

Exam Code: AB-620

Exam Name: AB-620: Designing and Building Integrated AI Solutions in Copilot Studio Training Course

Certification: Designing and Building Integrated AI Solutions in Copilot Studio

Vendor: AB

AB-620 Training Course

AB-620: Designing and Building Integrated AI Solutions in Copilot Studio Training Course

Structured Learning & Certification Preparation

Table of Contents

1. Introduction
 2. About This Training / Certification
 3. What We Offer (AAAdemy)
 4. Knowledge Overview
 5. Detailed Knowledge Explanation
 6. Learning Path & Study Advice
 7. Who This PDF Is For
 8. Call To Action
 9. Attachment: Answers by Knowledge Point
-

Introduction

This study pack is designed to support preparation for the Designing and Building Integrated AI Solutions in Copilot Studio exam through a clear, knowledge-point-driven structure. It brings the exam scope into one place so you can review Plan and configure agent solutions, Integrate and extend agents in Copilot Studio, Test and manage agents in the same order you are expected to master them.

The material is organized around 3 official blueprint domains, with each section keeping the detailed explanation content intact and pairing it with mapped practice questions. A practical way to use this pack is to move in a repeatable study, practice, and review cycle: study the explanation first, answer the related questions, then check the answer attachment to confirm where your understanding is already strong and where it still needs reinforcement.

About This Training / Certification

Designing and Building Integrated AI Solutions in Copilot Studio focuses on the ability to understand the core concepts, terminology, roles, operational practices, and decision-making patterns covered by the certification blueprint. The exam expects candidates to connect foundational knowledge with practical scenarios and choose actions that fit the stated business, technical, and operational context.

This training content supports that preparation by keeping the knowledge explanations structured and by pairing each exam domain with directly mapped practice questions. The result is a study pack that helps you connect key terms, domain concepts, practical trade-offs, and exam readiness in a format that is practical for steady exam preparation.

What We Offer (AAAdemy)

AAAdemy provides structured training resources designed to support certification preparation and skill development across a wide range of IT domains. Our learning materials are built around clear knowledge structures, practical study guidance, and exam-oriented practice to help learners progress with confidence.

We offer well-organized knowledge explanations that break down complex topics into clear, understandable sections aligned with official exam objectives and real-world skill requirements. Each topic is designed to support both conceptual understanding and practical application.

Our study plans and learning guidance help learners follow a logical progression, focusing on key concepts, common pitfalls, and effective preparation strategies. This approach enables learners to study efficiently while maintaining a clear view of their learning goals.

To reinforce understanding, AAAdemy also provides practice questions and exam-focused insights that reflect typical certification scenarios. These resources are intended to help learners evaluate their readiness and strengthen their confidence before taking an exam.

All content is designed for flexible, self-paced learning, allowing individuals to study independently or alongside their existing professional or academic commitments.

Knowledge Overview

- Plan and configure agent solutions
 - Plan responsible AI, identity, audience, channel, governance, and reusable component boundaries
 - Create and monitor agent flows including human-in-the-loop approval, inputs, outputs, connectors, and error handling
 - Configure topics with variables, Power Fx expressions, tools, generative answers, custom prompts, API calls, and adaptive cards
- Integrate and extend agents in Copilot Studio
 - Connect agents to enterprise knowledge sources including Copilot connectors, Power Platform connectors, and Azure AI Search
 - Add tools to agents with computer use, MCP tools, custom connectors, and REST APIs
 - Configure multi-agent collaboration with Foundry agents, Copilot Studio agents, Fabric data agents, and A2A protocol
 - Integrate agents with Azure AI Search, Foundry model catalog prompts, and Application Insights monitoring
- Test and manage agents

- Evaluate agent performance with test sets, evaluation methods, and result review
 - Implement ALM for Copilot Studio agents with solutions, environment variables, and Power Platform Pipelines
-

Detailed Knowledge Explanation

Plan and configure agent solutions

Plan responsible AI, identity, audience, channel, governance, and reusable component boundaries

Exam Radar

Core Priority: This is a planning-first topic. The exam usually gives a proposed audience, a sensitive action, or an enterprise data source and asks what must be decided before implementation. The correct answer protects identity, responsible AI review, channel exposure, Power Platform environment governance, and reusable component ownership before any topic or tool is built.

High Frequency: Expect questions about internal versus external audiences, identity strategy, deployment channels, environment strategy, DLP policy, responsible AI strategy, security review, and reusable agent components.

Confusion Alert: A planning question may mention topics or connectors, but the real conflict is usually exposure risk: external audience without identity, sensitive action without review, or reusable component without environment ownership.

Scenario Logic: A developer must decide whether an internal HR agent should use Microsoft Entra ID, Teams deployment, managed Power Platform environments, and reusable agent components before building topics. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: AB-620 now treats planning as broader than topic design: responsible AI strategy, identity, internal/external audience, deployment channel, security, governance, and reusable components are explicit exam scope.

Failure Trigger: Failure appears as a blocked channel publish, overexposed data source, unreviewed high-risk action, DLP-blocked connector, or duplicated component that drifts between environments.

Operational Dependency: The design must lock audience, identity, deployment channel, governance, and reuse before implementation because those choices control permission scope, data exposure, deployment model, and component ownership.

How the Exam Asks It: A question may describe a business scenario, a failing Copilot Studio configuration, or a deployment constraint and ask for the best design, first troubleshooting step, or required component.

How Distractors Are Designed: Distractors often solve a visible symptom while ignoring the owning object. They may suggest prompt tuning for a permission problem, a connector for a retrieval problem, a channel change for a flow exception, or manual deployment for an ALM dependency.

Why the Correct Answer Works: The correct option works because it chooses the design boundary that controls later implementation choices: audience drives identity, identity drives allowed data/action access, and responsible AI/governance rules decide what must be reviewed before publication.

Best Choice Rules: If the stem mentions sensitive actions or compliance, choose responsible AI controls, escalation, and auditability before implementation. If the stem mentions internal or external users, prioritize audience, identity, and channel compatibility before topic wording. If the stem mentions reuse across agents, prioritize solution-aware reusable components over copying topics manually.

Atomic Deconstruction - Operational Level

An agent solution plan is the control document for who can use the agent, which systems it can reach, which actions require review, and which components can be reused. Responsible AI strategy is not a late documentation task; it defines escalation, sensitive-content handling, traceability, and human review boundaries before tools are exposed to users.

Operationally, planning starts with a risk inventory. List every audience, channel, enterprise system, connector, human-review point, and reusable component. Then decide which Power Platform environment owns the asset, which policies can block or allow connectors, and which administrator role must approve publication. This keeps security and responsible AI decisions upstream of maker convenience.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | -----
- | ----- | ----- |

| Audience model | Internal or external users | Internal, external, mixed | Unset until channel chosen | Identity provider and channel policy | Wrong audience exposes private tools or blocks valid users |

| Identity strategy | Authentication boundary | Entra ID, connector identity, delegated user | Environment inherited | Power Platform environment and data connector policy | Tool calls fail with 401/403 or run under the wrong principal |

| Channel plan | Deployment surface | Teams, website, Microsoft 365 Copilot, supported channels | Draft agent only | Channel authentication and publishing policy | Agent cannot reach intended users or cannot satisfy compliance review |

| Reuse model | Component ownership | Topics, tools, flows, prompts, knowledge sources | Local to agent | Solution packaging and naming standards | Duplicated components create drift across agents |

| Responsible AI strategy | Human review and safety boundary | Refuse, confirm, escalate, audit | Implicit default behavior | Sensitive action classification and policy review | Agent performs or explains sensitive actions without review evidence |

Step-by-Step Execution Path

1. In Copilot Studio, inspect the agent audience, channel, authentication, and environment settings before editing topics. This establishes the control boundary that later tool calls must obey.
2. Classify each planned action by business risk and decide whether it needs confirmation, refusal, escalation, or human review. This places responsible AI strategy before tool exposure.
3. In the Power Platform admin center, verify the target environment, data loss prevention policy, and connector availability. This prevents a design that works in a maker sandbox but fails in a governed environment.
4. Map each planned enterprise system to an access method: Copilot connector, Power Platform connector, custom connector, REST API, MCP server, Fabric data agent, or Foundry agent. The map separates knowledge retrieval from action execution.
5. Document which components are reusable across agents and which are agent-specific. Reusable tools and flows should be solution-aware so ALM can move them predictably.

Evidence note: planning checks rely on design-review records, Copilot Studio settings, Power Platform environment configuration, DLP policy state, and administrator approval evidence rather than unverified CLI commands.

