** set.
Orphan	Lost a required relationship (for example, an app with no Runs on::Runs to a server)	**No orphan rules ship in the base system.** One orphan rule per class.
Stale	Not updated within a set period	One base rule on `cmdb_ci`: **60 days**, based on `sys_updated_on`. Applies to all child classes.

Duplicates: how the IRE handles them and how you fix them

- `glide.identification_engine.skip_duplicates` (true) and `...skip_duplicates.threshold` (5): when the IRE matches several duplicates under the threshold, it updates the **oldest** and tags the rest with **Duplicate of**. The IRE skips any CI whose Duplicate of is populated.
- **Duplicate of** replaced **Discovery source** as the duplicate marker, so Discovery source now keeps the last real source.
- De-duplication tasks are created **automatically**, with no configuration. Work them from the CMDB Workspace (Home, Important actions) or the De-duplication Tasks module.

Tool	Use
Duplicate CI Remediator	Wizard that resolves one task at a time. You choose the Main CI to keep.

De-duplication Templates	Resolve similar tasks in **bulk**, per class, using matching rules such as class, location, manufacturer, and model. They can run on a schedule and merge or delete automatically, so test them well before publishing.
Correctness scorecard	Shows how many duplicates exist
Now Assist for CMDB	Generative AI detection and suggested resolution

> **Main CI** (formerly Master CI) recommendation: the one that is **oldest**, has the **most relationships**, and was **most recently updated**. In practice, keep the record with related tasks and the one ServiceNow Discovery updates, and merge relationships from the others.

Reclassification

During identification, the IRE can move a CI to a better class. It's **automatic by default**. If the relevant property is false, the CI isn't touched and a **reclassification task** is created for review.

Type	Direction and example	Property	Data loss?
Upgrade	To a child: Server → Windows Server	`glide.class.upgrade.enabled`	No
Downgrade	To a parent: Windows Server → Server	`glide.class.downgrade.enabled`	**Yes**
Switch	Across branches: Linux Server → Windows Server	`glide.class.switch.enabled`	**Yes**

Granular control

- `glide.identification_engine.update_without_{upgrade|downgrade|switch}_enabled` are **false** by default. Set to **true**, the IRE applies the attribute updates but **skips the reclassification**. They take precedence over the legacy properties, but they're global (all or nothing).

- **Reclassification Restriction Rules** (`cmdb_ire_reclassification_restriction`) block a downgrade or switch between **specific** source and target classes while still allowing attribute updates. **Source/Target inheritance = true** extends the rule to child classes.

Worked examples from the Govern course

downgrade.enabled = true and update_without_downgrade_enabled = true. A Windows Server is identified as a Server.

It is NOT downgraded, and its attributes CAN be updated.

Downgrade enabled, update_without_downgrade false, restriction rule: Downgrade, Server (inheritance true) → Computer. A Windows Server is identified as a Computer.

It is NOT downgraded (Windows Server extends Server, and inheritance is on). Its attributes CAN be updated.

Same settings. An IP Router is identified as Network Gear.

It IS downgraded, because the rule only covers Server and its children. Attributes CAN be updated.

Remediation: task, rule, workflow

> By default, **no tasks** are created for failing health metrics, **except de-duplication tasks**. Turn them on in **Configuration > Health Preferences > Health Metrics**: pick the metric, toggle Create Task, set the assignee group (use **Managed By Group**).
>
> Flow: **Remediation Task → Remediation Rule → Workflow**. The rule references a classic workflow (IntegrationHub flows need custom work) and can run automatically or when someone clicks Remediate.

- Staleness, Orphan, Required, and Recommended all follow that same setup. First define staleness rules, orphan rules, and required/recommended fields in CI Class Manager.
- Two exceptions don't follow the flow: **compliance audit tasks** (configured in CI Class Manager > Health > Compliance) and **de-duplication tasks** (automatic).
- Example from the course: a stale CI triggers a workflow that re-runs Discovery. If it succeeds, the CI is no longer stale. If not, the workflow sets Non-Operational/Retired and assigns to Hardware.
- Now on Now: automating CMDB audits saved about 1,252 hours a year, a 60% improvement.

Life cycle management and field sync

Effective life cycle management: understand each data type's life cycle (containers live days, servers live years), identify data owners with the **Managed By Group**, set up life cycle fields, then automate with Data Manager policies.

Field group	Fields
Shared CI ↔ Asset	Asset tag, Department, Location, Support Group, Company, Serial Number, Model ID, Assigned to
CI legacy status	Operational status, Install Status, Hardware Status, Substatus
Asset legacy status	State, Substate
CSDM life cycle (both)	Life Cycle Stage, Life Cycle Stage Status

- Don't manage every status field. With legacy fields, pick Install Status **or** Hardware Status unless you have a business reason for both.
- Shared-field mappings live in **Asset-CI Field Mappings** and can be extended (for example, asset Comments → CI Short description). Baseline sync runs through business rules.
- Legacy sync is one-way in places: asset **State** → CI **Install Status**, but changing Install Status does **not** update asset State.

Object	What it is
`life_cycle_stage`	Meta-level state of the record
`life_cycle_stage_status`	Sub-level state
`life_cycle_mapping`	Pre-populated rules mapping each legacy value (by table) to a CSDM stage/status pair. This is the one you edit.
`life_cycle_object`	Uses the CI type (hardware, document, logical) to decide which status values apply

Click **Enable Life Cycle Sync** for a one-time sync from legacy to CSDM values, then resolve unmapped values in the **Discrepancy Report**. After that, legacy and CSDM fields stay in sync across the CI, the asset, and the related **Install Base Item**. End-of-Life statuses include Disposed and Donated.

CMDB Data Manager

A policy-driven framework for bulk end-of-life and certification work. The new Data Manager lives in the **CMDB Workspace** (legacy Configuration > CMDB Data Manager is deprecated by Yokohama). Old policies are compatible. Workspace 6.0.1 added certification policies. You create and publish policies from the Management tab; assignees work tasks from **My Work**.

Policy	Does	Example
Retire	Moves CIs to retired/End of Life	Retire ownerless computers created over a year ago
Archive	Removes records but keeps them restorable	Archive Seattle Linux servers not updated in six months
Delete	Permanently removes	Delete containers not discovered in the past week
Attestation	Confirms the CI still exists	Attest that servers in a location physically exist
Certification	A person validates specific attributes	Certify application services have the right owner

- Roles: `data_manager_admin` and `data_manager_user`.
- CSDM order: a CI must be **Retired** (End of Life) before it can be Archived or Deleted.
- Retire, Archive, and Delete policies need an **active life cycle rule** for every class the policy targets.
- Before using it, review the `life_cycle_mapping` entries: complete custom legacy mappings, give every mapping a lifecycle control, and activate them.
- Assign tasks with the **Managed By Group**. Discovery doesn't populate it, so maintain it through Technology Management Offerings with Dynamic CI Groups, or through CI Class Manager.

2. Gain Insight

20% of the exam

Tool	Cost	Where	Does
CMDB Query Builder	Free	CMDB Workspace > Management	Visual multi-table queries, no SQL
Unified Map	Free	CMDB Workspace (Home > Quick Links)	Map of CIs and relationships around a home node
Intelligent Search	Free	CMDB Workspace > Home	Natural language queries
Data Foundations Dashboards	Free Store app	Store install	35+ fixed data-quality metrics with playbooks

CMDB Query Builder

- Query types: **CMDB Query**, **Service Mapping Query**, and **Combination Query** (a saved Service Mapping query embedded in a CMDB query).
- Can mix CMDB and non-CMDB tables ("high-priority incidents for servers in Paris").
- Results can populate a CMDB Group, be exported, or be scheduled by email (daily, weekly, monthly, periodically, once, on demand).
- Move queries between instances with **Export query** (XML), then **Import query** from the Saved Queries tab.
- Interface: the **CMDB Classes**, **Non-CMDB Tables**, and **Saved Service Queries** lists to drag from; the **Filters** pane; the **Canvas** where you connect classes with relationships and AND/OR; the **Properties** tab (extra result columns) and **Applied Filters** tab; and **Query Results**.
- Its system properties are under **Configuration > CMDB Properties > Query Builder Properties**.

Unified Map

Combines **Dependency View and Application Service maps** into one view centered on a home node. Badges show incidents and changes, the Related items panel shows active incidents and planned changes, and you can filter (discovery source, CI type, levels) and

save filters as presets. It's used for impact analysis in Incident, Problem, Change, and Event Management.

Intelligent Search (NLQ)

Plain-language, multi-table queries without knowing the data model. `glide.cmdb.query.nlq.activated` (on by default) integrates it with Query Builder, so multi-table results show **View in Query Builder**. It keeps your five most recent searches. Sample searches live in `sn_cmdb_ws_nlq_sample_search`, editable by `sn_cmdb_admin`.

Data Foundations Dashboards vs. Health Dashboard

	Data Foundations Dashboards	CMDB Health Dashboard
Source	Store app	Base instance
Metrics	35+, fixed, not configurable	Three configurable scorecards
Scope	Can't customize	Filter by class, group, service, inclusion rules
Nature	Prescriptive best-practice checks with playbooks	Completeness, Compliance, Correctness
Audience	Internal CMDB managers	Senior IT leadership

- Two dashboards: **CMDB** (configuration quality and risk) and **CSDM** (framework adoption; it has Foundation, Crawl, Walk, Run tabs).
- They show nothing until the base scheduled jobs (names contain **"Get Well"**) run. Calculations live in **GetWell** Script Includes. The metrics are a subset of **HealthScan**, which you request through your solution consultant.
- Install the **CMDB and CSDM Data Foundations Dashboards** app from System Applications > All Available Applications. Open them from **Configuration > CMDB Data Foundations Dashboard** or **CSDM Data Foundations Dashboard**, or from the CMDB Workspace Management tab.
- To fill them right away, open each "Get Well" job in **System Definition > Scheduled Jobs** and click **Execute Now**.
- Focus on **Critical and High** indicators first.

Value to other products

Product	How a healthy CMDB helps
Incident	CI's **Support group → Assignment group**. Impacted services, priority by CI and service criticality.
Change	CI's **Change group → Assignment group**. Impact analysis through relationships.
Event Management	Binds alerts to CIs and services. Automates incident creation.
Asset Management	CI ↔ asset sync. Turn on the **Data Normalization plugin** before SAM (HP = Hewlett Packard).
Security Operations	Finds which servers run a vulnerable version and prioritizes by criticality

A well-configured CMDB can cut IT costs by up to 40%.

3. Ingest Data

19% of the exam

Ingestion tools

Tool	Also called	License	Know this
ServiceNow Discovery	Horizontal, agentless, IP-based	ITOM Visibility	Scheduled, **not real-time**. Uses a MID Server. Creates relationships.
Service Mapping	Top-down, tag-based, ML-powered	ITOM Visibility	Builds application service maps. Define service owners per CSDM.
Agent Client Collector	ACC, agent-based	Paid	**Real-time**. The only way to get **software utilization** (SAM). Best for intermittently

			connected devices and secure networks.
Service Graph Connectors	SGC	Paid, **SCCM is free**	Pre-built Store integrations. Built on IH-ETL and the IRE. Fastest time to value. Needs target credentials.
IntegrationHub ETL	IH-ETL	**Free**	RTE + IRE. Test with sample data, preview, roll back.
Import Sets + Transform Maps		Free	Oldest method. **Doesn't use the IRE by default.**

Aim for about 90% of CMDB data populated automatically and 10% maintained manually under governance.

ServiceNow Discovery

- Phases in order: **Scan → Classification → Identification → Exploration**.
- Types of horizontal discovery: **Network** (finds internal IP networks), **CI** (devices, computers, and applications; the one run most often), **Cloud** (AWS and Azure resources), and **Serverless** (applications on hosts without a host scan first, using patterns).
- Manage schedules from the **Discovery Admin Workspace**. Open the Configuration Console from **Discovery Definition > Configuration Console**; excluding a CI type disables its probe or classifier.
- The MID Server sits inside your network. The instance sends **probes**; results come back to **sensors**. Protocols: **WMI** (Windows), **SSH** (Unix/Linux), **SNMP** (network gear, printers).
- Start small. Keep schedules to **/16 or smaller** (about 65,000 IPs), and use ranges that exclude the network and broadcast addresses. Use the **Discovery Configuration Console** to switch off CI types you don't manage.
- Discovered CIs get Discovery source = **ServiceNow**. A record first created by Asset Management shows **SNAssetManagement** until Discovery finds it.
- Non-discoverable data to add afterward: Assigned to, Owned by, Support group, Change group, Managed by group.

Application service types

Type	Best for

Service Mapping (Top-down)	Mission-critical services that need precise pattern-based maps
Service Mapping (Connection Suggestion)	Custom-built apps, using fingerprinting
Tag-Based	Cloud-native, container, and VM environments
Dynamic CI Group	Small services easily grouped with filters and CMDB queries

Population method tables: discovered/manual = `cmdb_ci_service_discovered`, tags = `cmdb_ci_service_by_tags`, calculated = `cmdb_ci_service_calculated`, query-based = `cmdb_ci_query_based_service`. All extend `cmdb_ci_service_auto`.

