FRAMEWORK WHITE PAPER · VERSION 1.0

Agent Accountability Documentation Template

Intent Architecture Stack · All Three Layers

Sougata Roy | sougataroy.com | May 2026

i This template covers all three layers of the Intent Architecture Stack: Context (the environment), Intent (the purpose and scope), and Governance (the accountability structure). Complete one template per agent. Tier 1 agents may use class-level documentation for Section 1 and a simplified Section 2. Tier 3 agents require full completion of all three sections. Fields marked * are mandatory for all tiers.

How to use this template

This template is an organizational record, not a technical specification. **It is written before the agent enters production.** The single test for each section: could a new compliance officer, examiner, or board member read this document and immediately understand the agent's regulatory environment, its organizational purpose, and who is accountable for its actions, without asking the deployment team a single question? If not, the document is incomplete.

Complete in order: Context first, Intent second, Governance third. Intent cannot be written responsibly without Context. Governance without Intent means accountability for a boundary no one has defined.

SECTION 1 · LAYER 1

Context Document

Map the regulatory environment, stakeholder relationships, and system integrations before writing intent. This document must be dated before the Intent Document.

Agent Name *

Agent Display Name *	*The name as it appears in Agent 365 and Entra Agent ID*
Entra Agent ID	*The unique identifier assigned at registration*
Tier Classification *	*Tier 1 / Tier 2 / Tier 3 (see risk model)*
Context Document Version	*e.g. v1.0*
Document Date *	*Must predate the Intent Document*
Prepared By *	*Name, title, team*

Regulatory Environment * *(list all applicable frameworks)*

i Name every framework that applies to this agent's deployment environment. If the regulatory environment changes after deployment, update this document before the next agent review cycle.

Framework / Regulation	Applicability	Specific Obligation for This Agent

Stakeholder Groups and Data Touchpoints *

i Identify every group whose data or workflows this agent will touch, directly or indirectly. Include counterparties, customers, employees, and third-party systems that process data the agent generates.

Stakeholder Group	Data Type Accessed	Sensitivity Level	Consent / Notice Requirement

Stakeholder Group	Data Type Accessed	Sensitivity Level	Consent / Notice Requirement

System Integrations and Downstream Triggers *

i Map every downstream system this agent can trigger or access. If the agent can invoke an API, trigger a workflow, read a mailbox, or write to a connected system, that connection must appear here. An unmapped integration is an unauthorized action waiting to happen.

System / Service	Integration Type	Data Flow Direction	Access Scope	DLP / Label Applied

Context Completeness Test

Could a new security architect use this document to understand the full regulatory, stakeholder, and technical landscape this agent operates within, without asking the deployment team a single question? Answer yes or no, and if no, identify what is missing.

[Yes / No. If no: list what is missing and by what date it will be completed.]

SECTION 2 · LAYER 2

Intent Document

Document what the agent is supposed to accomplish and the explicit boundaries within which it must operate. This document is the governance standard against which the agent's production behavior is measured.

Agent Name *	*Must match Section 1 exactly*
Intent Document Version	*e.g. v1.0*
Document Date *	*Must postdate the Context Document*
Prepared By *	*Name, title, team*
Reviewed By *	*Name, title*

Purpose Statement *

i Describe what the organization intends this agent to accomplish, in plain language. This is not a capabilities description and not a technical specification. A compliance officer should be able to read this and evaluate whether a specific agent action is within the organization's intent. If the Purpose Statement could have been written by reading the configuration file, it is not an intent statement.

[Describe the organizational purpose of this agent in 2-4 sentences. What problem does it solve? Who does it serve? What does it produce? Example: "This agent reviews inbound customer service tickets and drafts initial responses for human review. It serves the customer operations team. It produces draft responses that must be approved by a named customer service representative before transmission."]

Authorized Scope — Permitted Actions *

i List the specific actions this agent is authorized to take. Be precise. "Summarize documents" is not a sufficient permission; "Summarize SharePoint documents in the Customer Operations library" is.

Permitted Action	Scope / Constraints	Systems Authorized

Permitted Action	Scope / Constraints	Systems Authorized

Authorized Scope — Explicit Prohibitions *

i This field is mandatory and must not be left blank. Explicit prohibitions define the boundary from the outside. An authorization document without explicit prohibitions cannot answer whether a specific agent action was in scope or out of scope. That determination is the one an examiner, a board member, or a court will ask for.