Technical Chain

A planning decision starts with the intended audience and data/action risk. Audience choice drives authentication and channel constraints; authentication drives connector identity and permission trimming; environment and DLP policy decide which connectors and flows are allowed; responsible AI rules decide when the agent must refuse, escalate, ask for confirmation, or require human approval. If this chain is skipped, the agent may appear functional in a maker test but fail security review, expose the wrong data, or perform actions without the review boundary the scenario requires.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | -----
----- |

| Validate environment boundary | Power Platform admin center > Environments > select environment > Policies and resources | Environment, region, security group, and DLP policy match the planned audience |

| Verify channel readiness | Copilot Studio > agent > Channels | Only channels approved for the audience and authentication model are enabled or planned |

| Review component reuse | Copilot Studio > agent > Tools, Topics, Knowledge; Power Apps solution view | Shared components are named, owned, and packageable inside the intended solution |

| Confirm identity dependency | Connector connection references and agent authentication settings | Tool execution identity aligns with least-privilege access and expected user delegation |

| Validate responsible AI boundary | Design review evidence for sensitive actions, escalation rules, refusal behavior, and audit trail | High-risk actions have an explicit confirmation, escalation, or review path before publishing |

Create and monitor agent flows including human-in-the-loop approval, inputs, outputs, connectors, and error handling

Exam Radar

Core Priority: This is an execution-contract topic. AB-620 may describe a failed action, a pending approval, or a vague completion message and expect the candidate to inspect input mapping, output schema, connector status, human-in-the-loop state, and run-history evidence.

High Frequency: Expect scenarios about creating agent flows, creating human-in-the-loop agent flows, adding input/output parameters, configuring connectors, monitoring flows, and implementing handled error paths.

Confusion Alert: A flow question may look like a prompt-quality problem, but the decisive clue is often a missing input, unhandled connector status, unresolved approval state, or absent run-history detail.

Scenario Logic: A support agent calls an agent flow to create a ticket. The flow works for happy-path values but silently fails when the API returns a validation error. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: The current scope calls out human-in-the-loop agent flows in addition to ordinary agent flows, so approval state, reviewer outcome, timeout behavior, and flow monitoring should be treated as first-class execution evidence.

Failure Trigger: Failure appears as no run history for the expected step, no approval state, a connector 4xx/5xx status, a timeout without mapped output, or a topic response that says completed when the flow is still pending.

Operational Dependency: Agent flows require a clear execution contract and observable failure handling; typed parameters and run monitoring expose whether the tool failed because of input shape, connector authentication, or downstream API validation.

Why the Correct Answer Works: The correct option works because it observes the flow as a contract: inputs enter, connector or approval steps execute, errors are handled, outputs return, and run history proves which stage failed.

Best Choice Rules: If the stem mentions approval, reviewer decision, or manual confirmation, inspect human-in-the-loop flow behavior before generic retry logic. If the stem mentions API validation errors, inspect input mapping, connector action details, error handling, and run history before prompts. If the stem mentions user-visible confirmation, verify the output parameter contract returned by the flow.

Atomic Deconstruction - Operational Level

An agent flow is the deterministic execution layer behind a conversational request. It receives typed values from the topic or orchestration layer, performs connector or approval work, and returns a structured result. A human-in-the-loop flow adds a review state: the agent must pause or continue based on approval outcome instead of pretending the action finished immediately.

Operationally, the flow should be read as a state machine. Required inputs must be present before a connector action runs; approval nodes must return approved, rejected, pending, or timed-out states; error branches must preserve the downstream status; outputs must be mapped back to topic variables so the agent can explain the result without inventing completion.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | -----
----- | ----- |

| Input parameter | Name and data type | Text, number, boolean, record-like values where supported | Not present | Topic variable mapping | Flow receives null or malformed values and connector action rejects request |

| Output parameter | Returned value contract | Status, ID, message, structured output | Not present | Agent response node | Agent cannot present confirmed result or recoverable error |

| Connector action | Connection reference | Standard, premium, custom connector | Maker-owned connection | DLP policy and authentication | Runtime failure from blocked connector or invalid token |

| Error branch | Failure path | Retry, fallback message, escalation, logged error | Unmodeled failure | Connector status and exception payload | User receives vague response while action did not complete |

| Human approval step | Decision state | Approved, rejected, pending, timed out | No approval gate | Approver identity and notification path | Agent reports completion before a human decision exists |

Step-by-Step Execution Path

1. Open Copilot Studio > agent > Tools or Topics and inspect where the agent flow is invoked. This confirms which user values are mapped into the flow contract.
2. Open the agent flow and verify input and output parameters before connector logic. Parameters must represent the exact values the topic can supply and the response can consume.

3. For human-in-the-loop scenarios, add an approval or review step before the irreversible connector action and map approved, rejected, and timeout states to output values. This prevents the topic from treating a pending human decision as a completed business transaction.
4. Inspect each connector action connection reference and DLP compatibility in the target environment. This excludes authentication and policy failure before debugging business logic.
5. Run a controlled test with a known invalid value and review flow run history. The error branch should capture connector status, validation message, and user-safe response text.

Evidence note: flow checks rely on Copilot Studio or Power Automate run history, approval state, connector action details, and mapped output values.

Technical Chain

The topic passes collected values into the agent flow. The flow binds those values to parameters, evaluates connector actions and optional human approval steps, then writes action-level run history. If the approval owner rejects or times out, the flow must return a status the topic can explain. If the connector returns a validation error, the error branch must preserve the status and message. Without typed outputs, the conversation layer cannot distinguish completed, rejected, pending, and failed states.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | -----
 ----- |

| Inspect flow run status | Copilot Studio or Power Automate run history for the agent flow | Run shows success, failed, or handled failure with timestamp and action-level details |

| Validate parameter mapping | Copilot Studio topic node > flow input mapping | Every required input is bound to a topic variable with the expected type |

| Check connector policy | Power Platform admin center > DLP policy > connector classification | Connector used by the flow is allowed in the environment's policy group |

| Confirm error response | Test pane conversation transcript and flow run details | User receives a specific recoverable message and the run records the downstream error |

| Validate approval state | Agent flow run history > approval or review action detail | Run records approved, rejected, pending, or timed-out state and returns a mapped output to the topic |

Configure topics with variables, Power Fx expressions, tools, generative answers, custom prompts, API calls, and adaptive cards

Exam Radar

Core Priority: This is a conversation-orchestration topic. The exam tests whether the candidate can place variables, Power Fx expressions, tools, generative answers, custom prompts, API calls, response formatting, and adaptive cards in the correct topic sequence.

High Frequency: Expect topic questions that combine trigger conditions, topic variables, global variables, Power Fx expressions, response formatting, tool nodes, generative answers, adaptive cards, and Send HTTP request behavior.

Confusion Alert: A topic question often hides the fault in state management: a stale global variable, a Power Fx expression returning the wrong type, or an adaptive card binding that runs before the API response exists.

Scenario Logic: A sales operations agent must answer policy questions from a knowledge source, call an API for customer status, and show a compact card for confirmation. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Topic configuration now spans variables, Power Fx-style logic, response formatting, generative answers, custom prompts, custom knowledge, API/Send HTTP requests, tools, and adaptive cards rather than only trigger phrases.

Failure Trigger: Failure appears as an empty topic variable, stale global variable, Power Fx type mismatch, missing tool input, unsupported adaptive card field, or Send HTTP request node that receives a malformed body.

Operational Dependency: The topic owns the conversation path: it decides when to ground, when to call a tool, when to call an API, and how to format the final response.

Why the Correct Answer Works: The correct option works because it places orchestration in the topic, where variables, Power Fx expressions, knowledge grounding, tool calls, API requests, and adaptive card output can be sequenced.

Best Choice Rules: If the stem mentions wrong values in a tool call, inspect topic variables, global variables, and Power Fx expressions before changing the connector. If the stem mentions response formatting, separate adaptive card/schema problems from retrieval or API execution problems. If the stem mentions generative answers, verify the knowledge source and grounding node before adding unrelated tools.

Atomic Deconstruction - Operational Level

A topic is a stateful conversation path. Variables hold collected facts, global variables carry shared context, and Power Fx expressions can transform values, branch conditions, or format data before the agent calls a tool or shows a response. Without variable discipline, a topic can call the right API with the wrong value or render an adaptive card before required data exists.

