

WHITE PAPER · AUGUST 2026

From enterprise architecture to agentic automation

How the TOGAF® framework gives enterprises the vocabulary to decompose their capabilities, and how that distillation reveals the workflows most ready for agentic AI.

Bartholdi Labs · Open Innovation Advisory

Access the world's innovation. Build tomorrow's champions.

The map precedes the machine

Enterprises do not lack agentic AI ambition. They lack an inventory of where autonomy belongs. Enterprise architecture already provides that inventory. This paper explains how.

That statement contains the entire argument of this paper. Agentic AI, meaning software agents that plan and execute multi step work under human direction, is moving from experiment to operating model. Gartner projects that by the end of 2026, 84 percent of enterprises will have increased their generative AI funding and 40 percent of enterprise applications will feature task specific agents. The question for leadership is no longer whether to deploy agents. It is where, in what order, and under what controls.

Most organisations answer that question anecdotally. A team pilots an agent where enthusiasm happens to live, not where the economics are strongest. The results are scattered proofs of concept, duplicated tooling, and no defensible answer when the board asks what comes next.

Enterprise architecture, and specifically the TOGAF Standard of The Open Group, offers the corrective. It gives the enterprise a shared vocabulary to describe itself as a structured system of capabilities, services, information flows, and technology. Once the enterprise is described in those terms, finding automation candidates stops being a guessing game and becomes a filtration exercise. We call that exercise distillation: decompose the enterprise into capabilities, refine capabilities into workflows, and qualify each workflow against explicit criteria for agentic readiness.

This paper walks through the method in four movements. First, what enterprise architecture and the TOGAF framework actually provide. Second, how the distillation process works in practice. Third, how the TOGAF Architecture Development Method maps onto a staged path to supervised autonomy. Fourth, the governance perimeter that makes any of this safe to run in production.

“

Enterprise architecture provides a framework and vocabulary by which to decompose the capabilities of an enterprise. Through the distillation process we find workflows which are candidates for agentic automation.

Travis Dayton, Chief Technology Officer

What enterprise architecture provides

An architecture framework is a foundational structure for developing a broad range of architectures. It describes a baseline state and a target state of the enterprise in terms of building blocks, shows how those blocks fit together, and plans the evolution from one state to the other. The TOGAF Standard, now in its 10th Edition, is the most widely adopted open framework of this kind, developed by consensus within The Open Group and used across sectors and geographies.

The TOGAF framework describes the enterprise across four architecture domains, and the discipline of the framework is that no domain is discussed in isolation. The Business Architecture defines strategy, governance, organisation, and key business processes. The Data Architecture describes the structure of logical and physical data assets. The Application Architecture provides the blueprint for individual applications, their interactions, and their relationships to core business processes. The Technology Architecture describes the software and hardware infrastructure capabilities required to support the other three, from cloud services to networks and middleware.

A vocabulary, not a bureaucracy

The practical power of the framework is semantic. It defines Architecture Building Blocks, which describe required capability, and Solution Building Blocks, which describe the components that implement that capability. It separates deliverables, artifacts, and building blocks so that catalogs, matrices, and diagrams can be reused rather than reinvented. It layers the work into four abstraction levels that answer four plain questions. Why is the architecture needed. What functionality must it deliver. How should that functionality be structured. With what assets shall it be implemented.

This vocabulary is precisely what agentic AI programmes are missing. An agent is, in TOGAF terms, a new class of Solution Building Block. It implements a capability that the business already performs, through a workflow that the architecture can already name. An enterprise that has decomposed itself into capabilities can therefore ask a precise question: which of our named workflows should this new class of building block execute, and under whose supervision. An enterprise that has not done this work can only ask a vague one: what could we do with agents.

Why this matters now

The economics of software delivery are shifting from instruction writing to decision supervision. As implementation becomes abundant, the scarce assets of the enterprise become intent, taste, and verification. Intent is exactly what an enterprise architecture encodes: a deliberate description of what the organisation does, why, and to what standard. Firms that hold a current architecture description enter the agentic era holding the map. Firms that do not will pay to redraw it under pressure.

The distillation process

Distillation is the movement from the whole enterprise to a ranked portfolio of agentic candidates. It proceeds in four refinements, each one narrowing the field and raising the purity of what remains.

1. Decompose

Model the enterprise as a capability map across the four TOGAF domains. Name what the organisation does, independent of who does it or which system supports it.

2. Refine

Within each capability, identify the concrete workflows and value streams that realise it. Attach volumes, cycle times, error rates, and cost to each.

3. Qualify

Filter workflows against explicit agentic readiness criteria. Most will fail. That is the point of a filter.

4. Select

Rank the survivors by value and risk. Pilot the top candidates as measurable experiments with defined success metrics and a named accountable owner.

What makes a workflow a candidate

A workflow qualifies for agentic execution when it satisfies conditions that can be stated and tested, not felt. The strongest candidates share five properties.

