

The Nine Skills Framework

A Complete Guide to Building AI Systems That Don't Fall Apart

Agentic AI Strategy for Builders, Strategists, and Executives

Terry Byrd

byrddynasty.com

Enhanced Framework: Essential Skills for Agentic AI Strategists in 2026

Author: Manus AI

Date: December 25, 2025

Version: 2.0 (Enhanced)

Executive Summary

The landscape of agentic artificial intelligence has evolved dramatically since the early frameworks emerged in 2024. Based on comprehensive parallel research across fifteen autonomous workflow frameworks and emerging standards, this document presents an enhanced skillset framework for the 2026 Agentic AI Strategist. While the original eight-skill framework provided an excellent foundation, this enhanced version incorporates critical insights from frameworks such as **Pydantic AI**, **LlamaIndex**, **Haystack**, **Semantic Kernel**, and others that have reshaped the field.

The fundamental shift proposed in this enhancement is a move from **framework-specific knowledge** to **principle-based competencies**. The rapid pace of innovation in agentic AI means that specific tools and frameworks will continue to evolve, merge, or become obsolete. However, the underlying principles of state management, observability, data-centric design, and security will remain constant. This enhanced framework therefore emphasizes transferable knowledge that will remain relevant regardless of which specific tools dominate the market in 2026 and beyond.

Key Enhancements

This enhanced framework proposes **nine core skills** organized into **three pillars**, representing a restructuring that better reflects the current state of the art:

Pillar I: Autonomous System Architecture (Skills 1-3) - Skill 1: Multi-Agent Orchestration and State Management Principles - Skill 2: Interoperability and Integration Engineering - Skill 3: Production-Grade Observability and MLOps for Agents

Pillar II: Data-Centric AI Engineering (Skills 4-6) - Skill 4: Hybrid Memory Architectures and Knowledge Engineering - Skill 5: Context Economics and Optimization - Skill 6: Data Quality, Governance, and Grounding

Pillar III: Security, Governance, and Capability Engineering (Skills 7-9) - Skill 7: Non-Human Identity and Access Management - Skill 8: Semantic Capability and Tool Engineering - Skill 9: Agentic Security and Adversarial Resilience

Part I: Autonomous System Architecture

The first pillar focuses on the fundamental architectural patterns and operational practices required to build, deploy, and maintain autonomous multi-agent systems at scale. This pillar emphasizes the control plane and the mechanisms that govern how agents exist, collaborate, and self-correct.

Skill 1: Multi-Agent Orchestration and State Management Principles

Original Skill: "Advanced Multi-Agent System (MAS) Topologies and Design Patterns" and "Orchestration Framework Proficiency and Hybridization"

Why This Change: The original framework split orchestration into two skills: one focused on topologies and another on frameworks. Research into modern frameworks like **Pydantic AI**, **LlamaIndex AgentWorkflow**, and **Semantic Kernel** reveals that the critical competency is not mastery of specific tools, but rather a deep understanding of the underlying principles of state management, control flow, and inter-agent communication. By 2026, the most successful strategists will be those who can design robust orchestration patterns regardless of the specific framework in use.

Enhanced Description:

The 2026 Agentic AI Strategist must possess a principle-based understanding of how to orchestrate complex, multi-step workflows involving multiple specialized agents. This skill encompasses three core competencies:

1.1 State Management Architectures

The foundation of any reliable agentic system is robust state management. Unlike traditional software where state is deterministic, agentic systems must manage probabilistic state transitions where the next state depends on LLM outputs that may vary between runs. The strategist must understand and implement multiple state management paradigms:

Stateful Graph Architectures: Systems like **LangGraph** pioneered the concept of representing agent workflows as explicit state machines, where each node represents a computational step and edges represent state transitions. This approach provides determinism and auditability, making it ideal for regulated industries. The strategist must understand how to design state schemas, implement conditional edges, and leverage checkpointing for fault tolerance.

Event-Driven State Management: Frameworks like **LlamaIndex AgentWorkflow** and enterprise event streaming platforms demonstrate the power of event-driven architectures. In this paradigm, state changes are represented as immutable events in a log (e.g., Kafka, Redis Streams), allowing for temporal replay, distributed processing, and natural integration with existing enterprise event systems. The strategist must design event schemas, implement event handlers, and manage eventual consistency.

Context-Based State: Emerging frameworks like **Pydantic AI** emphasize the use of dependency injection and context objects to manage state. This approach treats state as a first-class object that is explicitly passed between agent components, enhancing testability and modularity. The strategist must understand how to design context schemas, implement dependency injection patterns, and manage context lifecycle.

