

MARKETING CLOUD NEXT

Marketing Cloud Next Consultant Certification

THE ULTIMATE STUDY GUIDE

Exam content aligned to the Salesforce Summer '26 Release

What this guide covers

- Part 0 — Exam blueprint & how to use this guide
- Part 1 — Platform overview, setup, and consent management
- Part 2 — Data 360 foundations, CRM integration, and marketing flows
- Part 3 — Salesforce CMS, business units, and content personalization
- Part 4 — Contact point resolution, Agentforce/AI agents, predictive AI, and analytics
- Part 5 — Quick-reference cheat sheet, exam traps, and glossary

Part 0 — How to Use This Guide & Exam Blueprint

Read the guide top to bottom on a first pass, then use Part 5 revision or as a cheat sheet in the days before your exam.

0.1 Exam Format at a Glance

Attribute	Detail
Question format	60 scored multiple-choice questions
Unscored questions	Up to 5 additional unscored questions (randomly integrated, used for Salesforce performance data — do not affect your score)
Time allowed	105 minutes ($\approx$ 1 minute 45 seconds per question)
Passing score	72% — i.e., at least 44 of 60 scored questions correct
Exam type	Pass / fail (no benefit to scoring above the passing threshold)
Prerequisites	None formal, but a solid understanding of Salesforce CRM and Data 360 fundamentals is strongly recommended
Content basis	Salesforce Summer '26 release

0.2 Exam Objective Areas

The exam guide organizes content into six key objective areas, each carrying its own question weighting (the exact percentage breakdown is published in the official exam guide — always check it for the current weighting before you sit the exam). Therefore this document — covers the same material but breaks a few objectives down into more granular sections to make them easier to study:

- Platform overview, setup, and configuration
- Consent management
- Data 360 foundations, CRM integration, and marketing flows
- Salesforce CMS, business units, and content personalization
- Contact point resolution, Agentforce, and AI-powered features
- Analytics, reporting, and performance insights

Part 1 — Platform Overview, Setup & Consent Management

1.1 Agentforce Marketing vs. Marketing Cloud Next

A foundational exam concept is the distinction between the Agentforce Marketing brand and the Marketing Cloud Next product. Agentforce Marketing is not a single product — it is the broader brand for Salesforce's next-generation marketing portfolio. Marketing Cloud Next is just one product within that portfolio.



Figure 1.1 — The Agentforce Marketing portfolio and where Marketing Cloud Next fits within it.

The four products in the Agentforce Marketing portfolio

- **Marketing Cloud Next** — the next-generation marketing platform; the evolution of Marketing Cloud Engagement.
- **Salesforce Personalization** — the next-generation real-time personalization platform; the evolution of Marketing Cloud Personalization.
- **Marketing Intelligence** — an AI-powered analytics solution; the evolution of Marketing Cloud Intelligence.
- **Loyalty Management** — enables B2B and B2C loyalty and rewards programs through a no-code interface.

Exam tip

- If a question describes the overall brand/portfolio, the answer is Agentforce Marketing.
- If a question describes the specific next-gen marketing platform product, the answer is Marketing Cloud Next.

1.2 Growth Edition vs. Advanced Edition

Marketing Cloud Next currently ships in two editions. Advanced Edition is a superset — it includes everything in Growth Edition plus additional capabilities, most notably around predictive AI and account-level scoring.




Figure 1.2 — Growth vs. Advanced Edition capability split.

Capability	Growth Edition	Advanced Edition
Engagement scoring — individual records	Available	Available
Engagement scoring — account records	Not available	Available
Engagement Frequency (predictive AI)	Not available	Available
Send Time Optimization	Available	Available

Knowledge-check takeaway

- Scoring for individual records is available in both editions.
- Scoring for account records is Advanced Edition only.
- The predictive AI capabilities of Engagement Scoring and Engagement Frequency are exclusive to Advanced Edition.
- This is not an exhaustive feature list — the platform evolves every release.

1.3 Platform Setup & Configuration

Setting up Marketing Cloud Next is best thought of as configuration rather than implementation — the platform and Data 360 can both be enabled in a handful of clicks through a guided setup interface. There are six key configuration steps at minimum, plus optional steps that are not required for core functionality such as creating and automating marketing messages.

1.3.1 Required permission sets

Before configuration begins, a System Administrator profile user must assign two permission sets to whoever will configure Marketing Cloud Next:

- **Data Cloud Architect** — grants access to Data 360 setup and data modeling objects; allows installing data kits, deploying data streams, configuring identity resolution, and other admin tasks. (Renamed from Data Cloud Admin in the Spring '26 release.)
- **Marketing Admin** — allows configuration of most marketing settings in Salesforce Setup, plus publishing and activating campaigns and segments.

Exam trap

- "System Administrator" is a PROFILE, not a permission set — do not select it if a question asks which permission set is required.
- The correct Data 360 permission set as of Spring '26 is Data Cloud Architect (previously Data Cloud Admin).

1.3.2 The Setup Assistant — three guided tasks

Once permission sets are assigned, the Marketing Cloud Next Setup Assistant walks the admin through three button-click tasks, followed by two additional required manual tasks:




Figure 1.3 — End-to-end Marketing Cloud Next setup sequence.

1. Set up Data 360
2. Enable Marketing Cloud
3. Deploy the required data streams
4. Add your company's physical address (Setup → Company Information) — required for the footer of promotional emails to support regulatory compliance
5. Create an authenticated domain

1.3.3 Authenticated domains vs. authorized email domains

In Marketing Cloud Engagement, a Sender Authentication Package configures the sending domain. In Marketing Cloud Next, you instead add an authenticated domain, generate the required DNS records, and publish them through your DNS provider.



Figure 1.4 — Steps to create an authenticated domain.

From Setup, search "Authenticated Domains" in Quick Find → click Add Domain → Continue → define the subdomain and a default From email username → download the DNS zone file (or view individual records) → publish those DNS changes with your DNS provider → check the confirmation box → click Activate My Domain.

Authenticated domain vs. authorized email domain — know the difference

- An authenticated domain establishes complete DKIM/DNS sending configuration and is required for DMARC alignment; the From address domain must align with it.
- An authorized email domain only verifies ownership via a simple verification record — it does NOT provide full DKIM/DNS sending configuration.
- An authorized email domain is required if you want to use dynamic From or Reply-To addresses.
- Using a personal address on the root domain as the From address can cause DMARC alignment failures.

