ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

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

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

Why is it important to include examples in prompts?

A. Including examples should only focus on edge cases while ignoring typical scenarios
B. Examples should be omitted to allow AI to respond without guidance
C. Including examples guarantees perfect accuracy
D. Carefully chosen examples help guide the agent and improve generalization

Answer: D

Explanation:

Including well-designed examples supports few-shot learning by helping the model understand task structure, output style, and expected behavior. Examples enable pattern recognition and improve generalization across similar scenarios. Without examples, models must guess structure, reducing reliability. Examples enhance performance but still require evaluation and refinement.

Question: 2

Which persona typically models agentic processes in Maestro using BPMN and governs the full lifecycle?

A. Process operations teams and system administrators
B. Process excellence analysts
C. Automation developers in the Center of Excellence
D. Process owners in business teams

Answer: D

Explanation:

Process owners define and govern agentic workflows in UiPath Maestro. They use BPMN to model decisions, handoffs, escalation rules, and success criteria. Their business context enables ownership of the full lifecycle from design to optimization, supported by technical teams.

Question: 3

Which is an essential aspect of crafting a comprehensive agent story during the validation stage?

A. Brainstorming use cases without validating personas
B. Understanding daily pain points and inefficiencies of the selected role
C. Starting agent prototyping immediately
D. Generalizing automation opportunities across all roles

Answer: B

Explanation:

Agent stories must be grounded in real user pain points and inefficiencies. Understanding how a role spends time ensures automation targets real bottlenecks and delivers measurable value. Skipping persona validation leads to misaligned or low-impact agents.

Question: 4

An agent retrieves troubleshooting steps from the web, notifies users via Slack, and escalates unresolved tickets. Which evaluation strategy ensures comprehensive coverage without redundancy?

A. Random input sampling with LLM-as-a-Judge
B. Creating over 30 evaluations per tool
C. Grouping evaluations by success paths, failure paths, and edge cases
D. Equal evaluations for each individual tool

Answer: C

Explanation:

Grouping evaluations by real-world interaction patterns validates end-to-end logic. This includes success paths, fallback paths, and edge cases. It avoids redundant testing while ensuring functional coverage and traceable diagnosis of failures.

Question: 5

What is the primary purpose of a system prompt in a UiPath agent?

A. Define the agent's role, goal, constraints, and tool usage logic
B. Store enterprise context only
C. Control output formatting exclusively
D. Enumerate every possible dialog path

Answer: A

Explanation:

The system prompt provides persistent grounding by defining identity, goals, behavioral rules, escalation logic, and tool usage. It ensures consistent, safe, and goal-aligned agent behavior across dynamic scenarios.

Question: 6

What steps are required to create an evaluation set from scratch in UiPath?

A. Assign evaluators before defining inputs
B. Name the set, add inputs and expected outputs, save evaluations, then assign evaluators
C. Import JSON only from existing evaluations
D. Rely on automatic scoring only

Answer: B

Explanation:

Evaluation sets require defining inputs and expected outputs first, followed by evaluator assignment. This enables structured, repeatable validation of agent behavior before deployment.

Question: 7

Why do agents generate inconsistent output even when examples are provided?

A. Missing clear output-format instructions B. Agenda length constraints not defined C. Examples are too structured
D. Too much contextual grounding

Answer: A

Explanation:

Few-shot examples must be paired with explicit formatting instructions. Without clear structure, models infer formats inconsistently, breaking downstream automation reliability.

Question: 8

Which BPMN element maps to a decision point in agentic workflows?

A. Task
B. Swimlane
C. User Task
D. Exclusive Gateway

Answer: D

Explanation:

Exclusive Gateways represent branching decisions in BPMN. In agentic workflows, they allow agents to evaluate conditions and choose appropriate execution paths dynamically.

Question: 9

How does the “Number of results” setting affect context grounding?

A. Changes similarity threshold
B. Disables grounding
C. Controls number of retrieved context chunks
D. Selects Orchestrator folder

Answer: C

Explanation:

This setting determines how many context chunks are retrieved from the index. More results increase available context but also token usage, while fewer results improve focus and precision.

Question: 10

Why is a clear tool description required when adding tools to an agent?

A. To guide the agent on when and how to use the tool
B. To grant unrestricted system access
C. To prevent reasoning errors entirely
D. To automate escalation logic

Answer: A

Explanation:

Tool descriptions act as usage guidance for the LLM, helping it decide when to invoke a tool and how to supply inputs correctly. Poor descriptions lead to misuse or missed actions.