Operationally, the topic should be walked node by node. Confirm the trigger, inspect each variable assignment, check Power Fx expressions with sample values, verify the grounding node, then review tool/API input mapping and card binding. This order prevents the common mistake of debugging the API when the actual failure is a stale variable or formula result.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | -----
----- | ----- | ----- | ----- | -----
|

| Trigger and condition | Intent boundary | Phrases, variables, topic conditions | Broad match | User utterance and topic priority | Wrong topic intercepts request or misses valid intent |

| Generative answers node | Grounding source | Selected knowledge source or search connection | Unconfigured | Knowledge source availability | Ungrounded answer or no answer for known content |

| Tool/API action | Invocation point | Connector tool, agent flow, Send HTTP request, REST API | Not added | Authentication and input mapping | Tool runs with missing values or wrong permission |

| Adaptive card | Schema and variable binding | Supported Adaptive Cards schema subset | Plain text response | Topic variables and channel support | Card renders blank, invalid, or unsupported in channel |

| Topic variable | Scope and type | Topic-scoped value, global variable, system variable | Unset until collected or assigned | Question node, Power Fx expression, tool output | Tool receives stale, empty, or wrongly transformed value |

| Power Fx expression | Transformation or condition | Formula returning text, number, boolean, table-like value where supported | No formula | Variable type and function availability | Branch or API payload uses a technically valid but semantically wrong value |

Step-by-Step Execution Path

1. In Copilot Studio, open the topic and validate the trigger phrases or conditions. This makes sure the topic is reached only for the intended scenario.
2. Open variable management for the topic and identify topic variables, global variables, and any Power Fx expressions used for conditions or formatting. This excludes stale or incorrectly transformed values before blaming the connector.
3. Add or inspect the generative answers node and bind it to the intended knowledge source. This prevents the model from answering from unsupported context when policy accuracy matters.
4. Add the tool, agent flow, REST API, or Send HTTP request node after the topic has collected required inputs. Tool calls should happen after grounding or clarification when the downstream action needs exact values.

5. Use an adaptive card only after the result variables exist. Validate the card schema and variable binding in the test pane for each target channel.

Evidence note: topic checks rely on the authoring canvas, variable inspector, formula or condition editor, test-pane transcript, and target-channel rendering.

Technical Chain

The user utterance selects a topic, topic nodes collect or derive variables, Power Fx may transform or validate those variables, and the topic then chooses a grounding node, prompt, tool, HTTP request, or adaptive card. The response is reliable only when each variable has a known source, type, and scope. If a topic variable is empty, a global variable is stale, or a Power Fx expression returns an unexpected value, the downstream action may execute with a technically valid but semantically wrong payload.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

|-----|-----|-----|-----|

| Validate topic matching | Copilot Studio test pane > send scenario utterances | Expected topic is selected and no unrelated topic takes priority |

| Check grounding behavior | Topic node settings > generative answers knowledge source | Node points to the intended source and returns citations or grounded answer evidence where supported |

| Inspect API/tool inputs | Topic action node > input mapping | Required parameters are populated from confirmed variables before execution |

| Verify adaptive card rendering | Copilot Studio test pane and target channel preview | Card displays expected fields and actions without schema or channel rendering errors |

| Inspect variable scope | Copilot Studio topic authoring canvas > variables or node variable mapping | Topic, global, and system variables have expected source, type, and latest value |

| Validate Power Fx result | Copilot Studio formula or condition editor with test input values | Expression returns the expected type and value before the tool or response node uses it |

Practice Questions

1. Case study: A company is planning an external supplier-facing agent that answers order-status questions and can submit dispute requests to an internal system. The supplier audience is not in the tenant, and disputes require review. What should be defined before building topics?
 - A. Audience identity, channel exposure, responsible AI review, DLP policy impact, and reusable component ownership
 - B. Adaptive card color tokens and final response tone

- C. A single broad topic that calls every supplier API
 - D. Anonymous website access until the dispute workflow is complete
2. An internal HR agent will be published in Teams and reuse an approval flow from another agent. Which planning choice best supports governed reuse?
- A. Copy the flow into each agent so every maker can edit it independently
 - B. Package the shared flow, topics, and tools in a solution with clear ownership and connection references
 - C. Store the flow URL in the topic message so it is visible to makers
 - D. Create a separate unmanaged environment for every department
3. A finance agent can create refund requests, but the organization requires human review for refunds above a threshold. What should the solution include?
- A. A longer system prompt warning the user about refunds
 - B. A public channel so reviewers can access the conversation
 - C. A human-in-the-loop flow path that returns approved, rejected, pending, and timeout states
 - D. A generative answers node that summarizes refund policies
4. A topic calls an agent flow to create a case. The flow fails when one optional field is blank, and users see a generic failure message. What should be changed first?
- A. Republish the same agent to Teams
 - B. Move all case creation logic into a generative answer
 - C. Remove the connector from the flow
 - D. Validate input parameters and add an error branch that maps connector status to a user-safe response
5. A Copilot Studio topic sends an API request with an account ID from an earlier conversation after the user enters a new account ID. What should be inspected?
- A. Topic variables, global variables, and Power Fx expressions used to build the request body
 - B. Azure AI Search indexer history
 - C. The channel publishing policy
 - D. The solution import history
6. An agent displays an adaptive card after a REST API call, but several fields are blank even though the API call succeeds. What should be checked first?
- A. Whether the agent has been published to a different channel
 - B. Whether API response fields are mapped to the topic variables used by the adaptive card
 - C. Whether the knowledge source contains the same customer ID
 - D. Whether the flow has an approval owner
7. Order steps: A topic must answer policy questions and create a service request only after the user confirms request details. Which order is most appropriate?
- A. Call the service API, ground the answer, then ask for confirmation
 - B. Use global instructions, create the request, then collect variables

- C. Ground the answer, collect and validate required variables, confirm details, then call the service request tool
 - D. Render an adaptive card, create the request, then validate the API payload
8. A maker uses Power Fx expressions to compute a priority value before invoking a tool. What must be validated before the tool call?
- A. Whether the response contains citations
 - B. Whether Application Insights has exception telemetry
 - C. Whether the topic variable feeding the expression is populated by the previous node
 - D. Whether the expression returns the expected type and value for representative inputs
9. A planning review flags that a proposed agent can answer sensitive employee disciplinary questions without escalation. What should be added to the design?
- A. Responsible AI controls for refusal, escalation, review evidence, and auditability
 - B. A broader list of trigger phrases
 - C. A custom connector for every HR table
 - D. A different adaptive card layout
10. A topic should call a tool only when all required variables are available. In testing, the tool runs with a missing region value. Which fix is best?
- A. Increase the model creativity setting
 - B. Add a condition or clarification step before the tool node to collect and validate the region variable
 - C. Move the tool to a different environment without changing the topic
 - D. Replace the topic with a static FAQ

Integrate and extend agents in Copilot Studio

Connect agents to enterprise knowledge sources including Copilot connectors, Power Platform connectors, and Azure AI Search

Exam Radar

Core Priority: This is a retrieval-boundary topic. The exam may describe ServiceNow, SAP, Power Platform data, documents, or indexed engineering content and ask which knowledge connection gives the required freshness, schema, and permission behavior.

High Frequency: Expect comparison between Copilot connectors, Microsoft Power Platform connectors, Azure AI Search, RAG expectations, source-system permissions, searchable fields, filters, index freshness, and permission trimming.

Confusion Alert: A knowledge question may mention weak answers, but the cause can be stale indexer execution, non-retrievable fields, missing filters, or source permissions rather than prompt wording.

Scenario Logic: An agent must answer support-policy questions from ServiceNow articles, product documents, and indexed engineering content with permission-aware retrieval. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Enterprise knowledge scope includes Copilot connectors, Microsoft Power Platform connectors, and Azure AI Search, so the exam can test RAG-style retrieval, source freshness, and permission-aware access rather than file upload alone.

Failure Trigger: Failure appears as an indexer error, zero retrieved documents, missing searchable/retrievable fields, stale document count, or a permission-trimmed result that does not match the test user's access.

Operational Dependency: Knowledge integration is a retrieval and permission problem first; the correct source choice depends on where content lives, how it is indexed, and how access trimming is enforced.

How the Exam Asks It: A question may describe a business scenario, a failing Copilot Studio configuration, or a deployment constraint and ask for the best design, first troubleshooting step, or required component.

How Distractors Are Designed: Distractors often solve a visible symptom while ignoring the owning object. They may suggest prompt tuning for a permission problem, a connector for a retrieval problem, a channel change for a flow exception, or manual deployment for an ALM dependency.

Why the Correct Answer Works: The correct option works because retrieval quality and permission behavior are owned by the selected knowledge integration, not by a prompt-only workaround.

Best Choice Rules: If the stem mentions source-system permissions, validate connector identity and permission-aware retrieval. If the stem mentions stale or irrelevant RAG answers, check indexer freshness, schema, retrievable fields, and filters before prompts. If the stem asks for business action execution, do not answer with a knowledge-only retrieval source.

Atomic Deconstruction - Operational Level

Enterprise knowledge integration is not the same as action execution. A knowledge source feeds retrieval and grounding; a tool or API performs an action. Copilot connectors are useful for supported enterprise sources, Power Platform connectors can expose platform data and operations, and Azure AI Search is relevant when the solution needs controlled indexing, field schema, filtering, or vector and hybrid retrieval.