1.4 Consent Management

Consent is one of the most heavily tested areas of the exam. Marketing Cloud Next supports four levels (granularities) of consent for promotional communications:




Figure 1.5 — The four levels of consent granularity, from broadest (Level 1) to most granular (Level 4).

1. Level 1 — the entire individual record, regardless of channel, subscription type, or contact point value (a person-level permission/restriction covering ALL marketing communications).
2. Level 2 — an entire communication channel (e.g., email or SMS), regardless of which specific contact point is used within that channel.
3. Level 3 — a specific contact point value (e.g., one email address or phone number), regardless of subscription type.
4. Level 4 — a specific contact point AND a specific subscription type. This is the level Marketing Cloud Next enforces — the most granular option, giving precise control over what an individual receives.

Exam tip

- Marketing Cloud Next enforces LEVEL 4 consent: contact point + subscription type.

1.4.1 Where consent is stored — the four-object data model

Four Data Model Objects (DMOs) work together to manage Level 4 consent:




Figure 1.6 — Entity relationship diagram of the four objects that manage consent.

- **Communication Subscription (DMO)** — defines the TYPE of communication offered that an individual can opt into or out of (the purpose/category — e.g., "Marketing Events" or "Monthly Newsletter").
- **Communication Subscription Channel Type (DMO)** — defines the channel used to deliver a given subscription (e.g., the Email channel for a Marketing subscription).
- **Communication Subscription Consent (DMO)** — stores the individual's consent status for a subscription, at a specific contact point value, for a specific channel type (e.g., someone@example.com opted in to Marketing via Email).
- **Engagement Channel Type** — defines the engagement channel itself: Email, SMS, WhatsApp, or RCS.

1.4.2 Creating and updating consent




Figure 1.7 — Supported vs. unsupported methods for creating/updating consent.

Supported methods

- **File import** — upload a CSV of contact point values via the UI. Limitations: manual/user-driven; max 50,000 rows per file; each upload covers only ONE channel and ONE subscription; each upload can assign only ONE consent status (opt-in or opt-out).
- **Create Consent flow element** — available in Automation Event-Triggered Flows, On-Demand Flows, and Data Cloud-Triggered Flows. In Automation Event-Triggered and On-Demand Flows it is labeled "Consent Request."
- **Privacy Consent Status LWC** — a Lightning Web Component usable on Lead, Contact, and Person Account records to trigger the consent service.
- **Preference pages** — customer self-service consent management (see 1.4.4 below).

Unsupported methods (per Salesforce Help)

- Batch data transform
- Direct mapping from a data stream
- Bulk ingestion using the Bulk Ingestion API

- Messaging Consent v2 flow action

With custom configuration, these unsupported methods can technically materialize records from a Data Lake Object (DLO) into the Communication Subscription Consent DMO — but records ingested via data streams are NOT available in the DMO in real time (approximately a 15-minute delay to materialize). More importantly, using them can cause consent to be silently disrespected — see the cache scenario below.

1.4.3 The consent cache layer and read-after-write consistency

To provide read-after-write consistency, the consent service uses a caching layer. Each cached entry has a 90-day TTL (time to live), after which it expires automatically.




Figure 1.8 — How the consent cache, DLO, and DMO interact when consent is created through the supported service.

- When a consent record is created through the consent service, the cache is updated FIRST, making the latest status immediately available.
- The record is then written to the underlying DLO; when queried later, the Communication Subscription Consent DMO is hydrated with the record.
- At send time, the send engine checks the cache first. If there IS a cached value, it uses that value and does NOT check the DMO. If there is no cached value, it checks the DMO, and then updates the cache.
- There is currently no process that automatically refreshes an existing cached record when a status update arrives via an unsupported data-stream mapping.

Why this matters — the stale-cache scenario

- 1) A new opt-in consent record is created for someone@domain.com and materialized into the DMO via an UNSUPPORTED method (e.g., direct data stream mapping).
- 2) An email is sent to that address. Since there was no prior cache entry, the cache layer is populated with the current (opt-in) status.
- 3) An opt-out record is later streamed to the DMO for that same consent record via an unsupported method — the cache is NOT automatically updated.
- 4) The next email send checks the cache, sees opt-in, and SENDS the email — even though the individual is actually opted out.
- Lesson: only use SUPPORTED consent methods (Create Consent flow element, file import, preference pages, privacy consent LWC) to avoid this silent failure.

When the Create Consent flow element is used inside a Data Cloud-Triggered Flow (evaluating a condition on, e.g., the Contact Point Email DMO), the flow writes consent values immediately to BOTH the cache layer and the underlying DLO the moment the interview reaches that action — so subsequent reads return the correct, up-to-date status right away, without waiting on the DMO layer to refresh. The same consent element also exists in Automation Event-Triggered Flows and On-Demand Flows, just under the "Consent Request" label.

1.4.4 Preference pages

Custom preference pages were introduced in Spring '26; Summer '26 expanded the capability to support multiple preference pages (e.g., one per brand). When building a preference page, marketers can:

- Apply a brand to control colors, typography, buttons, spacing, and borders.
- Add content blocks: dividers, headings, lists, paragraphs, layouts, and image blocks.
- Designate a default preference page for use across all marketing message sends.

Preference page considerations (frequently tested)

- The message channel automatically determines which preference page type is shown — email messages link to the Email Preference Manager; SMS messages use a separate SMS-specific preference page.
- A single unified page covering BOTH email and SMS is not available out of the box; a true cross-channel experience requires custom development.
- Multilingual preference pages are NOT natively supported — this requires custom forms or custom development.
- You can publish multiple preference pages (e.g., one per brand); published pages appear in the merge field selector so you can choose the right one per send.
- The standard subscription block does not let you edit button labels/headings/subheadings or reorder the subscription list; as a partial workaround you can rename communication subscriptions from the Consent tab (full text customization is on the roadmap).
- Custom code — AMPscript, Apex, or Server-Side JavaScript — is NOT supported on preference pages.
- Pre-populating a page from URL parameters is not currently possible.

Building a preference page — walkthrough

1. From the Consent tab, select Subscriptions → New Subscription. Name it and assign it to a channel (and optionally scope it to specific business units).
2. Open a CMS workspace and create a new page, selecting the Subscription List block.
3. Choose the channel and define the business unit scope for the block.
4. Customize the page: add a header image, adjust container padding, add headings/paragraphs, and other components.
5. Save and publish. The page can then be added as a preference page link inside an email.