Question: 11

While configuring an Integration Service activity as a tool, how can an agent decide the value of a required field at runtime based on prompt instructions?

A. Set all fields as Arguments
B. Leave the field input method as Prompt

C. Define the field as an output argument
D. Convert the field into a variable

Answer: B

Explanation:

When a field is left as Prompt, the agent infers its value at runtime using instructions and context. This enables LLM autonomy and dynamic decision-making without hardcoded values or explicit argument passing.

Question: 12

What advanced feature does UiPath Context Grounding use to optimize retrieval from CSV files?

A. Structured SQL-style queries
B. Automatic conversion to XLSX
C. Real-time streaming ingestion
D. Semantic embedding of CSV data into JSON structures

Answer: D

Explanation:

CSV data is transformed into semantically rich JSON-style embeddings during indexing. This enables LLMs to retrieve relevant information using similarity search across structured tabular data.

Question: 13

Which limitation is commonly associated with traditional rule-based automation?

A. Too much flexibility across systems
B. Dependence on AI-powered agents
C. Inability to handle unstructured, judgment-based tasks
D. Poor performance for repetitive structured work

Answer: C

Explanation:

Traditional automation excels at structured tasks but struggles with unstructured inputs requiring reasoning or context. Agentic automation fills this gap using LLMs and contextual intelligence.

Question: 14

How long do key-value pairs stored in Agent Memory persist by default?

A. 3 months
B. 6 months
C. 12 months
D. Until agent version updates

Answer: C

Explanation:

Agent Memory retains key-value pairs for 12 months by default, enabling long-term context retention, preference tracking, and continuity across interactions unless manually cleared.

Question: 15

How do generative AI and agentic AI differ in a business workflow?

A. Generative AI creates content, while agentic AI automates decisions and actions
B. Generative AI handles sentiment, agentic AI writes content
C. Both only generate marketing material
D. Agentic AI creates content and generative AI executes tasks

Answer: A

Explanation:

Generative AI focuses on content creation, while agentic AI orchestrates decisions, workflows, and system interactions. Together they enable intelligent, end-to-end automation.

Question: 16

What type of agents can be invoked using “Start and wait for external agent” in UiPath Maestro?

A. UiPath robots only
B. Third-party SaaS agents
C. Agents configured within the same project
D. Agents without input or output parameters

Answer: C

Explanation:

This feature allows invocation of another agent within the same project, enabling agent-to-agent orchestration and modular workflow design with returned outputs.

Question: 17

What configuration options are available for setting up Context Grounding in UiPath?

A. Enable-only configuration without permissions
B. Indexes, folder permissions, LLM selection, and data syncing
C. Default settings only
D. Manual indexing without automation

Answer: B

Explanation:

UiPath allows full configuration of Context Grounding, including index creation, access control, LLM selection, and automated synchronization using dedicated activities.

Question: 18

What is the primary purpose of using Context Grounding with an ECS Index in Studio Web?

A. Store indexes in shared folders
B. Define static retrieval rules
C. Limit result count
D. Retrieve relevant data dynamically based on user input **Answer: D**

Explanation:

Context Grounding enables agents to dynamically retrieve relevant enterprise knowledge based on prompts or session input, improving accuracy and reducing hallucinations.

Question: 19

Why should the bpmn.uipath.com canvas primarily be used during early design stages?

A. It supports full execution and deployment
B. It functions as a sandbox without implementation capabilities
C. It imports templates into Studio Web
D. It replaces Automation Cloud

Answer: B

Explanation:

The BPMN canvas is a lightweight sandbox for modeling and collaboration. It supports ideation and discovery but lacks execution, tool integration, and runtime capabilities.

Question: 20

When should Web Search be used instead of Web Reader in an agent workflow?

A. For internal enterprise knowledge
B. For secure system data extraction
C. For summarizing information from multiple public sources
D. For parsing a known webpage

Answer: C

Explanation:

Web Search is ideal for open-ended discovery across public sources when no specific URL is known, whereas Web Reader targets content from a defined webpage.

Question: 21

When should Web Reader be preferred over Web Search in an agent workflow?

A. When summarizing information from multiple public sources
B. When extracting structured content from a known webpage
C. When searching news or public documentation
D. When grounding responses using ECS indexes

Answer: B

Explanation:

Web Reader is used when a specific URL is known and the agent needs to extract or summarize structured content from that page. It is not intended for open-ended discovery across multiple sources.