Operationally, start with source ownership and access behavior. A connector-backed source must prove connection health and permission behavior; Azure AI Search must prove indexer freshness, searchable/retrievable fields, filters, and vector or hybrid schema where used. Prompt changes come after retrieval evidence, not before it.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | -----
----- | ----- |

| Copilot connector | Source system connection | ServiceNow, SAP, supported enterprise sources |
Disconnected | Connector configuration and permissions | Agent cannot retrieve governed enterprise content |

| Power Platform connector | Data/action bridge | Standard, premium, custom | Connection required | DLP and
connection reference | Connector blocked or returns unauthorized data |

| Azure AI Search index | Schema and retrieval fields | Searchable, filterable, vector fields | No index attached
| Indexer and identity | Relevant documents are not returned or cannot be filtered |

| Permission trimming | User access enforcement | Delegated or source-aware permissions where supported |
Unverified | Identity mapping | Agent may over-share content or hide authorized content |

Step-by-Step Execution Path

1. Classify the source as enterprise connector content, Power Platform connector data, or indexed search content. This avoids using an action tool when the requirement is knowledge retrieval.
2. For connector-backed sources, verify connection status and identity behavior in Copilot Studio or the connector admin surface. The retrieval path must respect the intended user permission model.
3. For Azure AI Search, inspect the index schema, searchable fields, vector fields, filters, and indexer status in the Azure portal. Retrieval quality usually fails at schema or indexing before it fails at prompt wording.
4. Test with a user who should and should not see a known item. Permission-aware evidence is stronger than a generic successful answer.

Evidence note: retrieval checks rely on connector status, source permissions, Azure AI Search index schema, indexer execution history, and grounded answer evidence where supported.

Technical Chain

A user question is routed to a knowledge retrieval path. The connector or search integration uses its configured identity, source connection, index, filters, and retrievable fields to return candidate content. The agent then grounds the response in those retrieved passages or records. If the wrong integration type is chosen, the result may be stale, over-permissive, missing filters, or treated as an action call instead of retrieval evidence.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | -----
----- |

| Inspect connector status | Copilot Studio > Knowledge or Tools > selected connector configuration |
Connection is active and scoped to the intended source system |

| Check Azure AI Search index schema | Azure portal > AI Search > Indexes > selected index | Searchable,
retrievable, filterable, and vector fields match the retrieval design |

| Validate indexer freshness | Azure portal > AI Search > Indexers > execution history | Latest run succeeded
or exposes a specific source/indexing failure |

| Test permission-aware retrieval | Copilot Studio test pane with representative user access | Authorized
content is returned and restricted content is not exposed |

Add tools to agents with computer use, MCP tools, custom connectors, and REST APIs

Exam Radar

Core Priority: This is an action-surface topic. The question normally gives a target system and asks which tool mechanism matches the available interface: UI automation, MCP server tools, governed connector operations, or direct HTTP API calls.

High Frequency: Expect scenarios about computer use, MCP tools, existing custom connectors, REST APIs, authentication, operation schemas, response mapping, and monitoring of tool execution.

Confusion Alert: A tool question may name several integration options; the trap is choosing the most familiar option instead of matching UI automation, MCP schema, custom connector governance, or REST payload requirements to the target system.

Scenario Logic: A Copilot Studio agent must update a legacy system that has no modern API, call an internal API, and expose a governed reusable action. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Tool scope includes computer use, MCP tools, existing custom connectors, and REST APIs; the candidate must match the integration type to the external system surface.

Failure Trigger: Failure appears as an MCP tool not listed, rejected tool arguments, custom connector authentication error, REST 401/403/422 response, or computer use run stopping at the first unexpected screen state.

Operational Dependency: Tools are action surfaces with different contracts; selecting the tool type by integration boundary prevents unsupported automation and ambiguous execution.

Why the Correct Answer Works: The correct option works because it matches the tool mechanism to the real action surface, then validates schema, authentication, and runtime evidence.

Best Choice Rules: If the target has no API and must be operated through a UI, choose computer use with screen checkpoints. If a server exposes callable tool schemas, choose MCP and validate tool discovery and

arguments. If the API should be governed and reused in Power Platform, prefer a custom connector over ad hoc HTTP calls.

Atomic Deconstruction - Operational Level

A tool extends the agent from answering into doing. Computer use is a UI automation path, MCP exposes tool functions from a server, custom connectors package governed API operations, and REST API actions call an endpoint directly. The right choice depends on the external system surface, not on which option sounds most advanced.

Operationally, identify the external action surface first. UI-only targets need computer use checkpoints; MCP targets need tool discovery and argument schema; governed APIs need custom connector operations and connection references; direct HTTP targets need method, header, body, status, and response mapping checks.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | ----- |
----- | ----- |

| Computer use | Automation target | UI-driven application or website | Not configured | Credentials, session, UI stability | Action fails when screen state changes or authentication expires |

| MCP tool | Server-exposed function | Tool name, schema, transport, auth | No server attached | MCP server availability and schema | Tool not listed or arguments rejected |

| Custom connector | API operation definition | OpenAPI, auth, policy, connection reference | Unpublished connector | Power Platform DLP and connection | Operation blocked or called with wrong auth |

| REST API action | HTTP request contract | Method, URL, headers, body, response schema | No action | Endpoint auth and network reachability | HTTP 4xx/5xx or unusable response mapping |

Step-by-Step Execution Path

1. Identify whether the task is UI automation, tool-server invocation, governed API reuse, or one-off HTTP integration. This first split determines which runtime owns execution and authentication.
2. For MCP, inspect the tool schema and argument names before adding it to the agent. The agent can only call what the server exposes in its tool contract.
3. For custom connectors, validate the OpenAPI definition, authentication type, test operation, and DLP classification before using it in a topic.
4. For REST APIs, create a small request with explicit method, headers, body, and response handling. Test expected 2xx and known 4xx responses so the agent can distinguish user-fixable errors from system failures.

5. For computer use, record the UI path and define checkpoints after each screen transition. UI automation needs observable screen state, not only a desired business result.

Evidence note: tool checks rely on MCP tool discovery, custom connector test results, REST action response mapping, computer-use run state, and runtime transcript evidence.

Technical Chain

The agent selects a tool because the topic or orchestration layer has enough values to call it. The selected tool validates its schema, identity, endpoint, and runtime state. MCP rejects arguments that do not match its tool schema; custom connectors depend on OpenAPI operation definitions and connection references; REST APIs depend on method, headers, body, and status mapping; computer use depends on screen state and credentials. The observed failure tells which contract broke.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | -----
----- |

| Verify MCP tool exposure | MCP server tool list or Copilot Studio MCP tool configuration | Expected tool name, schema, and authentication state are visible |

| Inspect custom connector operation | Power Apps > Custom connectors > Test operation | Operation returns expected status and schema with the configured connection |

| Validate REST response mapping | Copilot Studio action test or API test console | Response status, body fields, and error payload map to agent variables |

| Check computer-use checkpoint | Copilot Studio computer use test run | Run reaches the expected UI state and records failure at the first broken screen transition |

Configure multi-agent collaboration with Foundry agents, Copilot Studio agents, Fabric data agents, and A2A protocol

Exam Radar

Core Priority: This is a delegation-architecture topic. The exam can ask when Copilot Studio should hand off to another Copilot Studio agent, a Microsoft Foundry agent, a Fabric data agent, or an A2A-compatible remote agent.

High Frequency: Expect questions about responsibility boundaries, handoff conditions, Foundry agent integration, existing Copilot Studio agents, Fabric data agents, A2A protocol, agent discovery metadata, and input/output contracts.

Confusion Alert: A multi-agent question may sound like a prompt-routing issue, but the failure is often an undefined handoff condition, missing specialist permissions, invalid A2A metadata, or a response schema the orchestrator cannot consume.

Scenario Logic: A customer support agent must delegate statistical data questions to a Fabric data agent, complex reasoning to a Foundry agent, and service workflow actions to another Copilot Studio agent. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Multi-agent scope includes Microsoft Foundry agents, existing Copilot Studio agents, Fabric data agents, and A2A protocol, so delegation contracts matter as much as single-agent topic design.

Failure Trigger: Failure appears as the local agent answering outside its boundary, delegation loops, Fabric workspace access denied, Foundry agent unavailable, invalid A2A agent card, or task response shape mismatch.

Operational Dependency: Multi-agent design succeeds when delegation has explicit responsibility boundaries, contracts, identity, and routing conditions; otherwise agents duplicate work or call the wrong specialist.

Why the Correct Answer Works: The correct option works because delegation becomes reliable only when each specialist agent has a defined responsibility, handoff condition, identity boundary, and response contract.

Best Choice Rules: If the stem mentions analytics over Fabric data, inspect Fabric workspace permissions and the Fabric data agent boundary. If the stem mentions remote-agent interoperability, inspect A2A metadata, endpoint trust, authentication, and task schema. If the stem mentions wrong specialist responses, fix delegation conditions and contracts before merging prompts.

