

AI Agent Architecture Blueprint

A governance-aligned operating design for safe, scalable AI agents

THBS helps organizations move from AI ideas to a controlled operating model. The blueprint defines which agents the business needs, how they connect to workflows and systems, what data they may access, when people must intervene, and how performance and risk will be managed.

ARCHITECTURE AT A GLANCE

1 BUSINESS Priorities and use cases	2 AGENTS Roles and authority	3 DATA Knowledge and access	4 CONTROL Human review and safeguards	5 OPERATE Monitor and improve
---	--	---------------------------------------	---	---

WHAT THE BLUEPRINT DEFINES

Agent Portfolio Agent purposes, roles, owners, responsibilities, and prohibited actions.	Human Oversight Approval gates, decision authority, escalation paths, and accountable owners.	Data and Access Approved sources, minimum-necessary permissions, retention, and sharing limits.
Workflow and Integration Triggers, handoffs, business-system connections, exceptions, and recovery.	Testing and Controls Accuracy, privacy, security, bias, prompt-injection, and failure testing.	Lifecycle Management Agent registry, monitoring, incidents, version control, review, and retirement.

CLIENT DELIVERABLES

• Architecture: future-state agent and integration map • Design briefs: purpose, boundaries, data, and authority for each agent • Responsibility model: human ownership, approvals, and escalation	• Control matrices: access, risk, testing, and release requirements • Implementation roadmap: prioritized pilot, dependencies, and deployment gates • Operating toolkit: agent registry, monitoring, incident, and review templates
---	--

THE THBS APPROACH



Your AI agents should operate with defined purpose, limited authority, trusted data, human accountability, and measurable controls.

Schedule an AI Architecture Discovery Consultation

ThresHold Business Solutions | San Antonio, Texas 78254 | 📞 726-203-4084

Bizstrategist@Thresholdbusinesssolutions.com | www.thresholdbusinesssolutions.com