Knowledge-check takeaways

- CommunicationSubscriptionConsent (a CRM/COM object) is NOT the same object as CommunicationSubscriptionConsentDMO — a common distractor.
- The HasOptedOutOfEmail field on Lead/Contact does NOT update consent in Marketing Cloud Next.

Part 2 — Data 360 Foundations, CRM Integration & Marketing Flows

2.1 Data 360 Foundations for Marketing Cloud Next

Marketing Cloud Next is a Lightning app built directly on top of Data 360, so the exam expects a working understanding of core Data 360 concepts — though not deep unstructured-data expertise. Salesforce explicitly recommends candidates consider earning the Data 360 Consultant certification before attempting the MC-Next exam.




Figure 2.1 — The Data 360 object model: from source system to activation.

- **Data Lake Object (DLO)** — raw, ingested data as it lands from a source system via a data stream.
- **Data Model Object (DMO)** — structured, mapped data (e.g., Individual, Contact Point) produced by mapping one or more DLOs into the canonical data model.

You should understand how data streams are created and how data is mapped from DLOs into DMOs, plus be broadly familiar with these supporting Data 360 capabilities:




Figure 2.2 — Supporting Data 360 capabilities referenced throughout the exam objectives.

- **Data spaces** — logical partitions of data (often aligned to business units).
- **Data kits** — pre-packaged data streams/mappings for common source systems (e.g., the Sales Data Kit).
- **Calculated insights (CI)** — derived/aggregated metrics computed from underlying data.
- **Data graphs** — pre-assembled, read-only views spanning multiple related DMOs.
- **Segmentation** — defining audience criteria against DMOs/data graphs.
- **Activations** — publishing a segment to a destination/target for use in engagement.
- **Identity resolution** — matching and unifying individual records into a single unified profile.

For deeper study, Salesforce recommends the Data 360 learning journey on Trailhead — a progressive, multi-part path built to develop this knowledge step by step.

2.2 Salesforce CRM Integration

Historically, any data used for audience segmentation, decisioning, or personalization needed to live in Data 360 first, because Marketing Cloud Next is built on that platform. Starting Spring '26, Salesforce began reducing this dependency, and Summer '26 extended it further.




Figure 2.3 — How Marketing Cloud Next now integrates directly with Salesforce CRM.

Spring '26 — CRM records as a direct audience source

- CRM records can now be used directly as an audience source — on a scheduled basis or triggered by an event (change data capture) — without first landing in Data 360.

Summer '26 — deeper integration

- **Campaign members** and **actionable lists** can now be used as audiences.
- **Sales Data Kit** — pre-configured data streams and mappings for common CRM objects: leads, contacts, accounts, prospects, and related objects.
- **Flow support** — flows can perform standard CRUD (create, read, update, delete) operations directly on Salesforce records.
- **Reporting** — standard Salesforce CRM reports can be used alongside Data 360 data, including reports built on DMOs, plus analytics experiences like the Unified Engagement History dashboard (covered in Part 4).

2.2.1 Actionable lists

An actionable list is a fixed collection of audience members — useful when you already know exactly who to engage (e.g., trade-show attendees, hand-picked leads/contacts, or a curated audience ready for activation).




Figure 2.4 — Building an actionable list from a CSV import.

Important consent caveat

- If an actionable list is created containing NEW leads or contacts, Marketing Cloud Next does NOT automatically create consent records for them.
- Consent records must be created or loaded separately using a supported method.

Demo walkthrough — creating an actionable list from a CSV import

1. From Leads/Contacts, click Import → Import from File.
2. Choose the CSV file and enable "Add to actionable list."
3. Select an existing actionable list, or create a new one.
4. Review field mappings and click Start Import.
5. Once complete, the new leads appear under the Actionable Lists tab, ready for activation and engagement in Marketing Cloud Next.

2.3 Marketing Flows — The Orchestration Engine

In Marketing Cloud Next, ALL outbound messaging is orchestrated through Flow — a significant shift from Marketing Cloud Engagement, where sends could originate from many different tools and execution contexts (Journeys, Automations, triggered sends, UI-based sends, Salesforce sends, Content Builder sends, and more).

These are called marketing flows, or high-scale flows, and they are fundamentally different from the "core" flows Salesforce admins typically know. Marketing flows run on a separate, high-scale processing engine, executing off-core — i.e., outside the core Salesforce transactional runtime. Because of this, many traditional concerns (governor limits, transaction management, scaling patterns) simply do not apply in the same way.




Figure 2.5 — The six marketing / high-scale flow types in Marketing Cloud Next.

2.3.1 Audience Flows (new in Summer '26)

Audience flows unify several previously-separate flow/audience models into a single orchestration experience — Segments, Actionable Lists, CRM records, and Campaign Members. Rather than choosing a different flow type per audience model, marketers choose the audience source directly within one Audience Flow's configuration. This includes support for list-triggered flows.

Audience flows will feel familiar to anyone who has used Journey Builder's contact entry mode in Marketing Cloud Engagement. When configuring a recurring schedule, you set both the execution schedule and the re-entry conditions, mirroring Journey Builder's three entry options:

- No re-entry
- Re-entry at any time
- Re-entry only after exiting

2.3.2 Activation-Triggered Flows

Introduced in Winter '26. These flows use a Data 360 Activation as the audience/target. Conceptually similar to segment-based flows, but with an important difference: when configuring an activation, you can add filters to activation membership, add contact point filters, limit audience size by attribute, and append additional related attributes from objects associated with a segment member. This means email merge fields and repeaters can reference enriched activation attributes directly, without requiring additional lookups at send time.

2.3.3 Automation Event-Triggered Flows

Also called simply "event-triggered flows." Unlike segment-triggered (scheduled/batch) flows, these launch automatically when a predefined event occurs. Event sources include:

- Standard engagement events — email opens, clicks, form submissions (no extra configuration needed).
- Engagement signals — custom-defined engagement events based on customer behavior (e.g., website interactions, PDF downloads).
- CRM record-triggered events (introduced Spring '26) — fire when a prospect, lead, contact, or related record is created or updated.

2.3.4 On-Demand Flows

On-Demand Flows let external systems trigger transactional messaging instantly via the REST API. Real-time personalization data can be passed directly in the API payload using an Apex-defined data schema — no requirement to first ingest/hydrate data into Data Cloud. Typical latency is ~1–3 seconds, making these ideal for use cases like one-time passcodes, order confirmations, password resets, and real-time service alerts.

2.3.5 Broadcast Flows

Designed for fan-out scenarios where a single event must trigger communication to a large group simultaneously (e.g., "a power outage occurred in sector X — SMS everyone in that area"). Also triggered via REST API, but the key difference from On-Demand Flows is the target: on-demand targets a single individual, broadcast targets ALL members of a dynamic segment.