Atomic Deconstruction - Operational Level

Multi-agent collaboration works only when each agent has a narrow responsibility and a clear contract. The orchestrating agent should know which requests it keeps, which requests it delegates, which values must be passed, and what response shape is expected. Fabric data agents are suited to data questions, Foundry agents to configured model-backed capabilities, and A2A to interoperable remote-agent task exchange when native integration is not the correct boundary.

Operationally, write the handoff contract before connecting agents. The orchestrator needs a clear trigger for delegation, a payload schema, an identity path, a retry or failure response, and a rule for resuming the conversation after the specialist returns a result.

Component Specifications

Object	Attribute	Value Range	Default State	Dependency	Failure State
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----

| Orchestrating agent | Routing responsibility | Intent, condition, skill boundary | Single-agent behavior | Delegate agent availability | Wrong agent handles request or loop occurs |

| Foundry agent | Reasoning or model-backed task | Connected Foundry agent | Not integrated | Foundry project, identity, model deployment | Delegation fails or returns untrusted output |

| Fabric data agent | Data question boundary | Semantic model or data source scope | Not connected | Fabric workspace permissions | Analytical answers are unavailable or unauthorized |

| A2A integration | Agent communication contract | Agent card, endpoint, auth, task exchange | No protocol configuration | Remote agent endpoint and trust | Remote agent cannot be discovered or invoked |

Step-by-Step Execution Path

1. Write a responsibility map before configuration: which agent owns retrieval, analytics, reasoning, or action execution. This prevents duplicated tools and cyclic handoffs.
2. In Copilot Studio, add the existing Copilot Studio agent, Foundry agent, or Fabric data agent using the supported integration path for that agent type. Native integration should be preferred when it satisfies the contract.
3. For A2A scenarios, validate agent discovery metadata, endpoint trust, authentication, and task input/output schema. A2A adds inter-agent protocol behavior and should be treated as a contract, not a generic connector.
4. Run handoff tests with boundary questions: one that should stay local, one that should delegate, and one that should be rejected or clarified. The transcript should show the decision point.

Evidence note: collaboration checks rely on integration settings, agent availability, A2A metadata, workspace permissions, handoff transcript, and returned task schema.

Technical Chain

A user request first reaches the orchestrating agent. The orchestrator evaluates intent, confidence, required data, and specialist ownership. When a handoff condition matches, it builds a task payload and invokes the specialist agent through the supported integration or A2A contract. The specialist returns a structured result or failure state, and the orchestrator resumes the conversation. If responsibility boundaries are vague, agents can duplicate answers, route in loops, or expose data through the wrong identity path.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | ----- |
----- |

| Validate agent routing | Copilot Studio test transcript for representative prompts | Expected agent receives the delegated task and local agent does not answer outside its boundary |

| Inspect Foundry integration | Copilot Studio > Tools or agent integration settings; Microsoft Foundry project status | Connected Foundry agent is available and authorized for the environment |

| Check Fabric data agent access | Fabric workspace and Copilot Studio integration settings | Workspace permissions and data agent selection match intended users |

| Verify A2A contract | A2A endpoint metadata and task response test | Agent card, endpoint, authentication, and task response schema are valid |

Integrate agents with Azure AI Search, Foundry model catalog prompts, and Application Insights monitoring

Exam Radar

Core Priority: This topic combines Azure-backed grounding, prompt/model configuration, and telemetry evidence. The scenario often describes weak answers after release and expects the candidate to separate retrieval quality, prompt/model behavior, and runtime monitoring.

High Frequency: Expect Azure AI Search with Microsoft Foundry for generative answers, custom prompts using the Foundry model catalog, and Application Insights monitoring through traces, requests, dependencies, and exceptions where available.

Confusion Alert: An Azure integration question may blame the model, but poor answers can originate from Azure AI Search freshness, Foundry prompt variables, or missing Application Insights correlation.

Scenario Logic: An enterprise agent uses indexed engineering content, a custom prompt backed by a selected model, and telemetry to diagnose poor answer quality after release. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Azure integration scope names Azure AI Search with Microsoft Foundry, Foundry model catalog prompts, and Application Insights monitoring, which makes retrieval, model choice, and telemetry separate exam dimensions.

Failure Trigger: Failure appears as stale Azure AI Search indexer status, non-retrievable grounding fields, missing Foundry prompt variable values, unsupported model selection, or no Application Insights dependency/exception correlation.

Operational Dependency: Azure integration has three observable planes: retrieval index health, model/prompt configuration, and telemetry. Diagnosing answer quality requires evidence across all three.

Why the Correct Answer Works: The correct option works because it separates three failure planes: Azure AI Search retrieval state, Microsoft Foundry prompt/model configuration, and Application Insights telemetry.

Best Choice Rules: If the stem mentions updated content not appearing, inspect Azure AI Search indexer history before prompt text. If the stem mentions custom prompt behavior, inspect required variables and

Microsoft Foundry model selection. If the stem mentions production diagnosis, choose Application Insights evidence over anecdotal transcript review alone.

Atomic Deconstruction - Operational Level

Azure integration gives the agent external evidence and observability. Azure AI Search contributes indexed retrieval, Microsoft Foundry model catalog prompts provide model-backed prompt execution, and Application Insights gives runtime evidence. The learner must avoid treating all quality problems as prompt problems.

Operationally, troubleshoot answer quality by separating evidence layers. First inspect search index state, then prompt variables and Foundry model selection, then telemetry correlation in Application Insights.

Changing prompts before checking index and telemetry evidence can hide the real failure.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | -----
-- | ----- | ----- |

| Azure AI Search index | Retrieval schema | Keyword/vector/hybrid fields where configured | Index unattached | Indexer and data source | Agent grounds on stale or irrelevant content |

| Foundry model catalog prompt | Model and prompt configuration | Supported model deployments/catalog selections | Default model behavior | Foundry access and prompt variables | Prompt cannot run or produces unsupported response shape |

| Application Insights | Telemetry sink | Traces, requests, dependencies, exceptions where available | Not connected | Instrumentation configuration | Failures cannot be correlated to tool or retrieval behavior |

| Correlation evidence | Conversation-to-operation trace | Session, operation, timestamp, dependency call | Fragmented | Telemetry fields and timestamps | Root cause cannot be distinguished from user phrasing |

Step-by-Step Execution Path

1. Inspect the Azure AI Search data source, index, and indexer status before changing prompts. If the index is stale or fields are not retrievable, prompt edits cannot recover missing evidence.
2. Review the custom prompt configuration and model selection in the Foundry-related workflow. Confirm required variables are supplied and the selected model supports the expected task.
3. Enable or inspect Application Insights monitoring for the agent path. Telemetry must capture enough request, dependency, trace, or exception evidence to correlate poor answers with retrieval or tool failures.
4. Replay one failing user question and compare retrieval evidence, prompt variables, model response, and telemetry timestamps. This sequence isolates whether the fault is data, prompt, model, or runtime dependency.

Evidence note: Azure checks rely on Azure portal index state, Microsoft Foundry prompt/model configuration, and Application Insights traces, requests, dependencies, or exceptions where available.

Technical Chain

A grounded answer begins when the agent sends the user query into the configured retrieval path. Azure AI Search returns documents according to index freshness, schema, filters, and retrievable fields. A custom prompt may then use a selected Microsoft Foundry model with variables supplied by the agent. Application Insights records runtime signals that can correlate the conversation with retrieval, dependencies, exceptions, or prompt/tool behavior. If telemetry is absent, the team cannot distinguish stale index content from prompt failure or dependency outage.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

|-----|-----|-----|

| Check index freshness | Azure portal > AI Search > Indexers > execution history | Recent indexer run succeeded and document count reflects expected source updates |

| Inspect retrievable fields | Azure portal > AI Search > Indexes > fields | Fields needed for answer grounding are searchable/retrievable and vector fields match embedding dimensions where used |

| Validate prompt variables | Copilot Studio custom prompt configuration with Foundry model selection | Required inputs are populated and model selection is supported in the environment |

| Query telemetry evidence | Application Insights > Logs > traces, requests, dependencies, exceptions | Timestamped records correlate the conversation with retrieval, tool, or exception behavior |

Practice Questions

1. An agent must answer from indexed engineering documents and filter results by product line and release. Which object should be validated first?
 - A. Computer use screen-state checkpoints
 - B. Adaptive card submit behavior
 - C. Azure AI Search index schema, retrievable fields, filters, and indexer status
 - D. Power Platform pipeline approval history
2. A legacy claims portal has no supported API, but an agent must perform a repetitive status update through the existing web UI. Which integration is the best fit?
 - A. Azure AI Search with a claims index
 - B. A custom connector without an API operation definition