Prohibited Action	Rationale	Enforcement Mechanism

Expected Outputs *

i Define what correct behavior looks like. If the agent produces an output that falls outside this range, that determination should be possible without requiring the original deployment team to be present.

[Describe the expected format, content, and recipients of the agent's outputs. Example: "Draft email responses in English, under 300 words, addressed to the customer ticket submitter. Responses must not contain pricing information, contract terms, or legal commitments."]

Human Review Triggers *

i Define the conditions under which the agent must pause and route to a named human for decision. These triggers are the primary monitoring standard for this agent's production behavior.

Trigger Condition	Routing Target	Response Time Requirement

Trigger Condition	Routing Target	Response Time Requirement

Intent Completeness Test

Could a new compliance officer use this document, on their first day, to evaluate whether a specific agent action was within the organization's intent? Could they determine, without investigation, whether the action was permitted or prohibited? Answer yes or no.

[Yes / No. If no: list what is ambiguous and by what date it will be resolved.]

SECTION 3 · LAYER 3

Governance Document

Assign the human accountability structure that owns the consequence when this agent acts. A platform sponsor field is not this document. This document gives that field its organizational meaning.

Agent Name *	*Must match Sections 1 and 2 exactly*
Governance Document Version	*e.g. v1.0*
Document Date *	*Date accountability was formally assigned*

Business Sponsor *

i The named individual who has authorized this agent's deployment and accepted organizational accountability for its purpose and lifecycle decisions. This is the name that goes in the Entra Agent ID sponsor field. The Business Sponsor must be an individual, not a team or a role title.

Full Name *	*[Name]*
Title *	*[Title]*
Reporting Line *	*[Reports to: Name, Title]*
Entra Agent ID Sponsor Field *	*[Confirm populated: Yes / No]*

Consequence Owner *

i The named individual accountable for board-level explanation and regulatory inquiry response when this agent takes a consequential action. For Tier 1 agents, the Consequence Owner and Business Sponsor may be the same person. For Tier 3 agents with external system access or customer-facing decisions, they should be different individuals at different levels of the accountability structure. The Consequence Owner is the answer to the board's question.

Full Name *	*[Name]*
Title *	*[Title]*
Reporting Line *	*[Reports to: Name, Title]*
Same as Business Sponsor?	*Yes / No*
Nested Accountability Structure *	*[Three-lines-of-defense / SMF / Risk Committee / equivalent]*

Escalation Path *

i The named individuals, in sequence, who are notified when this agent triggers an anomaly. This path is written before the first incident, not assembled during it. Mandatory for every Tier 3 agent. Required for Tier 2 agents operating in regulated environments.

Level	Name / Title	Contact	Trigger for This Level	Response Time

Review Cadence *

i Define when this document will be reviewed. The quarterly review is triggered by Microsoft's Agent 365 and Copilot capability releases. When Microsoft expands what this agent can do between deployments, this document must be reviewed before the expanded capability is used in production. Additional triggers are listed below.

Review Type	Frequency	Next Scheduled Date	Trigger Events

Governance Completeness Test

Could a board member, an OCC examiner, or an external auditor identify the Consequence Owner and the Escalation Path from this document, without investigation, in under two minutes? Answer yes or no.

[Yes / No. If no: list what is ambiguous and by what date it will be resolved.]

Accountability Assignment

The individuals below confirm they have read and accepted the accountability obligations described in this document, as of the dates signed.

Business Sponsor — Accepted by **Date**

Name / Title

Consequence Owner — Accepted by **Date**

Name / Title

Compliance / Risk Review — Reviewed by **Date**

Name / Title

Document Change Log

Version	Date	Changed By	Summary of Changes	Consequence Owner Sign-off

i This template is version 1.0, May 2026. Download the latest version at sougataroy.com/frameworks. Associated frameworks: Governance Readiness Matrix, Deployment Accountability Map, Tenant Agent Reconciliation Framework, Authorization Coverage Lifecycle. Free to use with attribution. Cite as: Roy, S. (2026). Intent Architecture Stack Documentation Template. sougataroy.com/frameworks.