Question: 22

Why is the “as-is” process map important in identifying agentic automation opportunities?

A. It defines the final automated workflow
B. It highlights current inefficiencies and bottlenecks
C. It assigns agents to optimized roles
D. It eliminates the need for a “to-be” process

Answer: B

Explanation:

The as-is process map documents how work is currently performed. It helps uncover delays, rework, and manual steps that can be improved or automated using agents.

Question: 23

Which configuration area defines how an agent should proceed after a human resolves an escalation?

A. Assignment recipient list B. Agent Memory settings C. Input description fields D. Outcome behavior section **Answer: D**

Explanation:

The Outcome Behavior section controls the agent’s next actions after receiving a human decision, such as continuing automation, branching logic, or triggering follow-up steps.

Question: 24

Why is mapping processes critical during agentic discovery?

A. It prioritizes ROI metrics first
B. It avoids analyzing individual workflow steps
C. It identifies specific tasks suitable for automation
D. It completes automation without further design

Answer: C

Explanation:

Process mapping enables teams to identify granular steps and decision points where agents or RPA can add value, reducing errors and improving efficiency.

Question: 25

What is the main benefit of using an impact and feasibility matrix?

A. Automating all feasible processes
B. Focusing only on high-impact ideas
C. Balancing business value with implementation effort
D. Avoiding stakeholder involvement

Answer: C

Explanation:

The matrix evaluates automation opportunities based on both potential impact and feasibility, ensuring teams prioritize high-value, realistic use cases.

Question: 26

In which scenario is deterministic evaluation preferred over model-graded evaluation?

A. Evaluating tone or helpfulness
B. Scoring open-ended reasoning
C. Validating fixed and known outputs
D. Measuring user satisfaction

Answer: C

Explanation:

Deterministic evaluation is ideal when outputs are exact and predictable, such as status values, field extraction, or schema validation, allowing rule-based scoring.

Question: 27

Why should every input and output argument in an agent include a description?

A. To auto-map Orchestrator triggers
B. To help the agent understand how to use the arguments
C. To force arguments to be mandatory
D. Only inputs need descriptions

Answer: B

Explanation:

Argument descriptions provide grounding context for the agent, improving accuracy, structure, and correctness when generating or returning data.

Question: 28

What is a key step to ensure an agent knows when to use an escalation?

A. Configure outcome behavior only
B. Assign escalation recipients
C. Use required input fields
D. Add a clear prompt explaining escalation conditions

Answer: D

Explanation:

Escalation prompts guide the LLM’s judgment by clearly defining when escalation is appropriate, preventing overuse or missed critical cases.

Question: 29

What hybrid automation approach does UiPath recommend?

A. Agents for all workflows
B. RPA for unstructured tasks
C. Agents for structured tasks and RPA for decisions
D. RPA for structured tasks and agents for unstructured tasks

Answer: D
Explanation:

UiPath recommends using RPA for deterministic, repetitive tasks and agents for unstructured, judgment-based work, combining both strengths effectively.

Question: 30

Which statement best describes how agents operate in dynamic environments?

A. They fail when processes change
B. They adapt using contextual reasoning
C. They rely only on static rules
D. They always require human intervention

Answer: B
Explanation:

Agents adapt to changing inputs and conditions using LLM reasoning, context grounding, and prompt-driven logic, unlike rigid rule-based automations.

Question: 31

Which prerequisite must be met before an existing UiPath process becomes selectable as a tool for an agent?

A. The process must expose at least one string input argument
B. The process must be published and deployed to an accessible Orchestrator folder
C. Any published process appears automatically
D. The process must expose at least one string output argument

Answer: B
Explanation:

An agent can only invoke processes that are published and deployed to an Orchestrator folder it has permission to access. This ensures proper governance, security, and execution visibility.

Question: 32

Why is an agent story important during the development lifecycle?

A. It highlights weaknesses in poorly designed agents
B. It is only needed for stakeholder demos
C. It helps developers focus on value-driving features
D. It replaces technical documentation

Answer: C
Explanation:

An agent story provides a clear narrative of what the agent does, for whom, and why. It keeps development aligned with business value and prevents scope creep.

Question: 33

What is the primary purpose of using Business Process Model and Notation (BPMN) for agentic workflows?