Imports and the IRE

- **IH-ETL flow:** source data lands in import set staging tables → the **Robust Transform Engine** uses the ETL transform map to build IRE payloads → the **IRE** identifies and reconciles → data is written.
- **IH-ETL terms:** a CMDB Application (Name + Discovery Source, for example SCCM) can have many ETL transform maps; each map has one Data Source.
- **SGC customization:** a customized script transform creates a skipped change on upgrade; a changed query means its updates are skipped; customized mappings aren't supported. Service Graph Connector Central is a CMDB Workspace tab.
- **SGC facts:** categories include partner cloud providers, endpoint and IT asset management, monitoring and observability, security, software/server/network apps, operational technology (OT), and partner-built connectors. Each connector is reviewed for IRE use and CSDM conformance, and all of them use IH-ETL. Using an SGC that goes through the IRE needs **ITOM Visibility** licensing, except SCCM. Customizing a connector or its mappings makes the customer fully responsible for supporting it.
- **IH-ETL guided setup** (ETL Transform Map Assistant), in order: **Basic Details** (data source and import set) → **Preview and Prepare Source Data** (cleanse and transform) → **Select CMDB Classes and Map Source Data** (plus relationships) → **Preview and Schedule** (test run, roll back if needed, then schedule).
- The **Integration Commons for CMDB** Store app (`sn_cmdb_int_util`) adds the **CMDB Integrations Dashboard**, with performance and error metrics for every SGC and custom IH-ETL integration.
- **Import sets into the CMDB:** IH-ETL is the more modern option. If you use a transform map, add an **onBefore** script that calls **CMDBTransformUtil** so rows go through the IRE; when it does, the transform map's coalesce values are ignored. Plan first: which classes, which attributes, how records are identified, which target fields are mandatory, and what custom logic you need.

- **Non-CMDB tables:** the IRE and IH-ETL can load only the tables in `glide.identification_engine.non_cmdb_tables` (`cmn_location`, `cmn_department`, `cmn_cost_center`, `cmn_building`, `sys_user`, `sys_user_group`). Each gets exactly **one** identification rule, managed from list views, not CI Class Manager.

Populating group data

Group field	Drives
Support Group	Incident Assignment group
Change Group	Change Assignment group
Managed By Group	Data Manager and remediation task assignment

Method	Sets	Scope	Precedence
Technology Management Service Offering	Support, Change, Managed By	CIs in the offering's Dynamic CI Group	**Priority 1 (wins)**
CI Class Manager	Managed By only	All CIs in that class (not extended classes)	Priority 2

Build chain: CMDB Group → Dynamic CI Group → Technology Management Service → Technology Management Service Offering → relate the Dynamic CI Group to the offering with **Contains::Contained by** → set the group fields. The business rule **CSDM - Sync Group Attributes** does the sync.

Where each piece lives: **Configuration > CMDB Groups**, **Configuration > Dynamic CI Groups**, **CSDM > Technology Management Service**, and **CSDM > Technical Service Offering** (the offering references the service as its parent). Add the relationship from the offering's Related Items with the green plus. View the sync script in System Definition > Business Rules, Advanced tab.

4. Configure the CMDB

15% of the exam

CI Class Manager

The central tree view for the class hierarchy (Configuration Item → Hardware → Computer → Server → Linux Server). Children inherit parent attributes.

Tab or setting	What it does
Basic Info	Display name, **Extensible** flag, **Principal Class** check box, table name, **Managed By Group**, default product model, icon
Attributes	View or add attributes, set mandatory. "Added" shows attributes created on this class.
Identification Rules	Identifier entries, independent vs. dependent
Reconciliation Rules	Which sources update which attributes, data refresh rules
Suggested Relationships	Recommended relationship types and directions
Health	Recommended/required fields, compliance audits, filters, templates
CI List, Pinned Classes	Records in the class; quick access to classes you use often

- **Principal Class** limits which CIs appear in the Configuration Item field on Incident, Problem, and Change. Adding or removing needs `sn_cmdb_admin` + `personalize_dictionary`.
- **New classes**: avoid when possible. First install the **CMDB CI Class Models** Store app, then extend an existing class, and only create net-new as a last resort. Roles: `admin`, or `sn_cmdb_admin` + `personalize_dictionary`. Custom tables start with `u_cmdb_ci` and need identification and reconciliation rules (created or inherited).
- To retire a custom class, find a baseline class in CI Class Manager and move the CIs over with attributes mapped.

The IRE and identification rules

Every source calls the IRE's APIs; nothing writes straight to the CMDB. Goals: fewer duplicates, fewer unwanted updates, no attribute flapping. Its four functions are

Identification, Reconciliation, De-duplication, and Reclassification. Logical CIs (Business Applications, Application Services, Services, Offerings) are created manually with wizards. Foundation data bypasses the IRE.

> The identification flow tries identifier entries in priority order (lower number first), for example Serial Number, then Name, then MAC Address. The first match **updates** the CI and hands off to reconciliation. Only when every entry fails does the IRE **create** a new CI.

- Default rules: **Hardware** on `cmdb_ci_hardware`, **Application** on `cmdb_ci_appl`. View them in **Configuration > Identification/Reconciliation > CI Identifiers** or CI Class Manager.
- **Independent** CIs can be identified alone. **Dependent** CIs (like applications) need a relationship to another CI.
- **Allow null attribute = false** means every criterion attribute must have a value or the entry is skipped. That's a common cause of unexpected duplicates.
- The **Hardware** rule tries, in order: serial number with serial number type, then serial number, then name, then MAC address with name. A match updates the CI; no match creates one. Since Washington, `product_instance_id` is the first entry, currently used only for healthcare CIs.
- Software usually has no serial number, so applications are identified by the process they run.
- Good identifiers are **unique and unchanging**: serial number, configuration file path. Poor ones: IP address, version, and name for multi-instance apps. If Version is a criterion, an upgrade creates a new CI.
- **Allow fallback to parent's rules**: if a child rule doesn't match, try the parent's.
- **Optional conditions** work on lookup fields by default. For non-lookup fields, set `glide.identification_engine.enable_identifier_optional_condition` to true.

Reconciliation

- **No reconciliation rules ship in the baseline.** Create one per authorized source. They govern **updates**, not inserts. Once a higher-priority source updates a record, lower-priority sources can't overwrite those attributes.
- **Lowest order number = highest priority.**
- To block **inserts** from a source, use **IRE Data Source Rules** (none in baseline).
- Reconciliation runs after identification, only on **updates**; it isn't part of creating a CI. Well-built rules handle about 90% of deciding which source is trusted.
- **Data Refresh Rules** let a lower-priority source update after the higher-priority one has been silent for a set number of days.
- `cmdb_datasource_last_update` (Data Source Histories) records the last update per source, but only once reconciliation rules exist.

CMDB 360 / Multisource CMDB

Without it, values from losing sources are thrown away. With it, every source's values (winners and losers) are kept in `cmdb_multisource_data`, so you can compare sources, change rules and **recompute** without a new load, and **revert** a source's contributions. It only sees data that went through the IRE.

To turn it on: activate the **ITOM Discovery License** plugin (`com.snc.itom.vis.license`) from System Applications > All Available Applications, then set `glide.identification_engine.multisource_enabled` to true. Each `cmdb_multisource_data` record (added in Paris) stores common attributes as columns and every attribute in a JSON blob. Before Paris, only the last source per record was tracked, in Data Source Histories [`cmdb_datasource_last_update`].

Property	Purpose
`glide.identification_engine.multisource_enabled`	Turns it on (false by default). Also adds the Static/Dynamic choice to reconciliation rules.
`glide.identification_engine.multisource_cmdb_ci_enabled`	Captures data for CIs in CMDB classes
`glide.cmdb.logger.source.cmdb_multisource`	Logging
`glide.identification_engine.multisource.recompute.max.ci.limit`	Recompute cap (100,000 CIs). Target specific classes.

View the data in CI Class Manager, the CMDB 360 Data Preview link, the Multisource Data related list, and the **CMDB 360** view in CMDB Workspace (dashboard showing number of sources, CIs not reported by sources, attribute gaps, and conflicting values). The view refreshes daily through the Multisource Dashboard Analytics Population job.

Recompute is run from CI Class Manager when reconciliation rules change or a source turns out to be bad. On a CI, the **CMDB 360 Data Preview** link is under Related Links. The CMDB 360 Saved queries card in CMDB Workspace shows up to **20** queries, newest first.

Static vs. dynamic reconciliation rules

	Static	Dynamic

Picks by	**Who** sent it (source priority)	**What** the value is
Available	Always	Only after Multisource is enabled
Types	Priority order	First Reported, Last Reported, Most Reported, Largest Value, Smallest Value (Largest/Smallest are numeric only)

Dynamic beats static for the same attribute, and a rule on the CI's own class beats a parent-class rule.

5. CSDM Fundamentals

11% of the exam

CSDM is a shared set of service definitions across ServiceNow products with prescriptive modeling guidance. It is **not** a product, SKU, process guide, set of reports, or code, and it doesn't fix past implementations automatically.

Seven domains (CSDM 5)

Domain	Contains
Foundation	Referential data: company, business unit, department, location, groups, users, product models, contracts, value streams. **Business Process** is the one CMDB table here. Personas: Process Owner, Data Steward, Product Owner, Contract Manager.
Ideation & Strategy	New in CSDM 5: ideas, strategy, planning items
Design & Planning	Business Capability, **Business Application**, Information Object
Build & Integration	SDLC Components, DevOps artifacts

Service Delivery	Application Services, Technology Management Services and Offerings, infrastructure CIs, Dynamic CI Groups
Service Consumption	Business Services, Business Service Offerings, request catalog
Manage Portfolio	Spans portions of all the other domains

Service types and tables

Type	Table	Published to
Business Service	`cmdb_ci_service_business`	Business users; underpins business capabilities
Technology Management Service (was Technical Service)	`cmdb_ci_service_technical`	Service owners; underpins business or application services
Application Service / Service Instance	`cmdb_ci_service_auto`	Logical representation of a deployed stack
Service Offering (business or technical)	`service_offering`	Every service should have at least one
Dynamic CI Group	`cmdb_ci_query_based_service`	Grouping of CIs by criteria

Yokohama relabels: Technical Service → **Technology Management Service**, Technical Service Offering → **Technology Management Service Offering**, and `cmdb_ci_service_auto` → **Service Instance**.

Foundation → Crawl → Walk → Run → Fly

Stage	Objects	Value

Foundation	Company, business unit, department, location, groups, users, CMDB group, product models, contracts	Reporting context for everything else
Crawl	Business Application, Application Service, Application, Server/Host (SDLC Component optional)	Minimum for Incident, Problem, Change
Walk	Technology Management Service and Offering, Dynamic CI Group	Auto-sync support and approval groups onto CIs
Run	Business Service Portfolio, Business Service, Business Service Offering	Business impact in Incident and Change; subscribers by user, location, department
Fly	**Business Capability**, **Information Object**, request catalog	Rationalization with EA; tracking PII, PCI, HIPAA data

Other CSDM facts from the practice bank

- Incident forms use **Service Offering** and **Application Service**. Prioritize by **Business Criticality** on the offering.
- Model a platform (like ServiceNow itself) as a **Platform Host** with Platform Applications that reference it.
- Track regulatory compliance by linking **compliance identifiers** to CSDM objects.
- Governance teams: Foundation = Enterprise Architects, Data Stewards, Process Owners, Platform Owner. Design = Enterprise Architects, Platform Owner. Technical = Technology and Application Service Owners, Platform Owner. Portfolio = Service Owners, Platform Owner.

What's new in CSDM 5

- A new **Ideation & Strategy** domain; **Software Bill of Materials (SBOM)**, from the Store as SBOM Core (`sn_sbom_core`); the **System Component Model** (with Software Component and Service Offering models) and Product Feature; **Value Streams**; **Teams**; new life cycle stages and statuses; the DevOps Change data model (Store); **Service Instance** and **AI** CI classes (model only); and the relabels Technology Management Service, Technology Management Service Offering, and Service Instance.
- **Service Instance** (formerly Application Service, `cmdb_ci_service_auto`) now has siblings: Data, Connection, Network, Facility, and Operational Process service instances. Application

services can be created manually, or with Service Mapping by entry point, tags, or dynamic query.

- The **Digital System Model** has four layers: (Business) Process, Service (Logical), Functional, and Infrastructure.
- Business services and technology management services are each **single-level, not a hierarchy**, and both are operational CIs used for impact analysis in Incident, Problem, and Change. Business services are also used for change approvals.
- **Teams** is a related list on CIs that tracks group assignments by type (Change Group, Managed by Group, Support Group) without adding more fields. Setting the group field on the CI syncs the Teams list.

Key CSDM tables

Object	Table
Business Capability / Business Application / Business Process	`cmdb_ci_business_capability` / `cmdb_ci_business_app` / `cmdb_ci_business_process`
SDLC Component / API	`cmdb_ci_sdlc_component` / `cmdb_ci_api`
Product models (all ten types)	`cmdb_model` and its extensions
Value Stream	`cmn_value_stream`
Product Idea / Planning Item	`sn_align_core_product_idea` / `sn_align_core_planning_item`
Strategic Plan / Strategic Priority / Goal / Target	`sn_gf_plan` / `sn_gf_strategy` / `sn_gf_goal` / `sn_gf_goal_target`
Common data: Building, Business Unit, Company, Department, Group, Location, User	`cmn_building`, `business_unit`, `core_company`, `cmn_department`, `sys_user_group`, `cmn_location`, `sys_user`

- **Common data elements are not CIs.** Before loading them, ask: is there a trusted source, are there several, how often does it change, is it deep enough for your CIs, and who maintains it? The location table has **Source ID**, **Location Type** (Region, Country, State/Province, City, Site, Building/Structure, Floor, Room), and **Managed by Group**.
- **Product models aren't CIs** either. CIs point to them with the **Model ID** field. The ten base types: Application, System Component (version-specific; includes Software Component, Service Offering, Firmware, and AI System Product models), Service, Software

(SAM entitlement only), Content, Contract, Facility, Hardware, Consumable, and Enterprise Good.