Broadcast flows require a dynamic segment — one evaluated at execution time using variables passed into the flow via the API request, enabling contextual audience resolution. You can combine static and parametrized criteria within the same segment. Common use cases: flight delay notifications, utility outage alerts, emergency notifications. Broadcast flows can run synchronously or asynchronously.

2.3.6 Data Cloud-Triggered Flows

Execute whenever a record within a Data Cloud DMO (or Calculated Insight Object) is created or updated. Under the hood, Data 360 monitors the specified DMO for changes; on a match, a data action publishes a platform event that triggers the flow. Technically NOT a marketing flow, but still considered a high-scale flow type — most commonly used to automate consent management via the Create Consent flow action (see 1.4.3).

Flow-type exam trap

- Record-triggered flows are core CRM flows and CANNOT be used to send emails from Marketing Cloud Next.
- Data Cloud-triggered flows also CANNOT be used to send emails from Marketing Cloud Next — they're for DMO-driven automation (like consent), not outbound messaging.
- Only marketing/high-scale flow types (Audience, Activation-Triggered, Automation Event-Triggered, On-Demand, Broadcast) can send messages.

2.3.7 Flow type comparison

Flow Type	Trigger	Typical Audience	Example Use Case
Audience Flow	Schedule / list / event	Segment, actionable list, CRM records, campaign members	Recurring newsletter, curated trade-show follow-up
Activation-Triggered	Segment refresh / publish	Data 360 Activation	Back-in-stock alert with enriched product data
Automation Event-Triggered	Engagement / signal / CRM event	Individual reacting to event	CSAT survey when a case closes
On-Demand	REST API (single call)	One individual	One-time passcode, order confirmation
Broadcast	REST API (single call)	All dynamic segment members	Flight delay notice, outage alert
Data Cloud-Triggered	DMO/CI record change	N/A (not a send)	Automate consent on a DMO change

2.3.8 Flow walkthroughs

Audience flow

Selected an Actionable List as the audience source → configured a recurring schedule with execution schedule + re-entry conditions → added a Send Email Message element from the flow canvas, configured via the properties panel.

Activation-triggered flow (back-in-stock alert)

Two data streams from an e-commerce platform: one imports the product catalog into the Goods Product DMO; another ingests back-in-stock registrations into a custom DMO related to Individual. An activation using the back-in-stock registration segment (Data 360 DMO target) adds related attributes from the Goods Product DMO for dynamic personalization. In Email Builder, an Activation Data Provider surfaces those attributes as merge fields. When the segment refreshes on schedule, the activation publishes automatically and qualifying members enter the flow per the configured re-entry rules — each receives a personalized email.

Automation event-triggered flow (CSAT email)

Configured a "prospect/lead/contact/related record change" action to start the flow whenever a Case record's status changes to Closed. The email uses content variables (see Part 3) to personalize with values from the flow interview — sent in near real-time.

Data Cloud-triggered flow (CSAT score capture)

Watches the Email Engagement DMO for a new record where the channel action equals "click" and the link matches the landing-page URL. Two formulas extract the case ID and score from URL parameters; the flow then updates the Case record with the CSAT score.

On-demand flow (transactional OTP email)

Business scenario: a bank sends a one-time-password email whenever a customer logs in. An Apex wrapper class defines the variables passed in the API payload; a transactional email uses an Apex Class Data Provider to personalize content from those properties. The On-Demand Flow contains an Apex-defined variable (marked "Available for input") plus a Send Email Message element referencing the transactional email. Triggered via Workbench's REST Explorer by appending the flow's API name to the endpoint route.

On-demand flow API notes

- Both the email address and the individual ID are required in the request.
- The individual ID does NOT need to correspond to an actual individual record in Data 360 — it simply cannot be null.
- The request will NOT create a new record in the Individual DMO, though the resulting engagement activity still appears in the Email Engagement DMO.

Broadcast flow (flight itinerary change notification)

Scenario: an airline sends a personalized notification to all affected passengers whenever a flight itinerary changes. A "passenger manifest" dynamic segment has parametrized values enabled on its rules, so values can be passed in dynamically via the API request (static and parametrized criteria can combine in the same segment). The email again uses an Apex Class Data Provider. Flow variables (marked "Available for input") map to the parametrized segment values. Triggered via the same Workbench REST Explorer pattern, appending the broadcast flow's API name — when executed: parameter values pass to the segment → the segment refreshes and evaluates all matching individuals (static + dynamic criteria) → every qualifying member is admitted into the broadcast flow → each passenger receives a personalized notification.

Knowledge-check takeaway

- To send an email from Marketing Cloud Next triggered by a single REST API call to a whole dynamic segment of recipients, use a BROADCAST FLOW (not record-triggered or data-cloud-triggered flows).

Part 3 — Salesforce CMS & Content Personalization

3.1 Business Units & Content Architecture

Business units in Marketing Cloud Next serve the same purpose as in Marketing Cloud Engagement and Marketing Cloud Account Engagement — separating content and data across teams, brands, or regions. All marketing content (images, landing pages, forms, brand assets, emails) lives in CMS content workspaces, and each workspace is associated with a specific business unit. This structure works alongside Data 360's data spaces.




Figure 3.1 — Relationship between data spaces, business units, and CMS content workspaces.

When a marketer creates an asset (e.g., an email), the system shows only content available to the associated business unit. The same filtering applies when selecting content within a campaign or flow — users see only content tied to that record's business unit or linked data space.

- A data space can be associated with either NO business units or exactly ONE business unit.
- A business unit can be associated with ONE OR MORE content workspaces.

Cross-business-unit content rule (frequently tested)

- By design, content CANNOT be accessed directly across business units.
- Each CMS content workspace must be associated with a specific business unit in Setup.
- EXCEPTION: content can be posted as a COMMON ASSET, making it available to all business units. Users in other business units can then copy that asset into their own workspace to use it locally.

3.1.1 Business unit roles




Figure 3.2 — Business unit access roles.

- **Marketing Standard** — access to both the data space AND CMS content.
- **Marketing Read-Only** — access to the data space only (not CMS content).

- Users without the Marketing Manager or Marketing Admin permission set can only be assigned the Marketing Read-Only role.
- Common assets can be enabled from the business unit settings page.