A. Design aesthetic diagrams only
B. Model static workflows without exceptions
C. Enable standards-based collaboration between business and IT
D. Avoid error handling in automation

Answer: C

Explanation:

BPMN provides a standardized, visual language that allows business users and technical teams to collaboratively design, govern, and manage dynamic agentic workflows.

Question: 34

Why is providing a focused description important when adding a tool to an agent?

A. It restricts the agent's reasoning
B. It guarantees automatic execution
C. It helps the agent understand when and how to use the tool
D. It replaces system prompts

Answer: C

Explanation:

Tool descriptions act as guidance for the agent, enabling it to correctly decide when to invoke a tool and how to use its inputs and outputs effectively.

Question: 35

What is the main benefit of grouping evaluations into functional sets?

A. Increasing evaluation volume
B. Testing tools independently only
C. Validating end-to-end agent behavior
D. Reducing the need for edge cases

Answer: C

Explanation:

Functional grouping validates how tools and logic work together across success paths, failure paths, and edge cases, providing comprehensive coverage without redundancy.

Question: 36

What is the role of Context Grounding in reducing hallucinations?

A. It disables LLM reasoning
B. It limits the number of prompts
C. It injects relevant enterprise context into the prompt
D. It replaces system prompts

Answer: C
Explanation:

Context Grounding supplies relevant indexed enterprise knowledge at runtime, anchoring LLM responses in real data and significantly reducing hallucinations.

Question: 37

Which feature allows agents to retain information across multiple runs?

A. Outcome behavior
B. Agent Memory
C. Tool arguments
D. Escalation prompts

Answer: B
Explanation:
Agent Memory provides persistent storage for key-value pairs, enabling context continuity, personalization, and stateful behavior across executions.

Question: 38

Why are exclusive gateways important in agentic BPMN workflows?

A. They group activities by role
B. They represent human approval tasks
C. They enable conditional branching and decisions
D. They store contextual data

Answer: C
Explanation:

Exclusive gateways define decision points where only one path is selected based on conditions, allowing agents to dynamically route execution.

Question: 39

What is a key advantage of agentic automation over traditional RPA?

A. Lower infrastructure requirements
B. Better handling of unstructured and judgment-based tasks
C. Elimination of governance controls
D. Faster execution of structured tasks

Answer: B
Explanation:

Agentic automation uses LLM reasoning and context to handle ambiguity and unstructured inputs, areas where traditional RPA struggles.

Question: 40

Why should prompts include both instructions and examples?

A. Instructions alone guarantee accuracy
B. Examples alone define behavior
C. Combined use improves consistency and reliability
D. They are only needed for testing

Answer: C

Explanation:

Combining clear instructions with examples provides structure and guidance, enabling consistent, reliable outputs and better generalization across scenarios.

Question: 41

What is the primary purpose of using Agent Memory in UiPath agents?

A. To store execution logs only
B. To persist context and decisions across multiple runs
C. To replace Context Grounding indexes
D. To control escalation routing

Answer: B

Explanation:

Agent Memory enables persistent storage of key-value data, allowing agents to recall preferences, prior decisions, and contextual information across executions for continuity and personalization.

Question: 42

When designing an agent to generate step-by-step troubleshooting guides, what prompt refinement improves clarity and usability?

A. Allow longer responses for completeness
B. Add random formatting examples
C. Provide clear instructions for concise, actionable steps
D. Avoid detailed explanations

Answer: C

Explanation:

Explicit instructions focusing on actionable, concise, and non-redundant steps significantly improve clarity and usability in generated troubleshooting guides.

Question: 43

What distinguishes deterministic evaluations from model-graded evaluations?

A. They rely on subjective judgment
B. They evaluate tone and sentiment
C. They validate exact, known outputs using rules
D. They require human feedback only

Answer: C

Explanation:

Deterministic evaluations are rule-based and ideal when outputs are fixed and predictable, such as exact matches, schema validation, or structured field extraction.

Question: 44

Why is Context Grounding considered a core enterprise feature in UiPath?

A. It eliminates the need for prompts
B. It enables governed, scalable access to enterprise knowledge
C. It automatically retrains LLMs
D. It replaces Orchestrator permissions

Answer: B

Explanation:

Context Grounding allows agents to securely access enterprise knowledge through governed indexes, improving accuracy, relevance, and safety at scale.

Question: 45

What is the main role of escalation prompts in agent design?