- Design & Planning tables (Business Capability, Business Application, Information Object) are used by Enterprise Architecture and are **not targets of Incident, Problem, or Change**.
- Ideation & Strategy: a **Product Idea** can be promoted into demand, project, epic, or story; a **Strategic Plan** holds the mission, vision, and values; **Goals** are qualitative and **Targets** measure them. Common persona: Product Owner.

CSDM life cycle stages

Life Cycle Stage (the phase) plus Life Cycle Stage Status (the step within it) together give an object's status. Added in Paris, provided by ServiceNow, and **not editable**; each object only shows the choices that fit it.

Process	Stages
Product (models only)	Ideation, Design, Operational, End of Life
Tangible / physical (assets and CIs)	Ideation (on the model), Purchase, Inventory, Deploy, Operational, Defective, Missing, End of Operation, End of Life
Intangible / logical	Ideation, Purchase, Design, Inventory, Deploy, Operational, End of Operation, End of Life
Document (contracts, business process)	Ideation, Operational, End of Life
Location	Design (Ideation), Operational, End of Life

- A discovered CI defaults to **Operational / In Use**. Tangible End of Life statuses include Disposed, Donated, Sold, RMA, Vendor Credit, Buy-out, and Lease Return.
- **Product Instance 2.0** (Xanadu) syncs stage and status across the CI, asset, and install base item, and **disables** the old legacy-status sync. Turning it off later is hard and can lose data, so test in non-production first.
- There's no rush to migrate, but some features (like CMDB Data Manager) require the life cycle mapping to be enabled. Check reports, scripts, business rules, and integrations that use the legacy fields first.

Migrating into CSDM

1. Back up your table data.
2. Map attributes: best practice, keep, refactor, or don't need.
3. **Run the dependency script** (the most important step). It finds the reports, business rules, and scripts tied to the old table. It doesn't move data.
4. Refactor attributes on the target tables.
5. Migrate the data: change the class, fix dependencies, reload, validate.

Properties and tables

The easiest points on the exam. Drill these last.

Property	Purpose and default
`glide.required.attribute.enabled`	Enforce mandatory fields in the IRE. True. Set false to stop failures.
`glide.identification_engine.multisource_enabled`	Enables CMDB 360. False.
`glide.identification_engine.skip_duplicates` / `.threshold`	True / 5. Update oldest, tag the rest.
`glide.class.upgrade|downgrade|switch.enabled`	Automatic reclassification. True. False creates tasks.
`glide.identification_engine.update_without_*_enabled`	Update attributes but skip reclassification. False.
`glide.identification_engine.enable_identifier_optional_condition`	Optional conditions on non-lookup fields
`glide.identification_engine.non_cmdb_tables`	Non-CMDB tables the IRE and IH-ETL can load
`glide.cmdb.query.nlq.activated`	Intelligent Search ↔ Query Builder. True.
`glide.cmdb.log.source.identification_engine`	Logs reclassification events

Table	Holds
`cmdb_multisource_data`	CMDB 360: every source's values
`cmdb_datasource_last_update`	Data Source Histories
`cmdb_health_metric_status`	Health job status
`cmdb_ire_reclassification_restriction`	Reclassification restriction rules
`life_cycle_mapping`	Legacy status → CSDM stage/status
`sn_cmdb_ws_nlq_sample_search`	Intelligent Search samples
`cmdb_ci_hardware` / `cmdb_ci_appl`	Default Hardware / Application identification rules
`cmdb_ci_information_object`	Information Objects

Off the top of your head

Answer each one before opening it.

Govern

What must you do before the Health Dashboard shows scores?

Activate its scheduled jobs; they're off by default.

Four types of horizontal discovery?

Network, CI, Cloud, and Serverless.

Hardware identification rule order?

Serial number with serial number type, serial number, name, then MAC address with name.

Two steps to turn on CMDB 360?

Activate the ITOM Discovery License plugin, then set glide.identification_engine.multisource_enabled to true.

What do Retire, Archive, and Delete policies need for each targeted class?

An active life cycle rule.

Metrics in each scorecard?

Completeness: Required, Recommended. Compliance: Audit. Correctness: Duplicate, Orphan, Stale.

Default staleness rule?

60 days on cmdb_ci, based on sys_updated_on.

Orphan rules in base, and per class?

Zero in base. One per class.

Which audit types count on the Compliance scorecard?

Desired State and Scripted.

Three objects for a Desired State audit?

Certification Filter, Certification Template, Audit.

Which health tasks are created by default?

Only de-duplication tasks.

Remediation order?

Task → Rule → Workflow.

Which reclassifications can lose data?

Downgrade and Switch.

Upgrade property is false and a Server is identified as Windows Server?

No upgrade; a reclassification task is created.

Recommended Main CI?

Oldest, most relationships, most recently updated.

Data Manager roles?

data_manager_admin and data_manager_user.

What must happen before Archive or Delete?

The CI must be Retired (End of Life).

Which legacy sync is one-way?

Asset State → CI Install Status. Not the reverse.

Insight

Unified Map combines which views?

Dependency View and Application Service maps.

Three Query Builder query types?

CMDB, Service Mapping, Combination.

Data Foundations Dashboards: base or Store? Configurable?

Store app. Fixed metrics.

Which dashboard is for senior leadership?

The CMDB Health Dashboard.

What field fills the Incident assignment group?

The CI's Support group. On Change, the Change group.

Ingest

Discovery phases?

Scan, Classification, Identification, Exploration.

Only method that captures software utilization?

Agent Client Collector.

Only free Service Graph Connector?

Microsoft SCCM.

How do Import Sets use the IRE, and what happens to coalesce?

onBefore script calling CMDBTransformUtil. Coalesce is ignored.

Offering vs. CI Class Manager for Managed By Group?

The Technology Management Offering wins.

Configure

Roles to create a class?

admin, or sn_cmdb_admin + personalize_dictionary.

Do reconciliation rules ship in baseline? What do they control?

No. Updates, not inserts.

How do you block inserts from a source?

IRE Data Source Rules.

Duplicate vs. overloaded?

Duplicate: many records, one thing. Overloaded: one record, many things.

Static or dynamic rule wins?

Dynamic.

CSDM

Role needed for scripted audits?

certification_admin.

IH-ETL guided setup steps, in order?

Basic Details, Preview and Prepare Source Data, Select CMDB Classes and Map Source Data, Preview and Schedule.

Five CSDM life cycle processes?

Product, Tangible/Physical, Intangible/Logical, Document, and Location.

Are common data records and product models CIs?

No. CIs reference product models through Model ID.

Seven domains?

Foundation, Ideation & Strategy, Design & Planning, Build & Integration, Service Delivery, Service Consumption, Manage Portfolio.

Stage for Business Capability and Information Object?

Fly.

Stage for Technology Management Services?

Walk.

Most important CSDM migration step?

Running the dependency script.

ServiceNow CIS-Data Foundations practice exam

286 questions grouped by exam domain. Answer key at the end.

Select exactly as many answers as each question asks for.

Configuration 54 questions

1. An admin ties several Technology Management Service Offerings to Dynamic CI Groups so the member CIs stay aligned with the right teams. Which group fields does the offering sync onto those CIs? (Choose two.)

A. Approval Group
B. Managed by Group
C. Owned by Group
D. Support Group

2. Pair each way of building an application service with the situation it suits best.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Service Mapping (Top-down)	1. Mission-critical services that need a precise, pattern-based approach
B. Tag-Based	2. Cloud-native, container-based, or virtual machine environments
C. Service Mapping (Connection Suggestion)	3. Custom-built apps, using application fingerprinting to produce maps quickly
D. Dynamic CI Group	4. Small services that are easy to group with filters and CMDB queries

3. You need a new custom CI class for an application your team now manages. Which roles are the minimum needed to create it?

A. itil_admin and personalize_form
B. data_classification_admin and personalize_dictionary
C. sn_cmdb_admin and personalize_dictionary
D. cmdb_inst_admin and personalize_form

4. Which relationship type correctly connects an application CI to the server it lives on?

A. Application>Runs on::Runs>Server
B. Application>Used by::Uses>Server
C. Application> Uses::Used by>Server
D. Application>Runs::Runs On>Server

5. Laptops imported into the CMDB fill in Assigned to on the Computer [cmdb_ci_computer] CI. The asset team wants the same value on the linked asset record, automatically. What does the configuration analyst do?

A. Write a business rule on the computer table that copies Assigned to onto the asset on insert or update
B. Create an Asset-CI Field Mapping rule that copies Assigned to between the asset and its CI
C. Hide Assigned to on the asset and add a new field that dot-walks to the CI

6. You need a new CI class for IoT sensors. What should you do? (Choose two.)

A. Delete an unused class and put the new one in its place
B. Install or update the CMDB CI Class Models store app and confirm the class doesn't already exist
C. Create the class under a suitable parent class
D. Repurpose an existing class

7. The CM team makes sure every managed CI has a Change Group so Change forms fill in automatically when a CI is picked. When a CI is picked on an incident, which base field fills in automatically?

A. Assignment group
B. Managed by group
C. Approval group
D. Change group

8. Each application server can host several applications, and an application can run on several application servers. What is the relationship between Application Server [cmdb_ci_app_server] and Application [cmdb_ci_appl]?

A. One-to-one
B. One-to-many
C. Many-to-one
D. Many-to-many

9. In CI Class Manager, which tab under Attributes shows only the attributes created on the Windows Server class itself?

A. All
B. Child
C. Added
D. Derived

10. An asset manager wants asset and CI records to stay in sync automatically. What should be in place? (Choose two.)

A. The active business rule that updates CI fields when the asset changes
B. The active business rule that updates asset fields when the CI changes
C. A strict one-to-one physical mapping between every asset and CI
D. Scheduled sync jobs that run after hours

11. The CMDB is healthy, discovered data is integrated, support groups are filled in, and Service Mapping makes it service-aware. When an agent selects a CI on an incident, which incident field fills in from the CI's support group?

A. Assignment Group
B. Change Group
C. Managed by Group
D. Approval Group
E. Support Group

12. You're proposing how to bring custom CIs into the CMDB. The approach must follow best practice, limit technical debt, and survive upgrades. Which options fit? (Choose two.)

A. Repurpose a base CI class and rename its attributes
B. Install or upgrade the CMDB CI Class Models store app and look for an existing class that fits
C. Extend an existing asset class to hold the new CI attributes
D. Extend an existing CI class to hold the new attributes

13. You want CI reference fields, such as the one on Incident, to offer only CIs from principal classes. Where do you set principal classes?

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

14. A class's identification rule needs changing. Where can you view and edit it? (Choose two.)

A. CI Class Manager
B. CI Identifiers module
C. IRE Application
D. API Integrations

15. Which CI Class Manager checkbox limits the CIs shown on Incident, Problem, and Change to the relevant managed classes?

A. Main Record
B. Principal Class
C. Independent
D. Extensible

16. Your team wants to start using CMDB 360 (Multisource CMDB). Which system property turns it on?

A. glide.identification_engine.multisource_cmdb_ci_enabled
B. glide.identification_engine.multisource_enabled
C. glide.identification_engine.multisource.query.max.limit
D. glide.identification_engine.multisource_non_cmdb_ci_ena

17. You're documenting several classes for stakeholders: their attributes, IRE rules, and suggested relationships. Where can you find all of this in one place?

A. CMDB Workspace
B. CI Class Manager
C. CMDB Data Manager
D. CI Identifiers

18. In CI Class Manager you see a Health Inclusion Rules tab on a class. What are these rules used for?

A. Limiting which CIs are included in the evaluation
B. Reconciling particular attributes by data source
C. Reducing how much data is ingested into the CMDB

19. The Windows team wants to group its CIs with CMDB groups. Which ways can you define the group? (Choose two.)

A. Encoded queries
B. Scripted queries
C. Saved CMDB queries
D. Tag-based queries

20. For one class, CI Class Manager sets Managed by Group = Enterprise IT Services. A Technical Service Offering for the same CIs sets Managed by Group = Windows Support and Change Group = Change Management Team. What Managed by Group do the CIs end up with?

A. Change Management Team
B. Enterprise IT Services
C. Windows Support

21. Which group fields on CIs are kept in sync automatically through Technology Management Service Offerings and Dynamic CI Groups? (Choose two.)

A. Change group
B. CMDB group
C. Approval group
D. Support group

22. You need a CI class for a brand-new type of equipment, and it must fit cleanly into the existing CMDB structure. What do you do?