Demo walkthrough — configuring a business unit

1. From the Business Units page in Setup, open an existing business unit (or create a new one).
2. Associate CMS workspaces, assign authenticated email domains, configure channel-specific settings and preference pages, and enable specific AI features.
3. Under Business Unit Members, click Add Users, choose users, and assign them a role.
4. Enable common assets from the business unit settings page if needed.

3.2 Content Workspaces & Content Types

Content created in Marketing Cloud Next is stored in a marketing workspace — either the default Marketing Cloud content workspace or additional workspaces you create to organize assets by campaign, initiative, brand, or team, and to support cross-team collaboration.

- Content can be shared between marketing workspaces, or with a general workspace used by other (non-marketing) users — e.g., managing brand assets centrally and sharing them out.
- Content within a workspace can be organized into folders.
- Content assets can be exported from one workspace and imported into another Salesforce org (e.g., moving content between production and sandbox).
- Access can be managed both at the business unit level AND at the individual workspace level.

Common content types to know

You do not need to know every content type Marketing Cloud Next supports, but you should be comfortable with the most common ones: images, landing pages, forms, and emails.

Demo walkthrough — building an email, a form, and a landing page

1. Email: open an existing email, select Edit, drag components (e.g., an image component) onto the canvas from the components panel, adjust properties (size, spacing, padding, color), save, then publish before it can be used.
2. Form: create a new form component → add a data source (choose the CRM object where submissions should create records, e.g., "Leads" → Lead object) → name the form and add fields (First Name, Last Name, Email) → adjust formatting → save.
3. Form handler flow: select New Flow — Flow Builder auto-adds a Create Records element with field mappings between the form and the target record. Publishing the form also activates its associated flow.

1. Landing page: create a page, add the published form plus other elements (e.g., a header image), configure the page title and SEO title, then save and publish. When a visitor submits the form, the flow runs and creates a new Lead record in CRM.

CloudCove.ai

3.3 Content Personalization — Methods & Data Sources

For the exam, you only need to know EMAIL content personalization — not other channels. In Marketing Cloud Engagement, personalization relies almost entirely on a single data source type: the data extension (sendable, with AMPscript retrieval for non-sendable, plus a few exceptions). Marketing Cloud Next offers a much broader set of options: data providers, content variables, and marketing objects.



Figure 3.3 — The personalization spectrum, from simple point-and-click to full scripting.

- **Merge fields** — insert a single value from a selected data source; the simplest option for field-level personalization.
- **Repeaters** — display a collection of related items (e.g., recent purchases, recommended products) by designing one item and connecting it to a data source.
- **Dynamic content** — swaps an entire component or field based on targeting rules (e.g., different images/offers by location, interest, or loyalty status). If a customer qualifies for multiple variations, you set the priority order.
- **Handlebars** — a templating language for advanced use: accessing nested data, applying conditional logic, formatting values, iterating over collections.
- **AMPscript** — used for complex personalization logic and data operations; NOT every AMPscript function available in Marketing Cloud Engagement is supported in Marketing Cloud Next.

In short: merge fields, repeaters, and dynamic content are marketer-friendly, point-and-click options; Handlebars and AMPscript give developers more power and flexibility.

3.4 Data Providers (know all eight for the exam)




Figure 3.4 — The eight data providers available for email content personalization.

- **Data graph** — a pre-assembled, read-only view spanning multiple related DMOs, typically starting from Individual/Unified Individual, connecting to contact points, accounts, orders, products, and engagement activity. Lets Marketing Cloud Next retrieve customer data quickly without joining underlying objects at render time.
- **Event data provider** — surfaces data from an event-triggered automation flow. Each event has a schema matching the underlying engagement DMO that triggered the flow (form submissions, email sends/opens, engagement signals). Supports direct attributes from the event DMO, but NOT related collections (e.g., an order event exposes Sales Order DMO fields, but not related Sales Order Product records).
- **Activation data provider** — makes an activation's direct AND related attributes available for personalization; supports segments based on Individual or Unified Individual. IMPORTANT: a single

message cannot combine an activation data provider with an event data provider — each requires its own flow orchestration pattern, and combining them would conflict.

- **Salesforce record data provider** — personalize with live data from a Salesforce CRM object (e.g., a Case, Lead, or another supported object) via merge fields.
- **Apex class data provider** — defines a custom schema (fields + data types) for personalization; ideal for On-Demand or Broadcast flows where data is passed in via the API payload, though usable with other flow types too.
- **Personalization recommender** — brings Salesforce Personalization recommendations (relevant products/services/content per recipient) into an email; typically displayed via a repeater.
- **Lookup graph data provider** — accesses non-profile data with no direct relationship back to a unified individual, by taking a value from the recipient's primary data graph as a key into a second, smaller data graph (e.g., using a product ID to retrieve catalog details).
- **Offer data provider** — used with Salesforce Loyalty Management to deliver tailored promotions (offer name, description, coupon code) via merge fields; combine with dynamic content rules to show different offers based on attributes like loyalty tier.

Data provider exam trap

- The SAME message CANNOT use both an Activation data provider and an Event data provider at the same time — they require conflicting flow orchestration patterns.

3.5 Content Variables

Content variables are custom placeholders you define within an email and populate at runtime from Salesforce Flow. Unlike a data graph, they do not retrieve data from predefined customer data structures — instead, the email defines what information it expects (e.g., order number, appointment date, discount amount, support rep's name), and the flow supplies the corresponding values at send time.




Figure 3.5 — Content variables: the content defines the placeholder, the flow supplies the value.

- A content variable can be mapped to ANY data source available to Salesforce Flow — Salesforce data, MuleSoft integrations, even HTTP connectors.
- Especially useful when the required information is calculated during the flow, or comes from a system outside Data 360.
- Supported value types: boolean, text, date, date/time, number, and a record ID that references a Salesforce record.

Example: an email defines variables for order number, delivery date, and a tracking URL. The flow retrieves those values from an external order-management system, maps them to the variables, and sends the personalized email.

3.6 Marketing Objects

A marketing object is a flexible, marketer-managed data store for marketing use cases in Marketing Cloud Next — the closest Marketing Cloud Engagement analogue is a data extension. Useful for storing product catalogs, reference tables, rewards balances, promotional codes, and similar reference data.




Figure 3.6 — Marketing objects: marketer-managed, CSV-driven reference data.

- Created by importing a CSV, defining its fields, and identifying a primary key — no admin needed to create a Salesforce object or modify the Data 360 data model.
- Data is queried at send time using Handlebars or AMPscript.
- Supported field types: text, number, and decimal.
- Data can be fully refreshed by uploading a new CSV, which replaces existing records.