They are digital at both ends, consuming and producing machine readable inputs and outputs. They are high in volume and repetition, so that supervision cost is amortised across many executions. They involve rule informed judgement rather than open ended creativity, which suits the stochastic nature of agents governed by verification. They have measurable outcomes, so success is observable in activation, cycle time, cost, or quality rather than in opinion. And they carry bounded consequences, meaning an incorrect execution is detectable and reversible before material harm occurs. Payment authorisation and safety critical control fail the last test and remain under human authorship. Invoice matching, tier one support triage, compliance evidence gathering, and release note assembly pass it comfortably in most organisations.

Distillation also respects a lesson from product practice: validated learning beats speculation. Each selected workflow is piloted as a minimum viable deployment, the smallest agent that delivers real value and tests the critical assumption, instrumented from day one. Every pilot ends in an explicit decision to scale, adjust, or stop.

The ADM as an adoption engine

The TOGAF Architecture Development Method, the ADM, is an iterative cycle for moving an enterprise from baseline to target architecture in a controlled way. It was designed for exactly the kind of change agentic AI represents: change that touches multiple systems with multiple interdependencies. Each phase of the cycle has a direct agentic reading.

Vision, Phase A	Define the ambition and scope of the agentic programme, identify stakeholders, and secure the mandate. The output is a target picture, not a tool list.
Business, Phase B	Build the capability decomposition and workflow inventory. This is where distillation begins, anchored in business value rather than technology enthusiasm.
Data and apps, Phase C	Assess whether the data an agent needs is accessible, governed, and of usable quality, and how agents will interoperate with the application estate.
Technology, Phase D	Establish the platform preconditions: versioned declarative infrastructure, automated security testing in the pipeline, and observability of agent reasoning and cost.
Opportunities, Phases E and F	Convert qualified workflows into a sequenced portfolio and migration plan, with each candidate carrying its business case and risk tier.
Governance, Phases G and H	Oversee implementation and manage continuous change. For agents this discipline has a name, AgentOps: the governance of probabilistic systems through verification gates and audit evidence.

Autonomy itself is granted in stages, not declared. Practice across early adopters converges on a staircase: agents first observe and report in read only mode, then remediate in non production environments, then propose actions in production that a senior human approves, and only then receive supervised autonomy over low risk tasks inside an isolated perimeter. The human holds the merge at every stage that matters.

The governance perimeter

Agents are stochastic systems. The same input can produce different outputs, and failure arrives as a plausible wrong answer rather than an error message. Governing such systems is a design problem, and the enterprise architecture is where that design lives. Four mechanisms form the minimum perimeter.

Risk tiered boundaries

Classify workflows and repository paths by consequence. Protected areas such as payments, authentication, and core data migrations remain closed to autonomous modification and require staged human approval.

Bounded review loops

Every agent task carries an iteration cap and a budget threshold. An agent that cannot converge escalates to a human instead of consuming resources indefinitely.

Verification as a stack

Automated static and dynamic testing, policy as code, and independent model based review sit between agent output and production. Verification, not generation, is where quality is decided.

Accountable roles

A named governance function, from chief AI officer to AI auditor, maps responsibility to regulation, including the EU AI Act and ISO 42001, and maintains the evidentiary record.

The regulatory calendar makes this urgent for European enterprises. High risk obligations under the EU AI Act took effect in August 2026, and ISO 42001 certification is becoming the reference for demonstrating a managed AI system. An enterprise architecture that already documents capabilities, information flows, and controls is the natural backbone of that evidence. Governance ceases to be paperwork produced after the fact. It becomes a property of the architecture itself.

There is a final cultural point. The purpose of this perimeter is not to slow the enterprise down. It is to make speed safe. Organisations with mature engineering foundations see agentic tooling multiply their throughput. Organisations with fragile foundations see the same tooling multiply their defects. The architecture determines which of the two amplification effects a firm will experience.

CONCLUSION

A leadership agenda

The rationale for this approach is simple. Agentic AI creates value only where work is precisely described, and enterprise architecture is the discipline of precise description. Anchoring agent adoption in the TOGAF framework converts a fashionable technology question into a governed portfolio decision.

The risks and trade offs deserve equal honesty. Architecture work takes time, and a capability map can become a ceremony if it is pursued for its own sake. The corrective is to keep distillation outcome driven: every artifact exists to qualify or disqualify a workflow, and every qualified workflow becomes a measurable pilot. The other risk is over delegation. Autonomy granted faster than verification matures converts productivity into liability. The staircase model and the risk tiers exist to prevent exactly that.

The suggested next action is a distillation sprint. In four to six weeks, an enterprise can map its priority value streams, qualify an initial set of workflows against the five criteria, and leave with a ranked candidate portfolio, a platform readiness assessment, and a governance baseline aligned to the EU AI Act. That is a decision ready foundation, produced at advisory speed.

4

refinements in the distillation process, from capability map to ranked agentic portfolio

5

readiness criteria that separate credible agentic candidates from enthusiasm

6

weeks, at most, for a distillation sprint that produces a decision ready portfolio

Bartholdi Labs runs this method with executive teams as part of its open innovation advisory practice, combining enterprise architecture discipline with hands on agentic engineering research. The framework is open. The judgement is the craft.

TOGAF is a registered trademark of The Open Group. This paper draws on The TOGAF® Standard, 10th Edition, Introduction and Core Concepts, The Open Group, 2022, and on Bartholdi Labs research into agentic software engineering, 2026.