A. Create it in CI Class Manager with no parent, since classes should stand alone
B. Add equipment-specific fields to an existing class
C. Define it in the Service Catalog, since CI Class Manager is only for services
D. Create the class in CI Class Manager under a relevant parent so it inherits properly

23. How do you mark a class as a principal class so fewer classes appear in an incident's Configuration Item field?

A. Select Principal Class on the class's Attributes tab
B. Select Principal Class on the class's Basic Info tab in CI Class Manager
C. Set a Principal Class attribute on each CI
D. Select Principal Class in CMDB Workspace

24. You clicked Enable Life Cycle Sync to adopt the CSDM life cycle field mappings. What should you do next?

A. Activate the CSDM life cycle field mappings
B. Change Life Cycle Stage values to match the legacy values
C. Fix any incomplete mappings shown in the discrepancy report

25. Which CI Class Manager location lets you mark a class as principal?

A. Class Info > Basic Info
B. Class Info > Attributes
C. Health > Attributes

26. Following best practice for new custom CMDB tables, what prefix do their names start with?

A. u_cmdb_ci
B. cmdb_ci
C. ci_cmdb
D. u_ci_cmdb

27. Duplicate server records keep appearing. The IRE identifies servers with identification rules. Where do you configure those rules?

A. CMDB Data Manager
B. CMDB Workspace
C. CMDB CI Class Manager
D. CMDB Health Dashboard

28. CIs are missing their support group. What can fill it in automatically and keep it current?

A. Technology Management Service (Technical Service)
B. CI Class Manager
C. Technology Management Service Offering (Technical Service Offering)
D. Dynamic CI group

29. The Tomcat identification rule uses Class and Install Directory. Which identifier entry setting lets the IRE try the parent Application rule when the Tomcat rule finds no match?

A. Criterion attributes
B. Applies to
C. Allow fallback to parent's rules
D. Independent

30. A class owner has to change the icon for the UNIX Server class. Which CI Class Manager tab does this?

A. CI List
B. Attributes
C. Basic Info
D. Suggested Relationships

31. Where do you review how legacy status values translate into CSDM life cycle stage and life cycle stage status values?

A. Life cycle choice list
B. Life cycle mappings
C. Life cycle properties

32. New IoT classes need to be added to, and outdated classes removed from, the principal class filter. Which roles are required? (Choose two.)

A. sn_cmdb_admin
B. sn_csdm_admin
C. personalize_dictionary
D. cmdb_query_builder

33. A Dynamic CI Group holds every Windows server due for an upgrade. What happens when that group is selected in a change's Configuration Item field?

A. Its CIs show in the Affected CIs related list
B. Its CIs show in the Impacted Services related list
C. Impact is calculated against the group itself

34. You're on a class's Basic Info tab in CI Class Manager. How do you make it a principal class?

A. Choose Yes in a Principal Class choice list
B. Select the Principal Class checkbox
C. Click a Principal Class UI Action

35. Data from Discovery, SCCM, AWS, and manual uploads classifies the same CIs differently, causing duplicates and inconsistency. What should the team do? (Choose two.)

A. Let each source define its own classes, even if they conflict
B. Write a custom script to fix data before it is stored
C. Use CI Class Manager to standardize classes and attributes across every source
D. Use identification and reconciliation rules to prevent duplicates and standardize classes

36. The owner of the Windows Server class wants quick access to it in CI Class Manager. Which feature helps?

A. Search CI Classes
B. CI Favorites
C. Pinned Classes

37. Change wants one change request to cover many CIs at once. Which field on the change gets the Dynamic CI Group?

A. Service Offering
B. Configuration Item
C. Business Service

38. A new application identification rule has to handle the same app installed more than once on one server. Which criterion attribute works best?

A. Configuration File Path
B. Version
C. Configuration File Name
D. Class
E. Port

39. You're setting up CMDB 360. Pair each system property with what it does. Some descriptions won't be used.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Sets the most CIs a CMDB 360 recompute can include	1. glide.identification_engine.multisource_enabled
B. Captures CMDB 360 data for CIs in CMDB classes	2. glide.identification_engine.multisource_cmdb_ci_enabled
C. Turns on CMDB 360	3. glide.cmdb.logger.source.cmdb_multisource
D. Captures CMDB 360 data for CIs in non-CMDB classes	4. glide.identification_engine.multisource.recompute.max.ci.limit
E. Turns on logging for CMDB 360	

40. A team at the Crawl stage of the CSDM journey wants to follow the model. Which domains does that stage cover? (Choose two.)

A. Foundation
B. Design
C. Manage Technical Services
D. Sell/Consume

41. Asset and CI life cycles are synchronized. When the Hardware Status field changes on a CI, which field on the linked asset updates automatically?

A. Substate
B. Installed
C. State
D. Asset tag

42. The CMDB now has several custom classes for near-identical devices (for example, u_cmdb_ci_custom_router and u_cmdb_ci_special_switch), causing fragmented data and reporting problems. What should the team do? (Choose two.)

A. Delete the custom classes and rebuild the CIs by hand
B. Keep every custom class to avoid migration risk
C. Move the CIs into standard classes and map their attributes correctly
D. Use CI Class Manager to consolidate the custom classes into standard classes

43. Only principal CIs should be selectable on change forms. Which class is the best candidate for principal class?

A. 'Disk' Class
B. 'Network Adapter' Class
C. 'Windows Server' Class

44. When adding a reconciliation rule in CI Class Manager, there's no option to choose Static or Dynamic. How do you make those options appear?

A. Change the glide.dynamic_reconciliation.allowed_classes property
B. Turn on the glide.identification_engine.multisource_enabled property
C. Elevate to security_admin
D. Enable a Dynamic Rules option in the Reconciliation Rules module

45. You create a custom child class of Application [cmdb_ci_appl] for a component no existing class fits, and all its CIs will come through the IRE. What must exist for the class, created or inherited? (Choose two.)

A. Dependent Relationship
B. Identification Rules
C. Reconciliation Rules
D. Custom Attributes
E. Data Refresh Rules

46. From a CI record, which tool is used to create and edit relationships with other CIs?

A. CI Class Manager
B. Suggested Relationships
C. Unified Map
D. Relationship Editor

47. In the baseline system, what keeps fields synchronized between CIs and their assets?

A. Business Rules
B. Scripts Includes
C. Scheduled jobs
D. UI Actions

48. The hardware identification rule is evaluating a new Windows server. Which order does it try identifier entries in (ignoring product_instance_id)?

A. Serial number with serial number type, serial number, name, MAC address with name
B. MAC address, IP address, name, serial number
C. Serial number, IP address, name, MAC address
D. Name, serial number, MAC address and name, serial number with type

49. Why are applications usually identified differently from hardware?

A. Applications can only be created manually
B. Applications always have unique serial numbers
C. Applications skip the IRE
D. Software rarely has a unique serial number, so it's identified by the process it runs

50. Which two steps turn on CMDB 360 (Multisource CMDB)? (Choose two.)

A. Activate the ITOM Discovery License plugin
B. Set glide.identification_engine.multisource_enabled to true
C. Install the CMDB CI Class Models app
D. Activate the Health Dashboard scheduled jobs

51. Before the Paris release, what did the CMDB track about data sources once reconciliation rules existed?

A. Nothing at all
B. Only the last source to update each record, in Data Source Histories
C. Every value from every source, in cmdb_multisource_data
D. Only rejected values

52. When does reconciliation happen in the IRE process?

A. After identification, on updates to existing CIs
B. Before identification, on inserts only
C. Only when a CI is created
D. Only during reclassification

53. Which group is described as the team that sets direction for the CMDB and decides whether it is ready to be used?

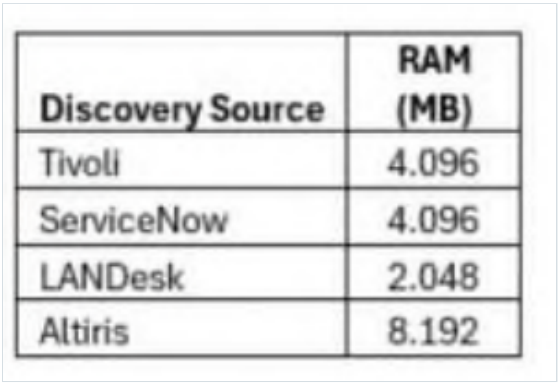
A. The Configuration Control Board
B. The Discovery team
C. The Change Advisory Board
D. The Service Desk

54. Who owns the data stored in the CMDB?

A. The data owners for each area, who must be part of CMDB governance
B. The ServiceNow platform owner only
C. Whoever loaded the data last
D. The CMDB team

Ingest 46 questions

55. Two dynamic reconciliation rules exist: the Server class uses Largest Value for RAM, and the Windows Server class uses Most Reported for RAM. Based on the multisource data shown, what RAM value is stored for a Server CI?



Discovery Source	RAM (MB)
Tivoli	4.096
ServiceNow	4.096
LANDesk	2.048
Altiris	8.192

A. 4.096 MB
B. 2.048 MB
C. 8.192 MB

56. You need to bring outside data into the CMDB without bypassing the Identification and Reconciliation API, to avoid duplicates and unauthorized updates. Which method is recommended?

A. Table API (REST API or SOAP API)
B. Import Sets and Transform Maps
C. IntegrationHub ETL

57. Several data sources are allowed to update the same class and attributes. How do you decide which source is the authoritative one?

A. Set data refresh rules with a time period
B. Give each data source a priority in the reconciliation rules
C. Run the data source updates by hand in the right order
D. Give each source a run order in the identification rules

58. With CMDB 360 turned on, dynamic reconciliation rules become available. Which dynamic rule types can you pick in the Create Reconciliation Rule wizard? (Choose two.)

A. Smallest Value
B. Last Updated
C. Last Created
D. Most Reported

59. Look at the identification rule shown. Two CIs are imported into the hardware class: CI1 matches an existing CI by name, and CI2 matches an existing CI by IP address. What happens?

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

A. CI1 is inserted as new and CI2 updates the match
B. Both update their matches
C. CI1 updates the match and CI2 is inserted as new
D. Both are inserted as new

60. Discovery moves existing Server [cmdb_ci_server] records into the more specific Linux Server and Windows Server classes. What kind of reclassification is this?

A. Class Upgrade
B. Class Switch
C. Class Downgrade

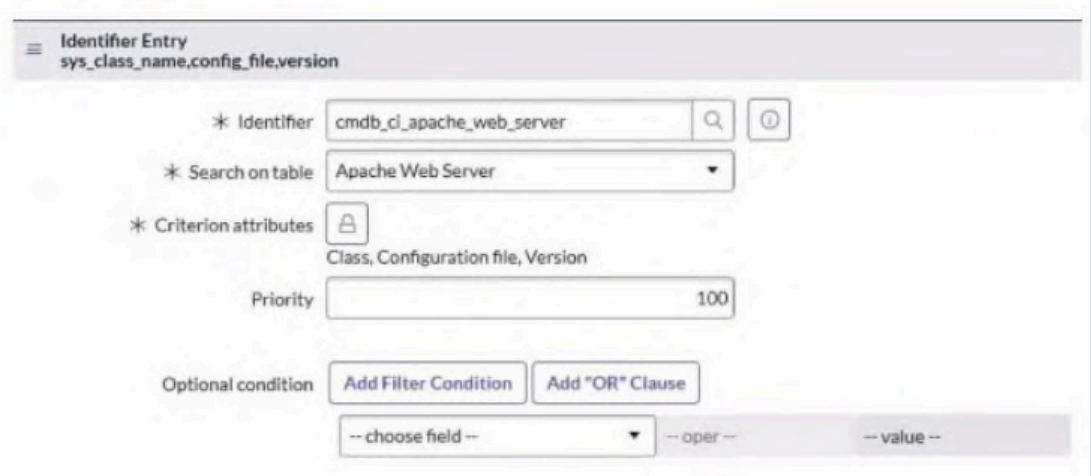
61. A data owner wants reconciliation rules across five discovery sources, but the company has no ITOM Discovery license for CMDB 360. What can they do?

A. Buy ITOM Discovery first to get multisource IRE rules
B. Use IRE reconciliation rules by discovery source, which work without CMDB 360
C. Use CMDB 360 right away since it's a platform feature

62. Which ServiceNow solution creates relationships automatically?

A. Discovery
B. Workflow Studio
C. IntegrationHub ETL
D. Service Mapping

63. The Apache Web Server identification rule uses Class, Configuration file, and Version. Yesterday Service Mapping discovered an Apache CI. Today the owner upgraded Apache and discovery ran again. What happens in the CMDB?



A. A duplication error occurs
B. The existing CI is reconciled and its version updated
C. The CI is reclassified as a generic Web Server
D. A new Apache Web Server CI is created

64. Which are typical reasons to use Agent Client Collector (ACC)? (Choose two.)

A. Servers in the data center
B. Network devices in the DMZ
C. Devices that connect to the network only now and then
D. Devices in secure environments

65. Warranty expiration dates on hardware CIs aren't collected by discovery. How do you keep them accurate for every CI?

A. Use a scheduled data import from an external source
B. Have the reconciliation engine update them
C. Create a separate class to hold non-discoverable attributes

66. Pair each CMDB ingestion tool with how it works. Some descriptions won't be used.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Third-party integrations from other vendors	1. Agent Client Collector
B. Automated agent-based collection running patterns	2. Import Sets
C. An in-house solution built on transform maps	3. Service Graph Connector
D. Prebuilt Store integrations that need little customization	4. ServiceNow Discovery
E. Automated agentless collection running patterns	

67. When you load CMDB data with import sets and transform maps, which transform script type sends the data through the IRE?