1.2 Control Flow Patterns and Orchestration

Beyond state management, the strategist must master the control flow patterns that govern how agents collaborate:

Sequential Pipelines: The simplest pattern, where agents execute in a fixed order, passing outputs to the next agent. This pattern is suitable for well-defined, linear workflows like document processing pipelines.

Parallel Execution and Fan-Out/Fan-In: For tasks that can be decomposed into independent subtasks, parallel execution dramatically reduces latency. The strategist must design systems that can dynamically spawn worker agents, aggregate their results, and handle partial failures gracefully.

Hierarchical Delegation: Complex tasks require hierarchical organization, where high-level "manager" agents decompose goals and delegate to specialized "worker" agents. This mirrors human organizational structures and is essential for scaling to enterprise-level complexity. The strategist must design clear delegation protocols, manage context propagation across hierarchy levels, and implement escalation mechanisms for failures.

Dynamic and Adaptive Topologies: The most advanced systems in 2026 do not use fixed topologies. Instead, they dynamically adapt their collaboration patterns based on task characteristics. For example, a system might use a hierarchical pattern for initial planning, switch to a network (mesh) pattern for creative brainstorming, and then use a supervisor pattern for execution. This requires the strategist to design meta-orchestration logic that can reason about which topology is most appropriate for a given sub-task.

1.3 Inter-Agent Communication Protocols

Agents must communicate effectively, and the strategist must design robust communication mechanisms:

Synchronous Request-Response: The simplest pattern, suitable for low-latency interactions where the calling agent can afford to block.

Asynchronous Message Passing: For long-running tasks, asynchronous patterns using message queues or event streams allow agents to continue working while waiting for responses. This is essential for building responsive, scalable systems.

Shared Memory and Blackboard Architectures: In some scenarios, agents collaborate by reading and writing to a shared memory space (the "blackboard"). This pattern is useful for collaborative problem-solving where multiple agents contribute partial solutions.

Handoff Mechanisms: A critical pattern identified in **LlamaIndex** research is the "handoff" mechanism, where an agent explicitly transfers control and context to another agent. The strategist must design clear handoff protocols that preserve context and prevent information loss during transitions.

Key Frameworks Demonstrating These Principles: LangGraph (stateful graphs), LlamaIndex AgentWorkflow (event-driven), Pydantic AI (context-based), Semantic Kernel (orchestration plugins), Haystack (pipeline architecture).

Skill 2: Interoperability and Integration Engineering

Original Skill: "Cross-Ecosystem Interoperability (Protocol Engineering)"

Why This Change: The original skill focused primarily on A2A and MCP protocols. While these are important, the research revealed that interoperability in 2026 extends far beyond these specific standards. The strategist must be able to integrate agentic systems with legacy enterprise infrastructure, handle multiple competing standards, and design for forward compatibility.

Enhanced Description:

By 2026, enterprises will have deployed agents from multiple vendors, each with different protocols and data formats. The strategist must be an expert in integration engineering, capable of building bridges between disparate systems.

2.1 Protocol Standards and Adaptation

Agent2Agent (A2A) Protocol: The A2A protocol, managed by the Linux Foundation, provides a standardized way for agents to discover each other's capabilities and collaborate. The strategist must understand the Agent Card specification, the JSON-RPC 2.0 transport layer, and the task lifecycle management (submitted, working, completed). Critically, the strategist must implement robust discovery mechanisms that allow agents to find collaborators based on semantic capability descriptions, not just hardcoded endpoints.

Model Context Protocol (MCP): MCP, introduced by Anthropic and now supported by major platforms including Google, standardizes how agents access tools and data sources. The strategist must understand the client-host-server topology and design MCP

servers that expose enterprise data in a secure, agent-friendly manner. A key insight from research is the importance of "Golden Skills" – curated, verified tool definitions that prevent agents from accessing dangerous or inappropriate APIs.

Beyond A2A and MCP: The strategist must also be familiar with other emerging standards, including OpenAPI specifications for tool definitions, the Open Agentic Schema Framework (OASF), and proprietary protocols from major cloud vendors (AWS Bedrock, Azure AI Studio, Google Vertex AI). The ability to translate between these standards and build adapter layers is critical.

2.2 Legacy System Integration

Most enterprises have decades of investment in legacy systems. The strategist must design integration patterns that allow agents to interact with:

Traditional REST and SOAP APIs: Wrapping legacy APIs with agent-friendly interfaces, including proper error handling and rate limiting.