- C. A Fabric data agent
 - D. Computer use with authentication planning and screen-state checkpoints
3. A team wants to expose functions from a tool server to a Copilot Studio agent. The tool does not appear in the agent's available tools. What should be verified first?
 - A. MCP server availability, authentication, tool discovery metadata, and exposed schema
 - B. Azure AI Search vector dimension
 - C. Whether the topic has enough synonyms
 - D. Whether the custom connector has a connection reference
 4. Several agents must call the same governed inventory API, and admins need DLP visibility and reusable authentication. Which approach is strongest?
 - A. Paste the endpoint into each topic's instructions
 - B. Create or use a custom connector with tested operations, connection references, and DLP-compatible policy
 - C. Store API responses in a static knowledge source
 - D. Use A2A protocol directly for the API
 5. A support agent should delegate analytics questions to a Fabric data agent but keep case-update actions local. What must be defined?
 - A. Only the final response format
 - B. One shared prompt for both agents
 - C. Handoff conditions, input/output contract, identity boundary, and Fabric workspace permissions
 - D. A second Azure AI Search index for case updates
 6. A Copilot Studio orchestrator integrates a Microsoft Foundry agent for specialized reasoning. Which validation is most important before release?
 - A. Whether the orchestrator has no local knowledge sources
 - B. Whether anonymous access works for the Foundry project
 - C. Whether the Foundry agent can be treated as an Azure AI Search index
 - D. Whether responsibility, authorization, input/output contract, and transcript evidence are defined
 7. A remote specialist agent will be invoked through A2A. Which evidence best proves readiness?
 - A. Valid agent metadata, trusted endpoint, authentication, task schema, and sample task response
 - B. A broader orchestrator greeting
 - C. A successful local flow unrelated to the remote agent
 - D. A Power Fx expression that formats the local response
 8. After a source content update, an agent still answers from old engineering documents. What should be checked before rewriting the prompt?
 - A. The channel authentication mode
 - B. Azure AI Search indexer history, document count, and indexed field availability
 - C. The REST action response schema
 - D. The A2A task response format

9. A production issue occurs only when an agent invokes a REST action after a grounded answer. Which monitoring evidence is most useful?
- A. The latest successful test transcript without dependency details
 - B. A list of topic trigger phrases
 - C. Application Insights traces, requests, dependencies, and exceptions correlated by timestamp
 - D. A connector inventory that does not include runtime status
10. A REST API action returns HTTP 403 when called from the agent, but the same endpoint works for an admin account. What should be inspected?
- A. The response message wording
 - B. The adaptive card action label
 - C. Whether the source should be indexed instead of called
 - D. The action's connection identity, permission scope, headers, and policy boundary
11. A computer use run fails after the target web app opens a modal that was not present during authoring. Which evidence should be reviewed?
- A. The recorded screen-state checkpoint and the first unexpected UI transition
 - B. Azure AI Search index statistics
 - C. The Fabric workspace role list
 - D. The deployment pipeline name
12. A custom connector operation succeeds in the maker test page but fails in the agent after deployment. What is the most likely object to validate?
- A. Whether the topic trigger condition changed during import
 - B. Connection reference binding and connector policy in the target environment
 - C. Whether the agent has a Fabric data agent
 - D. Whether A2A metadata is public
13. A prompt using the Microsoft Foundry model catalog returns incomplete output because a required customer segment value is missing. What should be checked first?
- A. Whether Azure AI Search returned enough documents
 - B. Whether the custom connector operation has a retry policy
 - C. Prompt variable mapping and the selected Foundry model configuration
 - D. Power Platform Pipelines deployment approval

Test and manage agents

Evaluate agent performance with test sets, evaluation methods, and result review

Exam Radar

Core Priority: This topic tests evidence-based quality control. The question usually asks how to prove answer quality, task completion, escalation behavior, or regression status before release.

High Frequency: Expect test sets, evaluation method selection, result review, transcript analysis, task success checks, failure categories, and regression comparison.

Confusion Alert: An evaluation question may describe a failed answer, but the trap is fixing the first visible response without classifying whether the owner is routing, grounding, tool execution, handoff, formatting, escalation, or deployment.

Scenario Logic: A team needs to prove that a production agent handles policy questions, tool actions, and escalation cases before publishing a new version. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: Testing scope is not just manual chat review: test sets, evaluation method selection, and result review are explicit management tasks.

Failure Trigger: Failure appears as unrepresented scenarios in the test set, no expected outcome, no failure category, no transcript-to-tool evidence, or no baseline to detect regression.

Operational Dependency: Agent evaluation must use representative scenarios and failure categorization because answer quality, grounding, tool completion, and escalation behavior fail in different ways.

How the Exam Asks It: A question may describe a business scenario, a failing Copilot Studio configuration, or a deployment constraint and ask for the best design, first troubleshooting step, or required component.

How Distractors Are Designed: Distractors often solve a visible symptom while ignoring the owning object. They may suggest prompt tuning for a permission problem, a connector for a retrieval problem, a channel change for a flow exception, or manual deployment for an ALM dependency.

Why the Correct Answer Works: The correct option works because it turns agent quality into repeatable evidence through test cases, evaluation methods, failed-result categories, and regression comparison.

Best Choice Rules: If the stem asks for readiness evidence, choose representative test sets and evaluation methods. If the stem mentions a failed answer, classify the failure before changing prompts. If the stem mentions release safety, compare against a baseline or previous evaluation run.

Atomic Deconstruction - Operational Level

An evaluation is a controlled measurement of agent behavior. The test set supplies representative questions and expected outcomes, the evaluation method decides how answers or tasks are judged, and result review turns failures into remediation categories. A single successful chat is not evidence of production readiness.

Operationally, evaluation begins with test-case design. Each case needs an expected outcome and an owning failure category. When a case fails, classify it as routing, grounding, tool execution, handoff, formatting, escalation, or deployment before selecting a fix.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | -----
----- | ----- |

| Test set | Scenario coverage | Questions, expected outputs, edge cases | Empty | Blueprint domains and production intents | Evaluation misses high-risk scenarios |

| Evaluation method | Scoring approach | Manual review, automated evaluation, task success checks | Ad hoc review | Expected answer and tool evidence | Results are subjective or not repeatable |

| Result category | Failure classification | Grounding, tool, topic routing, escalation, formatting | Unclassified | Transcript and telemetry | Team fixes the wrong layer |

| Regression baseline | Comparison point | Previous run, published version, expected threshold | No baseline | Versioned test set | Quality drift is not detected |

Step-by-Step Execution Path

1. Create a test set that covers each agent domain: planning assumptions, knowledge retrieval, tool execution, multi-agent handoff, and ALM-sensitive behavior. Coverage must match real user tasks, not only happy paths.
2. Select evaluation methods that fit each test case. Grounded answers need reference comparison, tool actions need completion evidence, and escalation needs route verification.
3. Run tests before publishing and record failed items with category, transcript, expected result, actual result, and likely owning object.
4. Review results by failure category before changing prompts. Topic routing, retrieval, tool schema, and identity failures require different fixes.

Evidence note: evaluation checks rely on test set coverage, evaluation-result details, transcripts, task-completion evidence, and comparison with a previous baseline.

Technical Chain

A test set triggers representative conversations. Each case produces a transcript, tool or flow evidence, and evaluation output. The reviewer compares expected and actual behavior, classifies the failure, and decides whether the owning object is topic routing, grounding, tool execution, handoff, formatting, escalation, or ALM. Without classification, teams change prompts for connector failures or rebuild tools for retrieval failures.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | ----- |
----- |

| Inspect test coverage | Copilot Studio evaluation/test set view | Test set includes normal, edge, negative, and escalation cases across exam domains |

| Review failed result | Evaluation result detail and conversation transcript | Failure is classified by route, grounding, tool, handoff, formatting, or escalation cause |

| Validate task completion | Transcript plus tool/flow run evidence | Expected external action completed or produced a handled failure response |

| Compare regression outcome | Current evaluation run vs previous run or baseline | No critical scenario regressed without an approved remediation plan |

Implement ALM for Copilot Studio agents with solutions, environment variables, and Power Platform Pipelines

Exam Radar

Core Priority: This topic is repeatable lifecycle management. AB-620 can ask how to move agents and dependencies across environments without hardcoded URLs, personal connections, or manual edits.

High Frequency: Expect solutions, adding existing agents to solutions, environment variables, connection references, Power Platform Pipelines, deployment history, and post-deployment smoke tests.

Confusion Alert: An ALM question may look like an import failure, but the clue often points to a hardcoded endpoint, missing environment-variable current value, unbound connection reference, or unmanaged dependency.

Scenario Logic: A developer must move an agent from development to test and production without hardcoding API URLs, connection values, or environment-specific settings. The best first move is the one that preserves identity, data boundary, execution contract, or evidence collection before optimizing the conversation wording.

Version Delta: ALM scope includes solutions, adding existing agents to solutions, environment variables, and Power Platform Pipelines, so deployment answers must be environment-aware.