A. onAfter
B. onStart
C. onComplete
D. onBefore

68. Look at the reconciliation rules shown for ServiceNow, Altiris, and SCCM. Which statements are true? (Choose two.)

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

A. SCCM can update any Windows Server record because it has the highest priority number
B. SCCM can insert new Windows Server records
C. ServiceNow can insert Windows Server records but can't update ones created by Altiris or SCCM
D. Altiris can update Windows Server records that SCCM inserted

69. You customized a script transform in an installed Service Graph Connector. What happens on a later upgrade that changes the default script transform?

A. A skipped change is logged and your script transform is left as is
B. The connector upgrade won't start
C. The upgrade stops with an error

70. The IRE moves a server from Server [cmdb_ci_server] to Windows Server [cmdb_ci_win_server]. What was this?

A. Class Switch
B. Class Change
C. Class Upgrade
D. Class Downgrade

71. Pair each product with what it does.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Service Mapping	1. Finds devices and applications on the network and keeps the CMDB current
B. Service Graph Connectors	2. Connects ServiceNow to external systems to import and sync data
C. ServiceNow Discovery	3. Maps a service's full topology and the infrastructure and apps behind it
D. Agent Client Collector (ACC)	4. Gives real-time endpoint configuration data to keep the CMDB current

72. A high-priority source updates server records weekly but stopped because of an integration problem. What lets a lower-priority source take over after a set period of inactivity?

A. Data Refresh Rules
B. Reconciliation Rules
C. Identification Rules
D. Health Inclusion Rules

73. After importing data with a Service Graph Connector, you want to see the field-to-field mappings. Where do you look?

A. Integration Hub
B. CMDB Integrations Dashboard
C. IntegrationHub ETL

74. What does IntegrationHub ETL do? (Choose two.)

A. Brings third-party data into CMDB or non-CMDB tables
B. Runs discovery and updates the CMDB
C. Sends data through the IRE
D. Imports SCCM and Intune data specifically

75. You need SCCM data in the CMDB as quickly as possible. Which method gets value fastest?

A. Import Sets
B. Service Graph Connectors
C. IntegrationHub ETL
D. Agent Client Collector

76. An admin changes the query in the SCCM Service Graph Connector. What's the effect?

A. Its data source is marked inactive
B. Its scheduled jobs must be reconfigured
C. Updates to that connector are skipped during upgrades

77. Vendor shipment files are imported into the CMDB with import sets, and duplicates keep appearing. How do you stop them?

A. Coalesce on two field maps
B. Set a property that makes transform maps use the IRE
C. Create comparison rules in the IRE
D. Call the CMDBTransformUtil API in the transform script

78. Which Service Graph Connector categories are available? (Choose two.)

A. Observability
B. DevOps
C. Workflow Automation
D. Cloud

79. Server uses a Lowest Value rule for Disk Space (GB) and Windows Server uses Most Reported. Based on the data sources shown, which Disk Space value is stored on a Windows Server record?

Data Source	Disk Space (GB)
Tivoli	75
ServiceNow	75
LANDesk	50
Altiris	80

A. 50
B. 75
C. 80

80. Which way of bringing third-party CIs into the CMDB carries the least technical debt?

A. Service Graph Connector
B. Vendor-provided integration
C. Table API
D. Import Sets and Transform Maps

81. Which solutions automatically create relationships between application CIs within an application service? (Choose two.)

A. IntegrationHub ETL
B. Data Manager
C. Discovery
D. Event Management
E. Service Mapping

82. What must be enabled before you can use dynamic reconciliation rules?

A. CMDB Data Manager
B. CMDB 360/Multisource CMDB
C. Reconciliation Rules
D. CMDB Workspace

83. You want the fastest time to value for bringing Windows computer data from SCCM into the CMDB. What do you use?

A. An import set with a JDBC data source and transform maps
B. The SCCM Service Graph Connector
C. The SCCM Usage Metering spoke
D. IntegrationHub ETL to SCCM with the Robust Transform Engine

84. You installed a Service Graph Connector and then changed its mappings. What's the consequence?

A. Fields filled by the customization get a special tag
B. The connector won't run without an approval on the changed mapping
C. ServiceNow doesn't support the customized mappings; you support them yourself

85. After installing the Service Graph Connector Central plugin, which workspace gets the new Service Graph Connector Central tab?

A. CMDB Workspace
B. Discovery Admin Workspace
C. Service Graph Connector Workspace

86. Discovery finds a web server instance and the server under it. What relationship is created between them by default?

A. Web Server > Runs::Runs On > Server
B. Web Server > Used by::Uses > Server
C. Web Server > Uses::Used by > Server
D. Web Server > Runs on::Runs > Server

87. Which ServiceNow feature offers prebuilt integrations designed for simple connections to external systems?

A. Identification and Reconciliation Engine (IRE)
B. Agent Client Collector (ACC)
C. Service Graph Connectors
D. Service Mapping

88. Which ServiceNow Store application imports data from different sources into the CMDB?

A. Integration Hub ETL
B. Agent Client Collector
C. Service Mapping

89. The IRE finds duplicates of one CI. Which one does it update?

A. The newest, by created date
B. The one updated most recently
C. The oldest, by created date

90. Facilities wants to keep location data [cmn_location] in ServiceNow but worries other sources will create duplicates. Why use the IRE? (Choose two.)

A. It blocks updates from unwanted sources
B. It uses relationships to decide insert or update
C. It blocks inserts from unwanted sources
D. It uses unique identification to decide insert or update

91. Why would a configuration manager use the Robust Transform Engine (RTE) for imports?

A. It inserts records into the right classes and avoids duplicates
B. It finds existing CIs and updates them with the latest data
C. It takes the definition of how data is converted from an object and passes it to the IRE
D. It loads external data into the staging table

92. When using CMDBTransformUtil() in an import of non-discoverable data, which transform map fields must be filled in? (Choose two.)

A. Serial Number field
B. Source field
C. Coalesce field
D. Target field

93. Which are types of horizontal discovery? (Choose three.)

A. Network discovery
B. Tag-based discovery
C. Cloud discovery
D. CI discovery
E. Top-down discovery

94. You want Discovery to stop detecting a type of CI your team doesn't manage. Where do you go?

A. System Import Sets > Data Sources
B. Discovery Definition > Configuration Console
C. CI Class Manager > Basic Info
D. CMDB Workspace > Insights

95. Which ingestion tool is free to use, according to the course? (Choose two.)

A. Agent Client Collector
B. IntegrationHub ETL
C. ServiceNow Discovery
D. Import Sets and Transform Maps

96. Which data is typically discoverable rather than non-discoverable?

A. IP address
B. Owned by
C. Support group
D. Change group

97. In the IntegrationHub ETL guided setup, which step comes right after Basic Details?

A. Preview and Prepare Source Data
B. Preview and Schedule
C. Integrate
D. Select CMDB Classes and Map Source Data

98. Which Store app adds a dashboard with performance and error metrics for every Service Graph Connector and custom IH-ETL integration?

A. CMDB Workspace
B. CMDB CI Class Models
C. Integration Commons for CMDB
D. Service Graph Connector Central

99. An admin customizes a Service Graph Connector's mappings. What changes?

A. Upgrades apply automatically over the changes
B. The customer becomes fully responsible for supporting the connector
C. ServiceNow keeps full support
D. The connector stops using the IRE

100. A transform map's onBefore script calls CMDBTransformUtil. What happens to the coalesce values on that transform map?

A. They take priority over the IRE
B. They create duplicates
C. They are required
D. They are ignored

Govern 79 questions

101. You want fewer stale CIs in the CMDB. Which CMDB Health Dashboard scorecard tracks them?

A. Completeness
B. Correctness
C. Compliance

102. Pair each CMDB Health Dashboard metric with its description.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Audits	1. Compare actual field values with expected values
B. Recommended fields	2. Keep their use to a minimum
C. Orphan CIs	3. Attribute values or relationships are missing
D. Stale CIs	4. Good to have filled in because they help with troubleshooting
E. Duplicate CIs	5. Haven't been updated in a while and may be out of date
F. Required fields	6. Found during identification and reconciliation, with built-in remediation tools

103. You've been asked to clean up duplicate CIs. What is the preferred place to manage this work?

A. The de-duplication dashboard in CMDB Workspace
B. A de-duplication task module
C. My Tasks in the Application Navigator

104. Two CIs are duplicates: one came from a manual import, one from discovery. The discovered CI has the current IP address; the imported one has the correct relationship to a critical business application. How do you resolve it with the Duplicate CI Remediator?

A. Keep the imported CI and delete the discovered one
B. Keep the discovered CI and merge in the relationship from the imported CI
C. Merge automatically, keeping every attribute from the discovered CI
D. Keep the discovered CI and delete the imported one

105. You want to remove every Linux server that hasn't been updated in six months. What does Data Manager recommend?

A. A scheduled job
B. An archive policy
C. A business rule

106. The team wants the Compliance scorecard to show meaningful results for servers that aren't on the latest patch. What must be configured?

A. Technical service offerings, Dynamic CI Groups, and CMDB groups
B. Stale, orphan, and duplicate rules
C. A certification filter, a certification template, and an audit
D. Certification policies, data filters, scheduled jobs

107. Leadership wants duplicate, orphan, and stale CI numbers. Which Health Dashboard scorecard has them?

A. Correctness
B. Compliance
C. Completeness

108. How do attestation and data certification policies differ?

A. Attestation requires fixing specific attributes; certification confirms the CI still exists
B. Attestation can be scheduled; certification can't
C. Attestation can go to a group or person; certification only to a person
D. Attestation confirms the CI still exists; certification requires validating specific attributes

109. Many hardware CIs are missing serial numbers, which could lead to duplicates. Which ServiceNow guidance helps you troubleshoot and fix this?

A. CSDM Data Foundations Dashboard playbooks
B. CSDM Now Create playbooks
C. CMDB Health Dashboard playbooks
D. CMDB Data Foundations Dashboard playbooks

110. Which data quality areas does the CMDB Health Dashboard cover? (Choose two.)

A. Upgraded CIs
B. Missing CIs
C. Stale CIs
D. Duplicate CIs
E. Downgraded CIs

111. For 60,000 servers, the Correctness scorecard shows Duplicate as 59,000/60,000 healthy/evaluated but Orphan as 45,000/50,000. What explains Orphan evaluating fewer servers?

A. A CMDB group is set for Duplicate on the Server class
B. A CMDB group is set for Orphan on the Server class
C. A health inclusion rule is set for Orphan on the Server class
D. A health inclusion rule is set for Duplicate on the Server class

112. Class owners see tasks on the My Work tab of CMDB Workspace. Which tool creates these tasks?

A. CMDB Data Manager
B. De-duplication templates
C. CMDB Health Dashboard

113. What keeps the amount of data in the CMDB manageable?

A. Archive Policies
B. Business Rules
C. Scheduled Jobs

114. What has to happen before the CMDB Health Dashboard produces health scores?

A. Install a CMDB health calculation plugin
B. Nothing; scores calculate by default
C. Activate the Health Dashboard's scheduled jobs

115. Group ownership of classes is set in CI Class Manager, and a Data Manager policy should use it. Which group field do you set on the policy?

A. Approval Group
B. Change Group
C. Support Group
D. Managed By Group

116. Where in CMDB Workspace do you find de-duplication tasks created from IRE-processed data?

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

117. Which Data Manager policy type creates tasks that let the assignee update fields on the CI?

A. Certification
B. Compliance
C. Audit
D. Attestation

118. A new customer is staffing its Configuration Management team. Pair each role with its description.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. CI Analyst	1. Read-only access to CMDB data and basic views like reports and dashboards
B. CMDB Process Owner	2. Accountable for every element of a portfolio across its full life cycle
C. Service or Product Owner	3. Manages assigned CI tables, keeps records current, and resolves CMDB tasks
D. Configuration Manager/CMDB Admin	4. Holds the highest level of CMDB privileges

119. Pair each duplicate-CI tool with what it does. Some tools won't be used.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. De-Duplication Tasks	1. Assigned to groups to resolve duplicates
B. CMDB Health Dashboard Correctness Scorecard	2. Shows how many duplicate CIs exist in the CMDB
C. Certification Tasks	3. Resolves de-duplication tasks in bulk
D. Duplicate CI Remediator	4. A wizard for resolving de-duplication tasks one at a time
E. De-Duplication Templates	

120. You want tasks created and assigned automatically to confirm CIs still exist. Which policy type do you use?

A. Retire
B. Attestation
C. Certification
D. Delete

121. What is CMDB Data Manager for?

A. Automating archiving and deleting records based on retention policies
B. Encrypting archived records
C. Enforcing relationship rules between CIs

122. Where does a data manager create, publish, and manage policies that govern CI life cycle operations?

A. The Management tab in CMDB Workspace
B. Service Operations Workspace
C. CI Class Manager
D. The CMDB 360 tab in CMDB Workspace

123. Short-lived CIs not discovered in the past week should be removed. After retiring them, what is the recommended next step?

A. Create a delete policy
B. Create a scheduled job
C. Create a business rule

124. The CMDB Data Foundations Dashboard shows several problems and links to remediation playbooks. How do playbooks help?