Enterprise Databases: Designing secure, read-only (or carefully scoped read-write) access patterns that allow agents to query databases without risking data corruption.

Message Queues and Event Streams: Integrating with existing Kafka, RabbitMQ, or proprietary message bus systems to allow agents to participate in enterprise event-driven architectures.

Human-in-the-Loop Systems: Designing workflows where agents can request human approval or input at critical decision points, integrating with ticketing systems, approval workflows, and collaboration platforms like Slack or Microsoft Teams.

2.3 Security and Trust in Interoperability

When agents from different vendors collaborate, security becomes paramount. The strategist must implement:

Mutual Authentication: Ensuring that agents can verify each other's identities before sharing sensitive information.

Data Lineage and Provenance Tracking: Implementing "Toxic Flow Analysis" to track how data moves through multi-agent systems and identify potential security vulnerabilities.

Capability-Based Access Control: Ensuring that remote agents can only access the specific capabilities they need, using scoped tokens and fine-grained permissions.

Key Frameworks and Standards: A2A Protocol, Model Context Protocol (MCP), OpenAPI, OASF, AWS Bedrock Agents, Azure AI Studio, Google Vertex AI Agent Builder.

Skill 3: Production-Grade Observability and MLOps for Agents

Original Skill: "Agentic Observability, Debugging, and Self-Correction"

Why This Change: The original skill introduced the concept of agentic observability but lacked concrete implementation guidance. Research into frameworks like **Pydantic AI** (with its native Logfire integration) and enterprise MLOps platforms revealed that observability must be treated as a first-class concern from day one, not an afterthought. This enhanced skill emphasizes structured, actionable observability using industry-standard tools.

Enhanced Description:

Agentic systems are inherently non-deterministic and opaque. Traditional debugging techniques (stack traces, breakpoints) are insufficient. The 2026 strategist must implement comprehensive observability that makes the "black box" transparent.

3.1 Structured Observability with OpenTelemetry

The industry has converged on **OpenTelemetry** as the standard for observability. The strategist must implement:

Distributed Tracing: Every agent interaction must be instrumented as a trace, with spans representing individual cognitive steps (planning, tool use, reasoning, output generation). This allows the strategist to visualize the entire execution path and identify bottlenecks or failures.

Structured Logging: Logs must be structured (JSON format) and include rich context (agent ID, task ID, user ID, timestamps, input/output samples). This enables powerful querying and analysis.

Metrics Collection: The strategist must define and collect key metrics including latency (time to first token, total execution time), cost (tokens consumed, API calls made), error rates, and semantic quality scores.

Example Implementation: Pydantic AI provides native integration with **Logfire**, an OpenTelemetry-based observability platform. This demonstrates the best practice of embedding observability into the framework itself, rather than bolting it on later.

3.2 Cost and Performance Monitoring

Agentic systems can be expensive to run. The strategist must implement:

Real-Time Cost Tracking: Monitoring the cost of each LLM call, aggregated by agent, task, and user. This enables budget enforcement and cost optimization.

Performance Profiling: Identifying which agents or steps are consuming the most time or resources, enabling targeted optimization.

Anomaly Detection: Implementing automated alerts for unusual patterns (e.g., an agent making an excessive number of API calls, indicating a potential loop or malfunction).

3.3 Semantic Quality Evaluation

Traditional metrics (latency, error rate) don't capture whether an agent is producing *useful* outputs. The strategist must implement:

LLM-as-a-Judge: Using a separate LLM to evaluate the quality of agent outputs on dimensions like helpfulness, accuracy, safety, and instruction adherence. This provides a quantitative measure of semantic quality.

Human Feedback Loops: Integrating mechanisms for users to provide feedback on agent outputs, which can be used to fine-tune evaluation criteria and improve agent performance over time.

Regression Testing: Maintaining a test suite of representative tasks and continuously evaluating agent performance to detect regressions when system prompts or models are updated.

3.4 Self-Correction and Autonomous Debugging

The ultimate goal is for agents to debug themselves. The strategist must design:

Reflection Loops: Implementing "actor-critic" patterns where one agent generates a solution and another agent reviews it. If the critic identifies issues, the actor iterates.

Automatic Retry with Validation: When an agent produces output that fails schema validation (e.g., using **Pydantic** models), the system automatically provides the validation error back to the agent and requests a corrected output.

(Content truncated due to size limit. Use page ranges or line ranges to read remaining content)