Example: an email uses the recipient's address (from the default data graph) to find their matching record in a marketing object, then retrieves their current rewards balance.

Demo walkthrough — creating a marketing object

1. From the Data Management tab in the Marketing app, select New.
2. Upload a CSV containing the reference data (e.g., product data) and click Next.
3. Name the marketing object; review/adjust detected fields and data types.
4. Select the field that will serve as the primary key. Done — the object is created and its records can be viewed.

Marketing object exam takeaway

- Marketing objects are useful for storing RELATED REFERENCE DATA (product catalogs, reward balances, promo codes) that a marketer manages directly via CSV — not for storing subscriber/contact records themselves.

3.7 Personalization in Action —Recap

Revisiting the back-in-stock alert email : the Back-in-Stock Alerts activation was already added as a data provider. From the email component, the Add Merge Field icon opens a menu listing the activation data provider alongside other options (e.g., global strings). Selecting "First Name" inserts the merge field, with an optional fallback value (e.g., "friend") for when the name is missing.

The same email demonstrated two Handlebars examples: retrieving the customer's first name from the activation and converting it to uppercase, and retrieving the price field from the related Goods DMO and formatting it as currency — illustrating the two levels of personalization: merge fields for simple, marketer-friendly insertion, and Handlebars/AMPscript for more controlled retrieval and formatting.

Dynamic content variations (emails and landing pages) require a data graph configured as a data source. Example: selecting a header image → Dynamic Content section → New Variation → named "US West States" → targeting rule set on the State/Province field from the individual's related Contact Point Address data, using the "Is In" operator with specific state values → a different header image assigned to the variation. If a recipient qualifies for more than one variation, a priority order determines which one is shown.

knowledge-check recap

- Content cannot be accessed directly across business units — but CAN be posted as a common asset to make it available to all business units.
- Marketing objects are best used for storing related reference data (not for subscriber-level personal data extensions in the MCE sense).

Part 4 — Contact Point Resolution, Agentforce/AI & Analytics

4.1 Contact Point Resolution at Send Time

Unlike a subscriber in Marketing Cloud Engagement (who can only have one email address), an Individual in Data 360 can have MULTIPLE email addresses, and multiple contact point values across other channels like SMS.




Figure 4.1 — One unified individual can have multiple contact points across source records.

Reconciliation rules vs. contact points — a key distinction

- Reconciliation rules determine which SINGLE value is chosen for a unified field that cannot hold multiple values (e.g., an individual's name).
- Reconciliation rules do NOT apply to contact points (emails, phone numbers) — ALL contact points are retained as part of the unified profile and remain available for segmentation and activation.
- When an email is sent to a unified individual, a SEPARATE message is sent to each UNIQUE email address. If the same address appears on multiple source records, it is treated as a single contact point — only one email is sent to it.

Because multiple contact points can exist, Marketing Cloud Next needs a way to decide which one to use at send time. You can use an activation or an activation template to control this.

4.1.1 How the contact point is chosen




Figure 4.2 — Decision logic for resolving the contact point used at send time.

- If a data graph was selected in the flow's automation properties when the flow was created, that data graph is used for contact point selection — and it takes precedence over other methods. The data graph must include a specific required structure and field (see the relevant Salesforce Help article) for messages to send successfully.
- On-Demand Flow: the contact point is typically derived from the value supplied in the API payload — UNLESS a data graph is configured in the flow's automation properties, in which case the data graph takes precedence.
- Activation-Triggered Flow: configure the activation's source priority order to determine which contact point is used when multiple sources are available.
- Audience Flow using a segment: create an activation template to define the source priority order and filter rules for segment members/contact points. Select that template when configuring the Send Message element. An activation template is REQUIRED when the segment's "segment on" object is not Unified Individual.

Demo walkthrough — creating an activation template

1. From the Activations tab in Data 360, click New → select Template as the activation type → Continue.
2. Select a data space and the DMO the segment criteria is based on (typically Individual or Unified Individual).
3. Choose Data Cloud as the platform → Continue.

4. Expand the Email channel → Edit → configure the source priority order (e.g., add Salesforce Contact, then Salesforce Lead, then "Any" as the catch-all for other sources) → Save.
5. Optionally add contact point filters (criteria applied to contact point DMOs) and activation membership filters (which segment members are included).
6. Name and save the template. In the audience flow's audience properties, select this activation template.

With the example priority order above, Marketing Cloud Next prioritizes an email address associated with a Contact record; if none exists, it looks for a Lead record's address, followed by an address from any other source.

Exam trap — phone-based activation templates

- For SMS/voice channel priority configuration, remember to select the telephone number field from the CONTACT POINT PHONE DMO (not a generic phone field on Contact/Lead) when configuring the source priority order.

4.2 Agentforce & AI Agents

The exam guide requires you to identify the different agents and predictive AI features and know when to use each. Marketing Cloud Next ships with several out-of-the-box AI agents:



Figure 4.3 — The five out-of-the-box AI agents in Marketing Cloud Next.

- **Segment Creation Agent** — uses natural-language prompts to create Data 360 audience segments, translating an audience description into editable segmentation criteria.
- **Campaign Creation Agent** — creates campaign briefs, flows, and multi-channel content from conversational prompts, accelerating campaign planning and development.
- **Content Creation Agent** — drafts and refines content for emails, landing pages, and SMS messages; can adjust tone, length, or audience relevance.
- **Journey Decisioning Agent** — selects the most appropriate journey/flow for an individual and can create personalized content for them.
- **Account Discovery Agent** — surfaces account insights and identifies potential buying-group members, helping teams discover stakeholders and guide nurturing/follow-up.

Campaigns walkthrough — Campaign Creation Agent + Einstein Segment Creation

1. From Marketing Home, select New Campaign → Draft with Agentforce.
2. Describe the campaign objective, choose to draft a brief, and ask Agentforce to generate the campaign.
3. Review the AI-generated campaign preview (flow + content); save it as a new campaign and open the campaign record.
4. Click Select Segment → Go to Segment Builder → Next → Create with Einstein Segment Creation.
5. Choose the DMO to segment on and click Next — Einstein has already drafted an audience definition based on the campaign objective, which can be refined or accepted as-is → Create Segment.

4.3 Predictive AI Features

Like Marketing Cloud Engagement, Marketing Cloud Next includes three predictive AI capabilities. Engagement Frequency and Engagement Scoring are Advanced Edition only; Send Time Optimization is available in both Growth and Advanced editions.