A. They track issues and produce metrics automatically
B. They install into the instance and fix issues automatically
C. They help you analyze and fix issues

125. What documentation does a Data Foundations Dashboard playbook give you? (Choose two.)

A. Root Cause
B. Problem Overview
C. Automated Remediations
D. Problem Analysis

126. What are benefits of the playbooks on the CMDB Data Foundations Dashboard? (Choose two.)

A. Remediation templates for low-scoring metrics
B. Insight into the downstream impact of low-scoring metrics
C. Automated scripts that fix low-scoring metrics
D. Remediation options to improve low-scoring metrics

127. You're starting to populate the CMDB and need a way to remove data that's no longer useful. Which governance tool does that?

A. De-duplication Templates
B. CMDB and CSDM Data Foundations Dashboard
C. CMDB Data Manager
D. CMDB Health Dashboard
E. CI Class Manager

128. Which field does the Staleness metric on the Correctness scorecard use to measure age?

A. First discovered (first_discovered)
B. Most recent discovery (last_discovery)
C. Updated (sys_updated_on)
D. Last modified on (last_modified)
E. Created (sys_created_on)

129. You want to confirm that every server in a data center really exists. Which Data Manager policy type do you use?

A. Promotion
B. Certification
C. Attestation
D. Verification

130. With the baseline Data Manager policies, what must be true of a CI before it can be archived or deleted?

A. It's marked inactive
B. It's marked critical
C. It's operational and in use
D. It's retired and at end of life

131. The team wants to verify that every server in the CMDB physically exists in the data center. Which policy type do they create?

A. Certification
B. Delete
C. Archive
D. Retire
E. Attestation

132. Pair each Health Dashboard metric with its definition.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Audit	1. Records that represent the same asset more than once
B. Recommended	2. Fields needed to create or update a CI
C. Stale	3. Records that have lost their relationships with other CIs
D. Required	4. Records no longer updated but still stored
E. Orphan	5. Fields that improve the accuracy, completeness, and usefulness of CIs
F. Duplicate	6. Actual field values compared with the expected values in a template

133. The Health Dashboard's three KPIs are Correctness, Compliance, and Completeness. Pair one sub-metric with each KPI. Some sub-metrics won't be used.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Stability	1. Completeness
B. Required	2. Compliance
C. Certify	3. Correctness
D. Audit	
E. Orphan	
F. Suggested	

134. How do the governance tools in the CMDB help organizations? (Choose two.)

A. Lower hardware costs
B. Better service management operations
C. Better visibility and control
D. Help choosing integrations

135. Which policy types can you create in CMDB Data Manager? (Choose two.)

A. Reconciliation
B. Retire
C. Archive
D. De-duplication

136. With Assignment type set to User Group Field on a Data Manager policy, which group fields are available by default? (Choose two.)

A. Owned by Group
B. Managed By Group
C. Assignment Group
D. Support Group

137. During implementation, one person is responsible for CI data accuracy and completeness, ongoing data quality, and resolving discrepancies. Which role is that?

A. Service Owner
B. Configuration Manager
C. CMDB Architect
D. IT Asset Manager

138. A certification policy uses Table = Server [cmdb_ci_server] and Filter = Operating System contains "Server" OR Operating System contains "Linux". Which operating systems are included? (Choose two.)

A. Linux CentOS
B. AIX
C. Windows 2019 Datacenter
D. Windows Server 2022 Datacenter

139. Which Health Dashboard scorecard tracks which CIs and classes are missing key data?

A. Completeness
B. Correctness
C. Compliance

140. You have many similar de-duplication tasks to resolve in bulk. How?

A. Create and fix each task by hand
B. Create and run a de-duplication template
C. Use the Duplicate CI Remediator wizard
D. Write a background script

141. You want tasks created when operational Linux servers are missing Support Group or Managed By Group. Where do you configure the recommended fields and task creation?

A. CMDB Workspace and scheduled jobs
B. Technical service offerings and Dynamic CI Groups
C. Dynamic CI Groups and CMDB groups
D. CI Class Manager and health preferences

142. A report lists services with no owner. Which governance play from a playbook keeps this from happening again?

A. Default the Service Owner field so it's never empty
B. Make Managed by mandatory on every CI
C. Make the Owned by field mandatory

143. By default, how old must a CI be to count as stale on the Correctness scorecard?

A. 24 hours
B. 30 days
C. 7 days
D. 60 days

144. A regional configuration manager only sees overall scores and scores by class on the Health Dashboard. How do they get a score just for their region's CIs?

A. Turn on a "group scores by region" setting
B. Create CMDB groups of type Health for each region
C. Customize the Health Dashboard jobs to group by region

145. Where on the platform does a CMDB admin work to keep CI data quality high?

A. CMDB Workspace
B. Data Quality Business Rule
C. CMDB Audit Business Rule
D. Data Quality Scheduled Job

146. The Correctness scorecard shows Duplicate for the Server class as 59,000/60,000 healthy/evaluated. How many server records are in scope?

A. 1,000
B. 60,000
C. 59,000 Linux and 60,000 Windows
D. 59,000

147. What do you need to do to start using the CMDB Health Dashboard and its KPIs?

A. Activate the dashboard's scheduled jobs
B. Download it from the ServiceNow Store
C. Nothing; it's on by default
D. Turn on a system property

148. Which are CMDB Data Manager end-of-life policy types? (Choose two.)

A. Archive
B. Disposed
C. Lost
D. Decommission
E. Retire

149. Some hardware CIs lack serial numbers. The team needs a structured way to find and fix the gaps. What fits?

A. CMDB Data Foundation Playbook
B. CI Class Manager
C. Service Graph Connectors

150. Delete and archive policies will target servers in one data center. Which life cycle state must the servers be in to be affected?

A. Inventory, Available
B. Missing, Stolen
C. End of Life, Retired
D. Any life cycle stage

151. What matters most for successfully operationalizing a CMDB?

A. Clear governance and continuous health monitoring
B. Relying only on discovery tools
C. Letting IT teams change records freely
D. Loading as much data as possible

152. After you configure Data Manager policies, where do users review and work the tasks they create?

A. CMDB Workspace, My Work tab
B. CMDB Workspace, Management tab
C. Health Dashboard, Duplicate CIs tab
D. Health Dashboard, Audit tab

153. Policy tasks should be assigned using the group attribute set on the class in CI Class Manager. Which field is recommended?

A. Support group
B. Approval group
C. Change group
D. Managed by group

---

154. The Data Foundations Dashboard shows low scores on several metrics. What should you do? (Choose two.)

A. Use the remediation playbooks next to each metric
B. Add exclusion rules to raise the scores
C. Focus first on Critical and High priority metrics
D. Remove non-operational and retired CIs

---

155. Which is the most complete list of CMDB Data Manager policy types?

A. Attestation, Retire, and Certification
B. Archive and Delete
C. Retire, Archive, Attestation, Certification, and Delete
D. Delete, Attestation, Retire, and Certification

---

156. You want the Health Dashboard to show only classes that are ready to be managed. What do you use?

A. Health Inclusion Rules
B. Reconciliation Rules
C. Data Refresh Rules
D. Identification Rules

---

157. On the first tab of the Duplicate CI Remediator you pick the main CI. Which attributes help choose it? (Choose two.)

A. Least Related Items
B. Newest Created
C. Oldest Created
D. Most Related Items

---

158. You need a Data Manager policy that removes records from a class but lets you restore them within a set time. Which type?

A. Delete
B. Certification
C. Archive
D. Retire

---

159. Which roles let a user configure or work Data Manager tasks? (Choose two.)

A. data_manager_user
B. cmdb_dm_user
C. cmdb_dm_admin
D. data_manager_admin

---

160. When a parent CI's Life Cycle Stage changes, its child CIs should update automatically. How?

A. Use a transform map to update stages
B. Create a CI class for life cycle management
C. Set up life cycle stage automation in Data Manager

---

161. Which CMDB Workspace tabs give access to Data Manager for setting up CI life cycle policies? (Choose two.)

A. Management
B. My Work
C. CMDB 360
D. Insights
E. Home

---

162. The company must track compliance requirements like GDPR or HIPAA against CMDB objects. What approach fits?

A. Create custom compliance identifiers in the CMDB
B. Associate compliance identifiers with CSDM objects
C. Request compliance tracking through the Service Catalog

---

163. Life cycle sync is on. An ITAM admin sets several server assets to retired. What life cycle stage do the linked CIs get out of the box?

A. End of life
B. End of operation
C. Disposed

---

164. Which are CMDB Data Manager use cases? (Choose two.)

A. Retiring computers not updated in six months
B. Reporting on duplicate Linux servers in London
C. Finding which source has the most accurate RAM value
D. Confirming every server in the Seattle data center exists

165. While moving legacy status fields to life cycle objects, you configure the reverse sync choice option. Which setting do you choose?

A. True, so the legacy state changes the life cycle value
B. False, so the life cycle value changes the legacy state
C. False, so the legacy state changes the life cycle value
D. True, so the life cycle value changes the legacy state

166. After using the Get Well playbook to fix missing serial numbers on hardware CIs, you want a monthly report to keep watching that number. Which play is it?

A. Analysis
B. Fix
C. Data Governance

167. Pair each life cycle activity with what it involves.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Archive	1. CIs no longer in active use but kept for history or compliance
B. Attest	2. Policies that regularly verify CI data meets required standards
C. Certify	3. A regular process confirming the CI record is accurate
D. Delete	4. Removing obsolete CIs according to retention and compliance rules

168. You want to retire operational CIs that have no owner and are over a year old. What does Data Manager recommend?

A. Create a scheduled job
B. Create a policy
C. Create a business rule

169. Which data foundation pillars help minimize technical debt when deploying a CMDB? (Choose two.)

A. Ingest
B. Govern
C. Plan
D. Insight

170. A new admin opens the CMDB Health Dashboard and every scorecard is empty. What is the most likely reason?

A. The Health Dashboard needs a paid license
B. Health inclusion rules block every class
C. The CMDB Workspace plugin isn't installed
D. Its scheduled jobs haven't been activated, since they're off by default

171. A regional manager wants Health Dashboard scores just for Windows and Linux servers in one region. Which CMDB Group setting makes the group appear on the dashboard?

A. Group type = Regional
B. Group type = Health
C. Group type = Dynamic
D. Group type = Query

172. A Retire policy in CMDB Data Manager won't run against the Linux Server class. What is required for each class the policy targets?

A. An active life cycle rule
B. A certification filter
C. A health inclusion rule
D. A reconciliation rule

173. Which attributes can a de-duplication template use as matching rules? (Choose two.)

A. Created by and updated by
B. Assigned to and comments
C. Location and manufacturer
D. Class and model

174. Why should you test a de-duplication template thoroughly before publishing it?

A. Templates can't be edited after publishing
B. Published templates merge or delete duplicates automatically on a schedule
C. Publishing deletes all open de-duplication tasks
D. Templates only run once

175. A team needs a compliance check that template conditions can't express, so they want to write the criteria in a script. Which role do they need?

A. personalize_dictionary
B. certification_admin
C. data_manager_admin
D. sn_cmdb_admin

176. Where do you create a scripted audit?

A. Data Manager > Policies
B. CI Class Manager > Health > Compliance
C. Compliance > Scripted Audits > Audits
D. CMDB Workspace > Insights

177. An audit has run, but the Compliance scorecard hasn't changed. What do you run next?

A. The CMDB Health Dashboard - Compliance Score Calculation scheduled job
B. The Multisource Dashboard Analytics Population job
C. The Get Well scheduled jobs
D. A Data Manager certification policy

178. Which tool has a person confirm that a CI's owner and support group are actually correct, rather than just filled in?

A. A health inclusion rule
B. A Desired State audit
C. A Data Manager certification policy
D. A de-duplication template

179. Serial Number lives on cmdb_ci, but it should be mandatory only for Hardware and its child classes. What do you use?

A. Make it mandatory on cmdb_ci
B. A reconciliation rule
C. A dictionary override on the Hardware class
D. A health inclusion rule

Insight 48 questions

180. A manufacturer runs Incident Management and wants to add more products. Pair each product with the value it adds to incident work.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Discovery	1. Asset actions and events for managing assets during incidents
B. Service Portfolio Management	2. IT and financial risk data to judge the business impact of incidents
C. Risk Management	3. Detailed operational data on hardware and application CIs to speed resolution
D. Hardware Asset Management	4. Life cycle information about services to tie incidents to service status and history

181. In CMDB Workspace, where do you view and create CMDB 360 saved queries?

A. The Coverage window on the CMDB 360 tab
B. The Saved queries window on the Insights tab
C. The Saved queries window on the CMDB 360 tab
D. CMDB Query Builder

182. An organization is moving data centers and wants to know the effect of planned changes. How does application service mapping help Change Management?

A. It shows which devices go offline first
B. It shows CIs' physical locations
C. It shows the business impact of CIs

183. Where can an admin run Natural Language Queries against the CMDB?

A. CMDB Health Dashboard
B. CMDB Data Manager
C. CI Class Manager
D. CMDB Workspace

184. A hospital's patient management system is down, and the team needs to know which users are affected. Which CSDM-related data helps?

A. Incident history for similar CIs
B. Service offerings by department or location
C. The Affected CIs [task_ci] related list
D. The application service environment attribute

185. Where do you open the Unified Map?

A. CMDB Data Manager
B. CI Class Manager
C. CMDB Workspace

186. A service owner uses the Unified Map and wants to filter out noise. Which filters are available? (Choose two.)

A. Managed by group
B. Business critically
C. Discovery source
D. CI type

187. You need a CMDB query for Seattle databases connected to application services, including their related incidents. What do you do? (Choose two.)