A. To define post-escalation actions
B. To select escalation recipients
C. To guide the agent on when escalation is appropriate
D. To store escalation outcomes

Answer: C

Explanation:

Escalation prompts guide the agent's judgment by defining conditions that warrant human intervention, ensuring escalations are used correctly and consistently.

Question: 46

Which practice best supports accurate intent classification using few-shot prompting?

A. Providing random unlabeled examples
B. Using structured input-output labeled examples
C. Leaving labels blank for inference
D. Using ranked category lists only

Answer: B

Explanation:

Clearly labeled input-output examples provide strong grounding, helping the model learn classification patterns and improving accuracy for new, unseen inputs.

Question: 47

What is the benefit of using outcome behavior after human escalation?

A. It assigns escalation tasks
B. It determines how the agent proceeds after resolution
C. It stores memory values
D. It configures context grounding

Answer: B
Explanation:

Outcome behavior defines the agent's next steps after a human decision, enabling intelligent continuation, branching, or follow-up actions in hybrid workflows.

Question: 48

Which automation component excels at handling repetitive, deterministic tasks?

A. Agentic AI
B. Generative AI
C. RPA robots
D. Context indexes

Answer: C
Explanation:

RPA robots are optimized for structured, repetitive tasks where speed, reliability, and deterministic execution are required.

Question: 49

Why should an agent not rely solely on examples without instructions?

A. Examples reduce creativity
B. Examples guarantee incorrect outputs
C. Lack of instructions causes inconsistent interpretation
D. Instructions override examples

Answer: C
Explanation:

Without explicit instructions, agents may interpret examples inconsistently, leading to format variation and unreliable outputs, especially in structured automation tasks.

Question: 50

How does agentic automation improve handling of complex workflows?

A. By enforcing static rules
B. By removing human involvement
C. By combining reasoning, context, and orchestration
D. By replacing all RPA bots

Answer: C
Explanation:

Agentic automation integrates LLM reasoning, contextual grounding, and orchestration to manage variability, ambiguity, and cross-system workflows effectively.

Question: 51

When should an agent escalate a task to a human reviewer?

A. Whenever a tool invocation fails
B. When confidence is low or business-defined conditions are met
C. Only after all tools are executed
D. Whenever Agent Memory is empty

Answer: B

Explanation:

Agents should escalate based on clearly defined business rules or low-confidence scenarios. Proper escalation ensures accuracy and safety while maintaining human-in-the-loop governance.

Question: 52

What is the purpose of folder-level permissions in Context Grounding?

A. To improve LLM creativity
B. To control which agents can access specific indexes
C. To define escalation recipients
D. To limit token usage

Answer: B

Explanation:

Folder-level permissions govern which agents or users can access specific Context Grounding indexes, ensuring secure and compliant access to enterprise data.

Question: 53

Why is business-led automation emphasized in agentic workflows?

A. It removes the need for developers
B. It ensures processes align with real business context
C. It simplifies BPMN diagrams
D. It eliminates governance

Answer: B

Explanation:

Business-led automation ensures agents are designed around real operational needs, enabling accurate decision-making and value-driven outcomes supported by technical implementation.

Question: 54

What advantage does agent-to-agent orchestration provide?

A. Reduced prompt size
B. Modular and reusable agent design
C. Elimination of escalation logic
D. Automatic model retraining

Answer: B

Explanation:

Agent-to-agent orchestration allows complex workflows to be broken into specialized, reusable agents, improving scalability and maintainability.

Question: 55

Why is consistent prompt structure important in agent design?

A. It limits agent adaptability
B. It improves determinism and reliability
C. It replaces evaluations
D. It removes the need for examples

Answer: B

Explanation:

Consistent prompt structure improves reliability by reducing ambiguity, ensuring predictable behavior across executions and easier evaluation of outputs.

Question: 56

What is the main role of evaluators in UiPath evaluation sets?

A. To execute tools
B. To score and validate agent outputs
C. To design BPMN flows
D. To store context data

Answer: B

Explanation:

Evaluators assess agent outputs against expected results, enabling qualitative and quantitative validation before deployment.

Question: 57

Why is grounding payload size important to tune?

A. It determines Orchestrator folder selection
B. It affects token usage and relevance of context
C. It controls escalation routing
D. It defines memory retention duration

Answer: B

Explanation:

Tuning grounding payload size balances context relevance with token efficiency, preventing noise while ensuring sufficient information for accurate responses.