Figure 4.4 — The three predictive AI capabilities and how their outputs are used.

To make any predictive AI attribute available for decisioning in a flow, you must first configure the required data graph. Once configured, fields such as Email Engagement Frequency or Engagement Score can be used to branch a Decision element, or as criteria when building a segment.

4.3.1 Engagement Frequency (Advanced Edition only)

Predicts how often each contact should receive an email, aiming to maximize engagement while reducing email fatigue and unsubscribes. Analyzes up to 90 days of email activity (sends, opens, clicks, bounces, unsubscribes, spam complaints, sending patterns), comparing each individual's behavior at different sending frequencies against similar individuals, then calculating an ideal frequency range. Trained only on your org's own data; scores refresh weekly.




Figure 4.5 — Engagement Frequency: inputs, model behavior, and the four classification buckets.

The Email Engagement Classification field (on the Email Engagement Frequency DMO) assigns each individual to one of four categories:

- Saturated
- Almost saturated
- On target
- Undersaturated

Marketers can use these predictions as segment criteria or within a Decision element to route individuals — e.g., reducing sends to saturated individuals, increasing sends to undersaturated ones.

4.3.2 Engagement Scoring (Advanced Edition only)

Estimates how likely each individual is to engage with FUTURE marketing communications, by analyzing up to 90 days of historical behavior (email sends, opens, clicks, bounces, unsubscribes, subscription history, website interactions), weighting recent activity more heavily. Generates predictions for: likelihood to open, likelihood to click, and likelihood to remain subscribed.




Figure 4.6 — Engagement Scoring personas, mapped by open/click likelihood.

- **Loyalists** — high likelihood to both open and click.
- **Selective Subscribers** — low likelihood to open, high likelihood to click.
- **Window Shoppers** — high likelihood to open, low likelihood to click.

- **Win-Back / Dormant Subscribers** — low likelihood on both open and click.

Each of the three likelihood predictions is also grouped into one of four tiers — Most likely, More likely, Less likely, Least likely — usable in a Decision element to route individuals down different flow paths. Example strategies: send surveys to Selective Subscribers, engage Window Shoppers through alternative channels, place Dormant Subscribers into a dedicated win-back campaign.

4.3.3 Send Time Optimization (Growth AND Advanced Edition)

Uses predictive AI to determine WHEN each individual is most likely to engage with an email, analyzing up to 90 days of engagement history (sends, opens, clicks, bounces, unsubscribes, spam complaints, and associated timestamps) instead of sending everyone at the same time.

- Must be enabled in Setup.
- Once enabled, add the Send Time Optimization DMO to your data graph and select the "hourly scores by week" field.
- The Send Email Message flow element uses this field to determine each recipient's optimal time.
- Within the element, enable send time optimization and define an optimization window ranging from 2 hours up to 1 week.
- Predictions refresh approximately once a week.

Marketing Cloud Next then sends each recipient's message at their predicted optimal time within the configured window — e.g., one person in the morning, another in the evening.

Predictive AI exam takeaways

- Engagement Scoring = likelihood to engage (open/click/stay subscribed).
- Engagement Frequency = ideal SEND CADENCE (avoiding fatigue/oversaturation).
- Send Time Optimization = the best TIME to send, within a configurable 2-hour to 1-week window.
- Engagement Frequency and Engagement Scoring: Advanced Edition ONLY. Send Time Optimization: both editions.
- All three analyze up to 90 days of history and refresh weekly.

4.4 Analytics & Performance Insights




Figure 4.7 — Marketing Cloud Next's analytics and reporting architecture.

4.4.1 Marketing performance dashboards

Available from the Marketing Analytics tab, giving marketers a consolidated view of performance across marketing activities — overall campaign engagement, top-performing campaigns/segments/content/channels — bringing together web, email, SMS, WhatsApp, and mobile activity. The same dashboards are also embedded directly on each campaign record, so you can evaluate an individual campaign's results in context.

4.4.2 Content performance dashboards

Help marketers identify top-performing content directly in their flow of work. Access by selecting the relevant flow element (e.g., a Send Email Message element) → Analytics tab → Details button, which surfaces pulse-card metrics relevant to that content type plus lower-level detailed data.

4.4.3 Deliverability dashboards

Available for SMS, WhatsApp, and mobile app messaging — provide visibility into message delivery health and campaign performance, and help you understand why a specific message failed to reach specific recipients, so you can resolve delivery issues and improve future reach.

4.4.4 The Semantic Data Model (powered by Tableau Next)

For anything beyond the built-in dashboards — custom reports, sharing data with stakeholders outside Salesforce CRM, or connecting external tools — use the Semantic Data Model, available from the Analytics tab in the Marketing app. Marketing Cloud Next includes the Marketing Intelligence Semantic Data Model, which combines data from multiple Data 360 objects and defines the relationships and metrics needed to report on campaign performance, engagement, and deliverability. Because these definitions are centrally applied, reports and dashboards calculate key metrics consistently.

- Build reports that answer specific business questions, then combine reports as widgets in a dashboard for a broader view.
- Example: a report compares email engagement by campaign; a dashboard brings together performance across campaigns, channels, content, and audience segments.

- Standard dashboards can be used as-is or customized (widgets, labels, layouts, filters).
- The Semantic Model can be extended to other Data 360 objects for org-specific reporting needs.

4.4.5 Three reporting capabilities — summary

Capability	What it provides	License required?
Campaign Performance Dashboard	End-to-end view of an individual campaign — sends, opens, clicks, bounces, and contact progression	No — included
Marketing Performance Dashboard	Broader cross-channel view over time — audience growth, program-level engagement trends	No — included
Semantic Data Model	Consistent business layer for custom reports, external stakeholder sharing, and connecting other analytics tools	No — included

4.4.6 Unified Engagement History dashboards

Pre-built dashboards giving sales and marketing teams a shared view of the activities performed by leads and contacts associated with an account. Sales reps use these to identify the most engaged people, understand their interests, and determine who to contact next. Powered by Data 360 and Tableau Next; can be embedded directly on Account, Lead, and Contact records by adding the Unified Engagement History dashboard component to the page layout.

knowledge-check recap

- Reconciliation rules select a single unified value (e.g., name) — they do NOT apply to contact points, which are all retained.
- For phone-channel activation templates, use the telephone number field from the Contact Point Phone DMO.
- Engagement Scoring estimates likelihood to engage with future communications.
- Content performance dashboards are accessed from the Analytics tab within the relevant flow messaging element.