A. Filter the database node on location = Seattle
B. Add property columns to the application service node
C. Add the Incident table to the canvas from the Non-CMDB Tables list
D. Set relationship level to "Up to 2nd level relationships"

188. For CMDB 360 to report on and analyze CI data, what must be true of that data?

A. It went through the IRE
B. It came from an authorized Service Graph Connector
C. It came from ServiceNow Discovery
D. Reconciliation rules with priorities exist

189. How does a well-maintained CMDB help Event Management? (Choose two.)

A. Alerts link to knowledge articles
B. Alerts bind to specific CIs
C. Mapped services show which users consume the service
D. Mapped services show an alert's business impact

190. You want to see how your CMDB setup compares with ServiceNow recommended practice. What do you use?

A. CMDB Data Manager
B. CMDB Health Dashboard
C. CMDB Workspace
D. CMDB Data Foundation Dashboard

191. How does the CMDB help security operations? (Choose two.)

A. It helps the security team assess and remediate incidents
B. It identifies infrastructure affected by a vulnerability
C. It enables audits and attestations across CIs
D. It resolves vulnerabilities automatically

192. With a healthy, service-aware CMDB that has support and change groups filled in, how does Change Management improve? (Choose two.)

A. It schedules and deploys changes with no human review
B. It guarantees no change causes downtime
C. It fills in the assignment group automatically to route changes
D. It shows the potential impact of a change

193. Why would a CMDB manager use CMDB 360?

A. To ingest data through Service Graph Connectors
B. To see CI attribute values reported by each data source
C. To populate the CMDB from many sources
D. To ingest data through import sets

194. What can the Unified Map do that the Service Mapping map can't? (Choose two.)

A. Zoom in and out
B. Show an application and its host
C. Filter nodes by user preference
D. Change how many levels are shown

195. A service desk manager wants to see active incidents and problems for a CI selected in the Unified Map. Which panel shows them?

A. Related items
B. Application services
C. Overview
D. Attributes

196. Which actions can you start from CMDB Workspace? (Choose two.)

A. Create a Data Manager certification policy
B. Run ServiceNow Discovery
C. Create a new CMDB class
D. Remediate duplicate CIs

197. What is a benefit of the CMDB Data Foundations Dashboard?

A. You can resolve certification tasks from it
B. It gives key health metrics to support decisions
C. You can configure health metrics in it

198. You need a CMDB Query Builder query showing every operational CI in one application service. Which steps produce that?

A. 1. Add Business Application, Application Service, and Configuration Item to the canvas. 2. Filter on the service name and operational status. 3. Set the relationships. 4. Run the query.

B. 1. Add Application Service and Configuration Item to the canvas. 2. Set the relationship between them. 3. Add Operational Status and Name as columns. 4. Run the query.

C. 1. Add Application Service and Configuration Item to the canvas. 2. Set the relationship between them. 3. Filter on the service name and the CIs' operational status. 4. Run the query.

---

199. What can CMDB 360 saved queries show you? (Choose two.)

A. Gaps in attribute data across sources

B. Duplicate CIs by source

C. Orphan CIs by source

D. Different attribute values reported by different sources

E. Unique CIs created by each source

---

200. What are benefits of the CMDB Data Foundations Dashboard? (Choose two.)

A. It offers playbooks to help remediate risks

B. It checks that important data is valid and set up correctly

C. It gives you a framework to build custom CMDB metrics

D. It fixes risks automatically

---

201. A service owner wants to see related items, such as active incidents and planned changes, for the home node of the Unified Map. Which area shows them?

A. Contextual side panel

B. Map

C. Content controls

D. Tool box

---

202. Where do you get the CMDB and CSDM Data Foundations Dashboard?

A. It's active by default

B. Free from the ServiceNow Store

C. Free from the ServiceNow Innovation Lab

D. Paid from the ServiceNow Store

---

203. In CMDB Workspace, which tool shows how CIs connect to each other?

A. Business Service Map

B. Unified Map

C. Relationship Map

---

204. In CMDB Query Builder, you want every application service that has a database with incidents, plus all the infrastructure in those services. What do you do? (Choose two.)

A. Add a non-CMDB table (Incident) to the query

B. Build a Service Mapping query that includes non-CMDB tables

C. Build a CMDB query that includes the application services and their infrastructure

D. Build a Service Mapping query to find incidents for the database

---

205. Where do you configure task creation for CIs flagged as orphans?

A. CMDB Health Dashboard

B. CI Class Manager

C. Health Metrics

D. System Properties

---

206. After Multisource CMDB is turned on, which CMDB Workspace tab starts showing data?

A. Management

B. CMDB 360

C. Insights

D. Home

---

207. What is a key benefit of the CMDB Data Foundations Dashboard?

A. It automatically fixes every CMDB data issue

B. It removes the need for CMDB audits

C. It improves visibility into CMDB health and data quality

---

208. Three duplicate records exist for Server A, created January 20, 21, and 22, 2025. After the Correctness calculation job runs, which records get their Duplicate Of field filled in?

A. The January 20 and January 22 records
B. The January 21 and January 22 records
C. The January 20 and January 21 records

209. The Completeness KPI is low. How do you get a KPI that reflects only the CIs you actually manage?

A. You can't; health is calculated for the whole CMDB
B. Write a business rule to filter health results
C. Configure health inclusion rules for the relevant classes

210. How does a CMDB Data Foundations Dashboard playbook help an admin improve CMDB health and align with best practice?

A. It gives a step-by-step approach to improving data quality
B. It replaces the Health Dashboard
C. It integrates third-party discovery tools automatically

211. Where can you create a CMDB 360 saved query?

A. CMDB Query Builder
B. CMDB Workspace
C. CI Class Manager
D. CMDB Reports

212. Every device should have a unique serial number to avoid duplicates. Which tool finds CIs without serial numbers?

A. CI Class Manager
B. CMDB Data Manager
C. CMDB Data Foundations Dashboard
D. CSDM Data Foundations Dashboard

213. You built a complex CMDB Query Builder query about business services, servers, and application services for a network service. Another team wants it in their instance. How do they get it?

A. Move it with an update set
B. Export it as a report and import it as a new report
C. Export it as XML and import it with Import XML
D. Export it from Query Builder and bring it in with Query Builder's import

214. Which metrics are on the Correctness scorecard?

A. Duplicate, Orphan, and Staleness
B. Audit
C. Recommended and Required

215. How does CMDB 360 help? (Choose two.)

A. It lets you define reconciliation rules for how sources change values
B. It stores the last record from each source for review and assignment
C. It provides a complete current assessment of CMDB quality
D. It lets the CMDB recompute after a reconciliation rule change without waiting for new data

216. A service owner set up filters for CI types and layers on the Unified Map. How do they keep those filters?

A. Make a favorite from the banner
B. Set them up again each time
C. Save them as a preset

217. In CMDB Workspace's intelligent search, what can you use instead of Natural Language Query?

A. Create report
B. Assign important actions
C. Use conditional search instead
D. Go to Query Builder

218. Which Health Dashboard metric makes Data Manager retirement policies more effective?

A. Recommended
B. Orphan
C. Audit
D. Staleness

219. With a healthy, service-aware CMDB that has support and change groups filled in, how does Incident Management improve? (Choose two.)

A. Incidents can be routed dynamically
B. You can see an incident's impact on a service
C. Knowledge articles are written automatically
D. Machine learning resolves every incident

220. What does the CMDB 360 dashboard display? (Choose two.)

A. CIs entered by hand
B. The number of discovery sources
C. CIs that no discovery source reports
D. CIs created through API calls

221. You're creating audits for the Compliance scorecard. Which audit type does the Health Dashboard allow?

A. Compliance
B. Desired State
C. Architecture Compliance

222. Every Windows server must have more than 4 GB of RAM. Which Health Dashboard metric do you configure?

A. Recommended
B. Orphan
C. Audit

223. While viewing a CI, how can you tell it's a duplicate?

A. Discovery Status is Duplicate
B. A Duplicate checkbox is checked
C. The Duplicate Of field has a value

224. When a Natural Language Query spans several tables, which button appears with the results?

A. View Report
B. View Data Manager Policy
C. View in Query Builder
D. View Unified Map

225. Which can you query with CMDB 360 in CMDB Workspace? (Choose two.)

A. Unique CIs by source
B. Gaps in attribute data across sources
C. Duplicate CIs by source
D. Different attribute values from different sources

226. In CMDB Query Builder, which tab adds extra columns to the query results for a selected class?

A. Canvas
B. Applied Filters
C. Properties
D. Saved Service Queries

227. You just installed the CMDB and CSDM Data Foundations Dashboards and want data now. What do you do?

A. Run the Health Dashboard jobs
B. Wait 30 days for baseline data
C. Execute the scheduled jobs whose names contain "Get Well"
D. Turn on a system property

CSDM Fundamentals 59 questions

228. Stakeholders listed use cases to decide the CSDM maturity they need. Which use case requires information objects?

A. Event Operations wants events turned into incidents automatically
B. Business Service Management wants to see operational impact on consumers
C. Customer Service wants proactive case management
D. SecOps wants to understand operational risk in the context of business applications
E. Asset Management wants asset life cycle compliance in the context of business applications

229. A manufacturer is aligning CIs to CSDM 5. Which domain do production line monitoring systems belong in?

A. Build and Integration (Build)
B. Foundation
C. Design and Planning (Design)
D. Service Delivery (Manage Technical)
E. Service Consumption (Sell/Consume)

230. An app holding customer PII and PCI data will be deployed to many servers, and audits will check compliance. Where in CSDM should the team store the information used for those audits?

A. Customer service offerings and databases
B. Technical service offerings and Dynamic CI Groups
C. Business applications and information objects

231. A CMDB is aligned to the CSDM Walk stage. What benefit does that bring?

A. Impact assessments on business services for incident, problem, and change
B. Faster APM Foundation rollout for later application rationalization
C. More detailed technical support structures, aligned to OLAs and commitments

232. Moving from legacy status fields to life cycle objects, pair each object or attribute with its description.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. life_cycle_stage	1. Pre-filled with mappings from legacy status values, by table, to the best-fit CSDM life cycle pair
B. life_cycle_mapping	2. A record attribute showing the top-level state of the record's life cycle
C. life_cycle_stage_status	3. A record attribute showing the sub-level state of the record's life cycle
D. life_cycle_object	4. Uses the type of CI (hardware, document, logical, and so on) to decide which sub-level states are available

233. You want to improve CSDM-related data quality. What do you do?

A. Run a ServiceNow Health Scan
B. Use the default CMDB Health Dashboard
C. Use the CSDM Data Foundations Dashboard

234. Under CSDM, a server team wants a catalog item for infrastructure upgrade requests. Which role starts the request and defines the requirements?

A. Application Service Owners
B. Technology Service Owners
C. Enterprise Architecture

235. Why keep monitoring the CMDB Data Foundations Dashboard? (Choose two.)

A. It reports every orphan CI
B. It gives metrics on CIs processed by the IRE
C. It lists every CI that failed a health audit
D. It gives metrics on active CIs updated in the last 90 days

236. Which benefits justify using the CMDB Data Foundations Dashboard? (Choose two.)

A. It automates change approvals
B. It finds and removes duplicate records
C. It gives actionable insight to improve data quality and completeness
D. It tracks CMDB health over time

237. How does Change Management benefit from CSDM? (Choose two.)

A. It routes the change dynamically
B. It shows the change's impact on services
C. It identifies blackout windows
D. It finds the root cause of a change problem

238. What are business values of a CMDB? (Choose two.)

A. Streamlining incident and change management
B. Maintaining CI relationships automatically
C. Collecting and managing financial data
D. Strengthening operational resilience

239. For Vulnerability Response or Security Incident Response, users need enough context to estimate risk and set priorities. Which CSDM Data Foundations Dashboard get-well playbook helps?

A. Percentage of Custom Status Values for CI Life Cycle Stages
B. Locations without a Parent Location
C. Application Services with Business Application Relationships
D. Named Product Models without Product Owners

240. You run the Services Have Owners Identified playbook to fix issues on the Data Foundations Dashboard. Which plays do you use? (Choose two.)

A. Fix Data
B. Report Data
C. Govern Data
D. Analyze Data

241. Moving from legacy CMDB status fields to CSDM life cycle fields, which table can you modify?

A. Life Cycle Stage Status [life_cycle_stage_status]
B. Life Cycle Controls [life_cycle_control]
C. Life Cycle Mapping [life_cycle_mapping]
D. Life Cycle Stages [life_cycle_stage]

242. The Data Foundations Dashboard shows a low Unique Locations result. What does the playbook recommend?

A. Keep both as valid alternates
B. Keep whichever matches your naming standard and delete the other without checking
C. Merge them with the Duplicate CI Remediator
D. Review both locations, update CIs to the correct one, and delete the duplicate

243. Pair each CMDB object with its CSDM domain.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Application Service	1. Design and Planning domain
B. Business Process	2. Foundation domain
C. Business Application	3. Service Delivery domain
D. Business Service	4. Sell / Consume domain

244. What rollout order does CSDM prescribe?

A. Initiate, Plan, Execute, Deliver, Close
B. Foundation, Crawl, Walk, Run, Fly
C. Initial, Developing, Defined, Managed
D. Architecture, Business, Technical, Governance

245. A customer is modeling business applications and wants to capture the kinds of data they hold, including PII. Which class do they use?