Question: 58

Which element supports conditional branching in agent workflows?

A. Swimlane
B. Task
C. Exclusive Gateway
D. Message Event

Answer: C

Explanation:

Exclusive gateways evaluate conditions and route execution down a single appropriate path, enabling dynamic decision-making in agentic processes.

Question: 59

What is a key benefit of using structured outputs in agent responses?

A. Improved creativity
B. Easier downstream automation and validation
C. Reduced prompt size
D. Elimination of context grounding

Answer: B

Explanation:

Structured outputs enable reliable downstream processing, validation, and integration with other automation components.

Question: 60

Why is testing edge cases critical in agent evaluations?

A. They are rarely encountered
B. They reduce evaluation coverage
C. They reveal failure modes and robustness gaps
D. They eliminate the need for success-path tests

Answer: C

Explanation:

Edge cases expose how agents behave under unusual or ambiguous conditions, helping identify weaknesses and improve robustness before production deployment.

Question: 61

What is the main purpose of assigning clear roles in an agentic workflow?

A. To reduce the number of BPMN elements
B. To clarify responsibilities between agents, robots, and humans
C. To eliminate the need for governance
D. To improve diagram aesthetics

Answer: B

Explanation:

Clear role definition ensures that agents, RPA robots, and humans each handle the tasks they are best suited for, improving accountability, efficiency, and governance.

Question: 62

Why is it important to define success criteria for an agent?

A. To restrict agent creativity

B. To enable measurable evaluation of agent performance
C. To avoid using Context Grounding
D. To remove escalation logic

Answer: B

Explanation:

Success criteria provide measurable outcomes that allow teams to evaluate whether an agent is delivering intended business value and behaving as expected.

Question: 63

What benefit does using few-shot examples provide when transforming unstructured text into structured data?

A. It increases randomness
B. It helps the model infer the desired structure
C. It replaces the need for instructions
D. It reduces token usage

Answer: B

Explanation:

Few-shot examples demonstrate the desired input-to-output mapping, helping the model learn how to consistently structure unstructured information.

Question: 64

Which factor most influences an agent's decision to invoke a tool?

A. Tool execution speed
B. System prompt and tool description
C. Agent Memory expiration
D. Orchestrator folder name

Answer: B

Explanation:

The system prompt and tool description guide the LLM's reasoning on when a tool is appropriate, shaping correct and consistent tool invocation behavior.

Question: 65

Why is separating system prompts from user prompts considered best practice?

A. It reduces total prompt length
B. It allows persistent behavioral guidance separate from user input
C. It eliminates the need for examples
D. It improves UI performance

Answer: B

Explanation:

System prompts provide stable, persistent instructions governing agent behavior, while user prompts vary per interaction. Separation ensures consistent behavior across sessions.

Question: 66

What is the purpose of defining agent boundaries in BPMN models?

A. To merge multiple processes
B. To identify where agent responsibility begins and ends
C. To replace swimlanes
D. To control memory retention

Answer: B

Explanation:

Agent boundaries clarify which steps are handled by agents versus humans or robots, ensuring proper orchestration and governance.

Question: 67

Why is governance important in enterprise agentic automation?

A. It slows down development
B. It ensures compliance, security, and controlled behavior
C. It replaces evaluations
D. It limits automation potential

Answer: B

Explanation:

Governance ensures agents operate safely, securely, and in compliance with organizational policies, especially when accessing enterprise systems and data.

Question: 68

What advantage does structured evaluation provide over ad-hoc testing?

A. Lower test coverage
B. Repeatable and measurable validation
C. Faster execution only
D. Elimination of human review

Answer: B

Explanation:

Structured evaluations allow consistent, repeatable assessment of agent behavior, making it easier to detect regressions and validate improvements.

Question: 69

Why should escalation criteria be aligned with business rules?

A. To reduce system prompts
B. To ensure escalations occur only when necessary
C. To automate all decisions
D. To avoid human involvement

Answer: B
Explanation:

Aligning escalation criteria with business rules ensures agents escalate appropriately, balancing automation efficiency with risk management.

Question: 70

How does Context Grounding support decision-making in agents?

A. By disabling LLM reasoning
B. By injecting relevant, real-time enterprise knowledge
C. By replacing system prompts
D. By storing execution history

Answer: B
Explanation:

Context Grounding enriches agent prompts with relevant enterprise knowledge, enabling informed decisions and reducing hallucinations.