Failure Trigger: Failure appears as a solution missing dependent flows or connectors, environment variables without target values, unbound connection references, failed pipeline stage, or a post-deploy smoke test calling a development endpoint.

Operational Dependency: Copilot Studio ALM depends on solution packaging, environment variables, connection references, and pipeline movement so environment-specific configuration is injected rather than hardcoded.

Why the Correct Answer Works: The correct option works because solutions, environment variables, connection references, and pipelines preserve dependencies while allowing target environments to supply their own values and connections.

Best Choice Rules: If the stem mentions different URLs, IDs, or settings per environment, choose environment variables. If the stem mentions connector authentication after import, inspect connection references. If the stem mentions controlled promotion, choose solutions and Power Platform Pipelines over manual export/import.

Atomic Deconstruction - Operational Level

ALM for Copilot Studio agents treats the agent, flows, connectors, environment variables, connection references, and dependent components as one deployable unit. Environment variables carry target-specific values; connection references bind connectors to target-environment connections; pipelines provide controlled movement and deployment evidence.

Operationally, ALM starts in the solution. The agent, flows, connectors, environment variables, and connection references must be included together. During deployment, target environments provide current values and connections, and pipeline history becomes the evidence that promotion followed the intended path.

Component Specifications

| Object | Attribute | Value Range | Default State | Dependency | Failure State |

| ----- | ----- | ----- | ----- | ----- | -----
----- | ----- |

| Solution | Package boundary | Agent, flows, connectors, variables, dependencies | Unmanaged local assets |
Power Platform environment | Missing dependency during import or publish |

| Environment variable | Configurable value | URL, IDs, feature flags, service endpoints | Unset in target |
Solution import and target environment values | Agent calls development endpoint in production |

| Connection reference | Connector binding | Per-environment connection | Maker connection | Connector
auth and DLP policy | Flow/tool cannot authenticate after deployment |

| Power Platform Pipeline | Deployment path | Dev, test, production stages | No pipeline | Managed
environments and permissions | Uncontrolled deployment and inconsistent versions |

Step-by-Step Execution Path

1. In Power Apps solution view, create or inspect the solution that contains the Copilot Studio agent and its dependent flows, custom connectors, connection references, and environment variables.
2. Replace hardcoded endpoints, IDs, and tenant-specific values with environment variables before deployment. This allows the same managed package to resolve different target values.
3. Validate connection references for every connector-backed flow or tool. The target environment must supply a valid connection that satisfies DLP and authentication requirements.
4. Use Power Platform Pipelines to move the solution through approved stages. Pipeline evidence gives a repeatable deployment record instead of a manual export/import trail.
5. After import, run a smoke test for topic routing, tool calls, knowledge access, and environment-specific endpoints before enabling the production channel.

Evidence note: ALM checks rely on solution contents, environment-variable current values, connection references, pipeline deployment history, and post-deploy smoke-test transcripts.

Technical Chain

A maker adds the agent and dependencies to a solution. During deployment, environment variables receive target-specific current values and connection references bind to valid target connections. Power Platform Pipelines moves the package through approved stages and records deployment history. If an endpoint or connection is hardcoded in a topic or flow, the imported agent may still call development resources or fail authentication in test or production.

Operational Skills Matrix

| Task | Precise Command or Path | Verification Standard |

| ----- | ----- | -----
----- |

| Inspect solution contents | Power Apps > Solutions > selected solution > Objects | Agent and all dependent flows, connectors, variables, and connection references are included |

| Verify environment variable values | Solution > Environment variables in target environment | Each required variable has a target-specific current value |

| Check connection references | Solution > Connection references | Every reference is bound to a valid connection in the target environment |

| Review pipeline deployment | Power Platform Pipelines deployment history | Deployment completed for intended stage with version and approver evidence |

| Run post-deploy smoke test | Copilot Studio test pane in target environment | Core route, tool call, and endpoint-specific response succeed or fail with handled diagnostics |

Practice Questions

1. Before publishing a new version, a team must prove policy answers, tool actions, and escalation paths still work. Which evaluation setup is best?
 - A. Review only the conversations that succeeded in the previous build
 - B. Run a small set of prompts without expected outcomes
 - C. Edit the system prompt and assume evaluation coverage is unchanged
 - D. A representative test set with expected outcomes, task evidence, and failure categories
2. Evaluation results show confident answers from the wrong knowledge source. Which category should be reviewed first?
 - A. Grounding, source selection, index, or retrieval configuration
 - B. Pipeline approval order
 - C. Adaptive card schema version
 - D. Environment variable current value

3. Choose two as one option: A tool action passes in development but calls the development endpoint after import to test. Which pair should have been configured?
 - A. Topic trigger condition and adaptive card schema
 - B. Environment variable for the endpoint and a target-environment current value
 - C. Static fallback answer and a broader generative instruction
 - D. Separate unmanaged agent copy and manual endpoint edit after import
 4. A solution import succeeds, but an agent flow cannot authenticate in the target environment. What should be inspected?
 - A. The topic trigger phrase list
 - B. Whether the agent has a published Teams channel
 - C. Connection references and target-environment connector connections
 - D. The evaluation test-set category names
 5. A team wants controlled promotion from development to test and production with deployment history. Which approach best fits?
 - A. Copy topics manually between environments
 - B. Edit production directly after every change
 - C. Export unmanaged files and fix values by hand
 - D. Use Power Platform Pipelines with solution packaging and approval history
 6. An evaluation run has many failures, but the team immediately edits the system prompt. What analysis step is missing?
 - A. Classifying failures by routing, grounding, tool execution, handoff, formatting, escalation, or deployment cause
 - B. Re-running only successful test cases after the prompt change
 - C. Deleting failed cases that use managed solution components
 - D. Reimporting the solution into the same environment without dependency review
 7. An agent and its dependent flows, custom connector, environment variables, and connection references must move together. What is the ALM boundary?
 - A. A transcript export containing successful validation chats
 - B. A Power Platform solution containing the agent and dependencies
 - C. A manually maintained list of topic names and flow URLs
 - D. A target environment with default connection values only
-

Learning Path & Study Advice

- Start with the Knowledge Overview so you can see the full exam scope and the exact order of the official domains, beginning with Plan and configure agent solutions, Integrate and extend agents in Copilot Studio, Test and manage agents.

- Read the Core Explanation in each knowledge point first to build a clean baseline understanding of the terminology, technologies, and customer scenarios.
 - Continue into the Advanced Explanation to deepen your understanding of design trade-offs, deployment planning, optimization options, and operational decision-making.
 - Work through the Practice Questions immediately after each knowledge point and answer them before checking the attachment section to strengthen retention.
 - Revisit the answer attachment to identify weak areas, then loop back into the corresponding knowledge-point section for targeted review.
-

Who This PDF Is For

This study pack is intended for learners preparing for the Designing and Building Integrated AI Solutions in Copilot Studio exam who want a structured, exam-aligned review resource. It is especially useful for professionals who need to connect the exam's knowledge points with practical responsibilities, business context, and operational decision-making.

It is also a good fit for self-paced learners who prefer to study from organized knowledge points, detailed explanations, and directly paired practice questions instead of jumping between multiple separate files.

Call To Action

This document provides an overview of structured learning and certification preparation approaches. For learners seeking clear knowledge organization, guided study planning, and exam-focused practice resources, AAAdemy offers a comprehensive platform to support independent and effective learning.

Explore additional training materials, study guidance, and practice resources at:

<https://www.aaademy.com/>

Attachment: Answers by Knowledge Point

Plan and configure agent solutions

Q1. Correct answer: A

Explanation:

A. Correct. External access plus dispute submission changes identity, channel, responsible AI, DLP, and component ownership decisions before authoring.

B. Wrong. Response styling is downstream and does not control external access or action risk.

- C. Wrong. A broad topic does not define the security or governance boundary for supplier actions.
- D. Wrong. Anonymous access removes the identity boundary needed for protected order and dispute data.

Q2. Correct answer: B

Explanation:

- A. Wrong. Manual copies drift and make ownership unclear.
- B. Correct. Solution-aware components and connection references support reuse, ownership, and environment movement.
- C. Wrong. A URL in response text is not a reusable ALM component.
- D. Wrong. Separate unmanaged environments increase drift and do not package dependencies.

Q3. Correct answer: C

Explanation:

- A. Wrong. Prompt warnings do not enforce reviewer decisions.
- B. Wrong. Public channel exposure weakens governance and does not model approval state.
- C. Correct. The approval state is the control object for high-risk refund actions.
- D. Wrong. Policy summarization does not execute or validate approval.

Q4. Correct answer: D

Explanation:

- A. Wrong. Publishing channel does not change the flow contract.
- B. Wrong. Generative answers cannot replace deterministic case creation.
- C. Wrong. Removing the connector removes the action rather than fixing the failure.
- D. Correct. The failure signal points to input mapping and connector error handling.