A. API Component
B. Data Classification
C. Business Capability
D. Information Object

246. Which CSDM Data Foundations Dashboard tab shows how well product models, locations, and business units follow best practice?

A. Foundation
B. Run
C. Fly
D. Walk
E. Crawl

247. A service portfolio manager wants to see which application services the business service offerings depend on. Which CSDM stage maps that relationship?

A. Fly
B. Walk
C. Run
D. Crawl
E. Foundation

248. An enterprise architect wants to track the company's capabilities. Which CSDM 5 domain covers that?

A. Design and Planning (Design)
B. Service Consumption (Sell/Consume)
C. Service Delivery (Manage Technical)
D. Build and Integration (Build)
E. Foundation

249. Pair each service type with its definition.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Business Service	1. A logical representation of a deployed system or application stack
B. Application Service	2. Published to service owners; underpins one or more business or application services
C. Technology Management Service (Technical Service)	3. Published to business users; underpins one or more business capabilities

250. A platform owner is building CSDM governance teams. Pair each domain with the roles on its team.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Design Domain	1. Service Owner(s), Platform Owner
B. Foundation Domain	2. Technology Service Owner(s), Application Service Owner(s), Platform Owner
C. Portfolio Domain	3. Enterprise Architect(s), Platform Owner
D. Technical Domain	4. Enterprise Architect(s), Data Steward(s), Process Owner(s), Platform Owner

251. Pair each CSDM domain with its benefit.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Foundation	1. See CIs tied to business applications and capabilities to find redundancy, watch costs, and plan investments
B. Service Consumption	2. See technical services and offerings, their support, and their links to technology CIs
C. Design and Planning	3. See business services, their ownership, cost, and scope, and let consumers request them
D. Service Delivery	4. Use base system tables to get the most value from ServiceNow products and the platform

252. A health organization must track regulated patient data and how it relates to business applications. What does CSDM recommend?

A. Have DBAs classify each database, then relate databases back to business applications
B. Add a "contains patient data" field to business applications and have owners fill it in
C. Create an information object for the patient data and relate it to the business application, working with the app owner

253. Which stakeholders should help decide how to populate the CMDB at the CSDM Crawl stage?

A. Application owner and application service owner
B. Business service manager and technology service owner
C. Customer service manager and infrastructure manager

254. A hospital creates a catalog item doctors use to book patients on a new CT scanner. Which CSDM domain does the catalog item map to?

A. Design and Planning (Design)
B. Build and Integration (Build)
C. Foundation
D. Service Consumption (Sell/Consume)
E. Service Delivery (Manage Technical Service)

255. Which CSDM Data Foundations Dashboard tab measures how well technical services and technical service offerings follow best practice?

A. Crawl
B. Walk
C. Fly
D. Run

256. Which tables make up the CSDM Sell/Consume domain?

A. Application Service, Technical Service, Technical Service Offering
B. Business Service Offering and Business Service
C. Business Capability, Information Object, Business Application

257. Application Portfolio Management wants to model platforms as business applications. How would the ServiceNow platform, Incident, and Change be modeled in CSDM?

A. All three are web-based architecture type, with Depends On relationships from the platform to Incident and Change
B. All three are web-based architecture type, and Incident and Change reference the ServiceNow platform
C. The ServiceNow platform is a Platform Host architecture type; Incident and Change are Platform Applications that reference it
D. The ServiceNow platform is a Platform Host; Incident and Change are Platform Applications linked with Depends On relationships

258. Which CSDM domain explains how to relate infrastructure CIs, technical service offerings, and application services?

A. Service Delivery (Manage Technical Services)
B. Build and Integration (Build)
C. Foundation
D. Design and Planning (Design)
E. Service Consumption (Sell / Consume)

259. A retailer wants incidents on customer-facing services resolved fast to protect revenue. Which CSDM attribute helps prioritize them?

A. Business criticality on the service offering
B. Service classification on the technical service
C. Affected CIs on the incident
D. Assignment group on the CI

260. Who should be involved in adding CSDM Design and Planning components to the CMDB? (Choose two.)

A. Application Owner
B. Enterprise Architect
C. Application Service Owner
D. Business Relationship Manager

261. What is the purpose of a CSDM information object?

A. It describes data passed between an API and an application
B. It describes general data for a group of CIs
C. It describes the logical data a business application holds

262. At which CSDM stage do you use technical service offerings to fill in the support group on related CIs?

A. Crawl
B. Foundation
C. Fly
D. Walk
E. Run

263. You're building end-to-end service models and want status and playbooks to improve quality. Where do you look?

A. CSDM Data Foundation Dashboard
B. Service Mapping Data Foundation Dashboard
C. CMDB Workspace
D. CMDB Data Foundation Dashboard

264. For Integrated Risk Management entity scoping, you use the CSDM Data Foundations Dashboard playbooks on the Run tab. Which relationships do they cover? (Choose two.)

A. Locations with parent records
B. Business applications related to information objects
C. Business applications related to application services
D. Logical CIs related to information objects

265. Which CSDM classes are used on the incident form? (Choose two.)

A. Service Offering
B. Application Service
C. Service Portfolio
D. Business Application

266. A business relationship manager wants to set up Service Portfolio Management and present offerings to business consumers. Which CSDM domain fits?

A. Service Delivery
B. Build and Integration (Build)
C. Service Consumption (Sell/Consume)
D. Design and Planning (Design)

267. How does the CMDB support business processes? (Choose two.)

A. It improves decision-making and efficiency across the organization
B. It gives one central view of CIs and their relationships
C. It extends service delivery to every department
D. It lets the CFO or CIO track software licenses

268. After an acquisition, a company wants to unify both companies' deployed applications. Which CSDM record type helps?

A. Incident [incident]
B. Business Application [cmdb_ci_business_app]
C. Application [cmdb_ci_appl]
D. Application Service [cmdb_ci_service_auto]

269. Pair each product with what it achieves using the CMDB.

Match each item on the right to a choice from the list. Some choices may not be used.

Choices	Items
A. Change Management	1. Tracks a CI's life cycle
B. Asset Management	2. Creates incidents automatically based on thresholds
C. Event Management	3. Uses the CMDB to find and fix issues faster
D. Incident Management	4. Shows the impact on downstream CIs when planning an upgrade

270. How would a process owner use business application records to drive business decisions?

A. Application availability
B. Application rationalization
C. Finding impacted applications on a change

271. You set a server's life cycle stage to End of Life while retiring it. Which life cycle stage status values can you pick? (Choose two.)

A. Donated
B. Broken
C. Disposed
D. Charity

272. Groups defined on a technical service offering sync to its CIs. Which ones? (Choose two.)

A. Managed by Group
B. Owned by Group
C. Support Group
D. Approval Group

273. How is an application service created along with its relationships to the infrastructure beneath it?

A. Event Management
B. Service Mapping
C. Agent Client Collector
D. ServiceNow Discovery

274. An IT operations team needs changes to an application environment routed to the right group. Which CSDM data do they use?

A. The Change Group attribute on the CI
B. The approval group on the business service offering
C. The service offering attribute for impacted environments
D. Service subscriptions by user

275. Which of these are good examples of business capabilities to model in CSDM? (Choose two.)

A. Managing finances
B. Providing IT support services
C. Running compliance training
D. Securing personal information

276. Why monitor the CMDB Data Foundations Dashboard regularly? (Choose two.)

A. It gives metrics on active CIs updated in the last 90 days
B. It lists every CI that failed a health audit
C. It gives metrics on CIs the IRE processed
D. It reports all orphan CIs

277. A startup's asset manager follows the change process to move development resources to a new location. Which CSDM domain do those resources belong to?

A. Design and Planning (Design)
B. Service Consumption (Sell/Consume)
C. Build and Integration (Build)
D. Foundation
E. Service Delivery (Manage Technical Service)

278. Which statements about common data in CSDM are true? (Choose two.)

A. Common data is stored in cmdb_ci
B. Common data includes company, department, location, groups, and users
C. Common data elements are CIs
D. Common data elements are not CIs

---

279. How does a CI reference its product model?

A. Through the Model ID field
B. Through a CI relationship
C. Through the Class field
D. Through the Teams related list

---

280. Which life cycle stages belong to the Document life cycle process?

A. Purchase, Inventory, Deploy, Operational
B. Design, Build, Operational
C. Ideation, Operational, End of Life
D. Ideation, Purchase, End of Operation

---

281. What is the default life cycle stage and status for a newly discovered CI?

A. Inventory / Available
B. Ideation / Pilot
C. Operational / In Use
D. Deploy / In Stock

---

282. What should you know before turning on Product Instance 2.0? (Choose two.)

A. It syncs life cycle stage and status across the CI, asset, and install base item
B. It can be turned off later with no impact
C. It removes the life cycle fields
D. It disables the old sync between legacy statuses and life cycle fields

---

283. Which are sibling classes of Application Service under Service Instance in CSDM 5? (Choose two.)

A. Business Service Instance
B. Technical Service Instance
C. Network Service Instance
D. Data Service Instance

---

284. Which related list on a CI tracks its Support, Change, and Managed by groups by type?

A. Group Members
B. Related Items
C. Affected CIs
D. Teams

---

285. Which table stores strategic plans with a mission, vision, and values?

A. sn_gf_goal
B. cmn_value_stream
C. sn_gf_plan
D. sn_align_core_planning_item

---

286. Which CSDM 5 statement about Design & Planning tables is true?

A. They are the main targets of Incident, Problem, and Change
B. They store infrastructure CIs
C. They are used by Enterprise Architecture and are not direct ITSM targets
D. They were added in CSDM 5

Answer key

1. B, D
2. 1-A 2-B 3-C 4-D
3. C
4. A
5. B
6. B, C
7. A
8. D
9. C
10. A, B
11. A
12. B, D
13. B
14. A, B
15. B
16. B
17. B
18. A
19. A, C
20. C
21. A, D
22. D
23. B
24. C
25. A
26. A
27. C
28. C
29. C
30. C
31. B
32. A, C
33. A
34. B
35. C, D
36. C
37. B
38. A
39. 1-C 2-B 3-E 4-A
40. B, C
41. C
42. C, D
43. C
44. B
45. B, C
46. D
47. A
48. A
49. D
50. A, B
51. B
52. A
53. A
54. A
55. C
56. C
57. B
58. A, D
59. C
60. A
61. B
62. A
63. D
64. C, D
65. A
66. 1-B 2-C 3-D 4-E
67. D
68. B, D
69. A
70. C
71. 1-C 2-B 3-A 4-D
72. A
73. C
74. A, C
75. B
76. C
77. D
78. A, D
79. B
80. A
81. C, E
82. B
83. B
84. C
85. A
86. D
87. C
88. A
89. C
90. A, D
91. C
92. B, D
93. A, C, D
94. B
95. B, D
96. A
97. A
98. C
99. B
100. D
101. B
102. 1-A 2-F 3-C 4-B 5-D 6-E
103. A
104. B
105. B
106. C
107. A
108. D
109. D
110. C, D
111. C
112. A
113. A
114. C
115. D
116. B
117. A
118. 1-A 2-C 3-B 4-D
119. 1-A 2-B 3-E 4-D
120. B
121. A
122. A
123. A
124. C
125. B, D
126. B, D
127. C
128. C
129. C
130. D
131. E
132. 1-F 2-D 3-E 4-C 5-B 6-A
133. 1-B 2-D 3-E
134. B, C
135. B, C
136. B, D
137. B
138. A, D
139. A
140. B
141. D
142. C
143. D
144. B
145. A
146. B
147. A
148. A, E
149. A
150. C
151. A
152. A

153. D
154. A, C
155. C
156. A
157. C, D
158. C
159. A, D
160. C
161. A, E
162. B
163. B
164. A, D
165. D
166. C
167. 1-A 2-C 3-B 4-D
168. B
169. B, C
170. D
171. B
172. A
173. C, D
174. B
175. B
176. C
177. A
178. C
179. C
180. 1-D 2-C 3-A 4-B
181. C
182. C
183. D
184. B
185. C
186. C, D
187. A, C
188. A
189. B, D
190. D
191. A, B
192. C, D
193. B
194. C, D
195. A
196. A, D
197. B
198. C
199. A, D
200. A, B
201. A
202. B
203. B
204. A, C
205. C
206. B
207. C
208. B
209. C
210. A
211. B
212. C
213. D
214. A
215. A, D
216. C
217. C
218. D
219. A, B
220. B, C
221. B
222. C
223. C
224. C
225. B, D
226. C
227. C
228. D
229. D
230. C
231. C
232. 1-B 2-A 3-C 4-D
233. C
234. B
235. B, D
236. C, D
237. A, B
238. A, D
239. C
240. A, D
241. C
242. D
243. 1-C 2-B 3-A 4-D
244. B
245. D
246. A
247. C
248. A
249. 1-B 2-C 3-A
250. 1-C 2-D 3-A 4-B
251. 1-C 2-D 3-B 4-A
252. C
253. A
254. D
255. B
256. B
257. C
258. A
259. A
260. A, B
261. C
262. D
263. A
264. B, C
265. A, B
266. C
267. A, B
268. B
269. 1-B 2-C 3-D 4-A
270. B
271. A, C
272. A, C
273. B
274. A
275. A, B
276. A, C
277. D
278. B, D
279. A
280. C
281. C
282. A, D
283. C, D
284. D
285. C
286. C