4.5 Topics Mentioned but Not Deep-Dived

This guide deliberately does not cover every Marketing Cloud Next feature. However the following areas are worth independent research before the exam, since they may still appear in the exam guide's broader objectives:

- Scoring rules
- Convergent features shared with Marketing Cloud Engagement
- Other channels: RCS and WhatsApp
- Retail triggers
- Einstein Metrics Guard
- Distributed marketing

Review the official exam guide's help-documentation links for these topics, and treat this document as a strong foundation rather than exhaustive coverage.

CloudCove.ai

Part 5 — Quick-Reference Cheat Sheet, Exam Traps & Glossary

Use this section in the final days before your exam. It distills every number, rule, and commonly confused pair from Parts 1–4 into fast-recall form.

5.1 Key Numbers Cheat Sheet

Number	What it refers to
60	Scored multiple-choice questions on the exam
Up to 5	Additional unscored questions (randomly integrated; don't affect your score)
105 minutes	Total exam time (~1 min 45 sec per question)
72% (44/60)	Passing score
Summer '26	Salesforce release the exam content is based on
4	Levels of consent granularity (MC Next enforces Level 4: contact point + subscription)
4	DMOs that make up the consent data model
90 days	TTL of a consent cache entry; also the lookback window for Engagement Frequency, Engagement Scoring, and Send Time Optimization
~15 minutes	Typical delay for a record ingested via data stream to materialize in the Communication Subscription Consent DMO
50,000	Max rows per CSV file import for consent
1	Channel + 1 subscription + 1 consent status per CSV consent import
1–3 seconds	Typical latency for an On-Demand Flow
2 hours – 1 week	Configurable Send Time Optimization window
Weekly	Refresh cadence for all three predictive AI features
6	Marketing / high-scale flow types (Audience, Activation-Triggered, Automation Event-Triggered, On-Demand, Broadcast, Data Cloud-Triggered)
8	Data providers available for email content personalization
5	Out-of-the-box AI agents (Segment Creation, Campaign Creation, Content Creation, Journey Decisioning, Account Discovery)
3	Predictive AI features (Engagement Frequency, Engagement Scoring, Send Time Optimization)
3	Built-in reporting capabilities (Campaign Performance, Marketing Performance, Semantic Data Model) — no extra license
2	Required permission sets for setup (Data Cloud Architect, Marketing Admin)

5.2 Editions at a Glance

Feature	Growth	Advanced
Engagement Scoring — individual	Yes	Yes
Engagement Scoring — account	No	Yes
Engagement Frequency	No	Yes
Send Time Optimization	Yes	Yes

5.3 Commonly Confused Pairs — Exam Traps

This...	...is NOT this
System Administrator (a profile)	Data Cloud Architect / Marketing Admin (permission sets needed for MC Next setup)
Authenticated domain (full DKIM/DNS + DMARC alignment)	Authorized email domain (ownership verification only; enables dynamic From/Reply-To)
CommunicationSubscriptionConsent (a CRM/COM object)	CommunicationSubscriptionConsentDMO (the actual Data 360 consent object)
HasOptedOutOfEmail field on Lead/Contact	A valid way to update consent in Marketing Cloud Next
Reconciliation rules (pick ONE value for a unified field, e.g. name)	Contact point handling (ALL contact points are retained, not reconciled)
Data Cloud-Triggered Flow (DMO change automation, e.g. consent)	A flow type that can send marketing emails
Record-Triggered Flow (core CRM flow)	A flow type that can send marketing emails from MC Next
Broadcast Flow (targets ALL members of a dynamic segment via REST API)	On-Demand Flow (targets ONE individual via REST API)
Activation data provider	Combinable with an Event data provider on the same message
Marketing object (marketer-managed CSV reference data)	A Data 360 DMO or a subscriber's personal profile record
Engagement Scoring (likelihood to engage)	Engagement Frequency (ideal send cadence) or Send Time Optimization (best send hour)
Common asset (available to all business units)	Default cross-business-unit content access (content is normally siloed per BU)

5.4 Glossary of Key Terms

Term	Definition
Agentforce Marketing	Salesforce's broader next-gen marketing portfolio brand (Marketing Cloud Next, Salesforce Personalization, Marketing Intelligence, Loyalty Management).
Marketing Cloud Next	The next-generation marketing platform product; a Lightning app built on Data 360.
DLO (Data Lake Object)	Raw, ingested data as landed from a source system via a data stream.
DMO (Data Model Object)	Structured, canonical data mapped from one or more DLOs (e.g., Individual, Contact Point).
Data graph	A pre-assembled, read-only view spanning multiple related DMOs. It connects related data to provide context for personalization, segmentation, and activation.

Term	Definition
Activation	Publishing a segment to a destination/target for use in engagement or personalization.
Activation template	A reusable activation definition (source priority order, filters) selectable when configuring a flow's audience/send properties.
Identity resolution	Matching and unifying individual records into a single unified profile.
Marketing flow / high-scale flow	A flow that runs off-core on a separate high-scale engine; orchestrates all outbound engagement in MC Next.
Actionable list	A fixed collection of audience members, ideal for known/curated audiences.
Consent service / cache layer	The mechanism providing read-after-write consistency for consent, with a 90-day TTL cache in front of the DLO/DMO.
Preference page	A customer-facing page for self-service subscription/consent management.
Business unit	A structure separating marketing content and data by team, brand, or region.
CMS content workspace	Where marketing content (images, pages, forms, emails) is stored, tied to a business unit.
Common asset	Content posted so it's available across all business units (with local copy required to edit).
Data provider	A structured source of data usable for email content personalization (8 types — see Part 3.4).
Content variable	A placeholder defined in content and populated at runtime by Salesforce Flow.
Marketing object	A marketer-managed, CSV-driven reference data table queried via Handlebars/AMPscript.
Engagement Scoring	Predictive AI estimating likelihood to open/click/stay subscribed (Advanced Edition only).
Engagement Frequency	Predictive AI recommending ideal send cadence per individual (Advanced Edition only).
Send Time Optimization	Predictive AI recommending the best send hour per individual (Growth + Advanced).
Semantic Data Model	The Tableau Next-powered reporting layer underlying MC Next's reports and dashboards.
Unified Engagement History	Pre-built dashboards showing lead/contact engagement activity on Account/Lead/Contact records.

Final reminder

- This guide is a study aid built from the available information as of now and is not a substitute for the official exam guide.
- Always verify current feature availability, edition mapping, and objective weightings against the live Salesforce exam guide and Help documentation before sitting the exam.
- Good luck!