Q5. Correct answer: A

Explanation:

- A. Correct. A stale value in an API body is usually a variable scope or expression assignment issue.
- B. Wrong. Search index freshness does not control API request payload variables.
- C. Wrong. Channel policy does not decide which account ID is sent.
- D. Wrong. Import history may matter for deployment, not a runtime stale variable in a topic.

Q6. Correct answer: B

Explanation:

- A. Wrong. Republish may expose the card elsewhere but does not populate blank bindings.
- B. Correct. Blank card fields after a successful API call point to response mapping or variable binding.
- C. Wrong. Knowledge retrieval is not the source of these API response fields.
- D. Wrong. Approval owner matters for human review, not card field binding.

Q7. Correct answer: C

Explanation:

- A. Wrong. Calling the API before confirmation risks creating the wrong request.
- B. Wrong. Global instructions do not replace state collection or validation.

- C. Correct. The topic should ground, collect, validate, confirm, and then execute the action.
- D. Wrong. Rendering before validation can show or submit incorrect data.

Q8. Correct answer: D

Explanation:

- A. Wrong. Citations relate to grounded answers, not computed tool inputs.
- B. Wrong. Telemetry helps after runtime execution but does not validate the formula result first.
- C. Wrong. Source-variable population matters, but the asked validation is the expression result that will be passed to the tool.
- D. Correct. The expression result must match the tool's expected input contract.

Q9. Correct answer: A

Explanation:

- A. Correct. Sensitive HR scenarios require response boundaries and review evidence before release.
- B. Wrong. Trigger coverage does not control sensitive content handling.
- C. Wrong. Connectors may be needed later, but the review issue is responsible AI and escalation.
- D. Wrong. Card layout does not address sensitive-answer governance.

Q10. Correct answer: B

Explanation:

- A. Wrong. Creativity does not populate missing required variables.
- B. Correct. The topic must collect and validate required state before execution.
- C. Wrong. Environment movement does not fix a missing runtime variable.
- D. Wrong. A static FAQ cannot perform the required action.

Integrate and extend agents in Copilot Studio

Q1. Correct answer: C

Explanation:

- A. Wrong. Computer use is for UI automation, not document retrieval.
- B. Wrong. Card submit behavior does not control search fields or filters.
- C. Correct. Indexed retrieval depends on schema, filters, field retrievability, and freshness.
- D. Wrong. Pipeline history does not prove search relevance or index freshness.

Q2. Correct answer: D

Explanation:

- A. Wrong. Search retrieves content; it does not operate a transactional UI.
- B. Wrong. A custom connector requires a callable API contract; the scenario says no supported API exists.
- C. Wrong. A Fabric data agent supports data questions, not UI automation.
- D. Correct. Computer use matches UI-only action surfaces and needs screen checkpoints.

Q3. Correct answer: A

Explanation:

- A. Correct. Missing tool visibility points to MCP discovery, auth, or schema exposure.
- B. Wrong. Search vector settings do not expose MCP tools.
- C. Wrong. Synonyms may affect topic routing but not server tool discovery.
- D. Wrong. Connection references matter for connector execution, not MCP tool discovery.

Q4. Correct answer: B

Explanation:

- A. Wrong. Hardcoded endpoints are difficult to govern and reuse.
- B. Correct. Custom connectors provide reusable operations, policy visibility, and connection handling.
- C. Wrong. Static knowledge does not execute inventory updates or live queries.
- D. Wrong. A2A is for agent-to-agent task exchange, not direct API governance.

Q5. Correct answer: C

Explanation:

- A. Wrong. Formatting does not decide delegation ownership.
- B. Wrong. A shared prompt blurs specialist boundaries.
- C. Correct. Delegation requires a routing rule, contract, identity, and data-agent permission path.
- D. Wrong. Search indexing does not route analytics questions to Fabric.

Q6. Correct answer: D

Explanation:

- A. Wrong. Local knowledge may still be valid for local tasks.
- B. Wrong. Anonymous access is not an enterprise authorization strategy.
- C. Wrong. A Foundry agent is a delegated reasoning component, not a search index.
- D. Correct. Specialist delegation needs ownership, auth, contract, and evidence.

Q7. Correct answer: A

Explanation:

- A. Correct. A2A readiness depends on discoverability, endpoint trust, auth, and task exchange behavior.
- B. Wrong. Greeting text does not prove remote invocation.
- C. Wrong. An unrelated local flow does not validate A2A.
- D. Wrong. Local formatting does not prove remote task schema compatibility.

Q8. Correct answer: B

Explanation:

- A. Wrong. Channel authentication affects access to the agent, not whether new documents were indexed.
- B. Correct. Stale answers after a content update point to index freshness or schema.
- C. Wrong. REST response schema matters for action calls, not retrieval freshness.
- D. Wrong. A2A task response format matters for agent delegation, not search indexing.

Q9. Correct answer: C

Explanation:

- A. Wrong. A transcript without dependency details cannot isolate REST failures.
- B. Wrong. Trigger phrases do not explain downstream REST failures.

- C. Correct. Correlated telemetry distinguishes retrieval success from action dependency failure.
- D. Wrong. Inventory alone does not show runtime dependency failures.

Q10. Correct answer: D

Explanation:

- A. Wrong. Response wording cannot grant API permissions.
- B. Wrong. Card labels do not affect authorization.
- C. Wrong. Indexing changes retrieval design and does not fix an action authorization failure.
- D. Correct. 403 indicates an authorization or policy boundary for the action call.

Q11. Correct answer: A

Explanation:

- A. Correct. UI automation fails at screen-state changes; the first unexpected transition is the key signal.
- B. Wrong. Search statistics do not affect browser UI automation.
- C. Wrong. Fabric roles matter for data agents, not web UI state.
- D. Wrong. Pipeline names do not explain runtime screen transitions.

Q12. Correct answer: B

Explanation:

- A. Wrong. Topic routing can decide whether the action is reached, but the symptom points to connector execution after deployment.
- B. Correct. Deployment can break connection references or DLP compatibility even when maker testing worked elsewhere.
- C. Wrong. Fabric data agents are unrelated to this connector action.
- D. Wrong. A2A metadata does not control custom connector execution.

Q13. Correct answer: C

Explanation:

- A. Wrong. Retrieval volume may affect grounded content, but the symptom names a missing prompt variable.
- B. Wrong. Connector retry policy does not supply prompt inputs.
- C. Correct. Missing prompt variables and model configuration directly affect prompt execution.
- D. Wrong. Pipeline approval does not explain a missing runtime prompt value.

Test and manage agents

Q1. Correct answer: D

Explanation:

- A. Wrong. Successful prior conversations do not cover regression risk in the new version.
- B. Wrong. Prompts without expected outcomes cannot classify pass/fail behavior.
- C. Wrong. Prompt edits can introduce regressions and still require evaluation.
- D. Correct. Representative test sets provide repeatable quality and regression evidence.

Q2. Correct answer: A

Explanation:

- A. Correct. Wrong-source answers are a grounding/retrieval failure signal.
- B. Wrong. Pipeline approval may matter for deployment, not source selection.
- C. Wrong. Card schema affects rendering, not the source used for the answer.
- D. Wrong. Environment values can affect endpoints, but this symptom names answer grounding.

Q3. Correct answer: B

Explanation:

- A. Wrong. Trigger and card settings do not set target endpoint values.
- B. Correct. The variable definition and target current value keep endpoint values environment-specific.
- C. Wrong. Fallback wording hides the issue but does not set configuration.
- D. Wrong. Manual edits and unmanaged copies create drift.

Q4. Correct answer: C

Explanation:

- A. Wrong. Trigger phrases affect routing, not connector authentication.
- B. Wrong. Channel publication does not bind target-environment connector credentials.
- C. Correct. Connection references supply the target environment's connector identity.
- D. Wrong. Test-set category names do not control runtime authentication.

Q5. Correct answer: D

Explanation:

- A. Wrong. Manual copy causes drift.
- B. Wrong. Direct production edits bypass controlled promotion.
- C. Wrong. Hand-fixed unmanaged exports are not repeatable.
- D. Correct. Pipelines and solutions provide controlled movement and evidence.

Q6. Correct answer: A

Explanation:

- A. Correct. Failure classification identifies the owning layer before remediation.
- B. Wrong. Re-running only successful cases hides the failing behavior.
- C. Wrong. Deleting failed cases removes evidence instead of explaining it.
- D. Wrong. Reimporting without dependency review does not identify the failing layer.

Q7. Correct answer: B

Explanation:

- A. Wrong. Transcripts are evidence, not deployment packages.
- B. Correct. Solutions package the agent and dependent Power Platform components.
- C. Wrong. A manual list does not deploy or bind dependencies.
- D. Wrong. Default connection values do not define the package boundary.