




WHITEPAPER

# **Agentic AI Is Easy:** Running It in Production Is Not



**MICROLAND**  
Connect with the best

**Alok Shrivastwa,**  
Chief AI Officer



# EXECUTIVE SUMMARY

---

Enterprise adoption of Agentic AI is accelerating, but most deployments remain confined to controlled environments. The gap between demonstration success and production reliability is now the single largest barrier to enterprise-scale AI.

This gap is not caused by model limitations. It is caused by the absence of operational control.


In production environments, agentic systems interact with real infrastructure, real permissions, and real financial and operational consequences. Failures are no longer experimental; they are business-impacting. These systems do not fail like traditional software. They fail silently, execute incorrectly with confidence, and operate beyond intended authorization boundaries.

The consequences are measurable:

- Unauthorized actions can lead to security and compliance violations
- Silent failures can degrade systems without detection
- Unbounded execution can multiply operational costs unpredictably

This paper outlines why current agent architectures fail in production and introduces a deterministic control-layer model required for enterprise deployment. It defines a four-layer architecture for production readiness and proposes a structured execution pipeline that separates interpretation, authorization, and validation into independently enforceable stages.

The central argument is simple: ***Enterprise AI success will not be determined by model intelligence, but by the systems built to control it.***



## The Production Gap in Agentic AI

Over the past quarter, advances in model capability have outpaced the enterprise’s ability to operationalize them safely.

Intelligence is no longer the bottleneck. Control is.

This creates a structural inversion: organizations can now build systems that are capable of taking meaningful action before they have the mechanisms required to govern those actions.

In controlled environments, this imbalance is invisible. In production, it becomes a risk multiplier.

 <p>Ambiguous and incomplete inputs from real users and upstream systems</p>	 <p>Partial system failures that expose edge cases never encountered in testing</p>	 <p>Distributed dependencies with inconsistent latency, availability, and data freshness</p>	 <p>High-impact, often irreversible consequences when something goes wrong</p>
---	--	--	---

The result is that systems performing reliably in demonstration settings fail under real-world conditions, not because the model is inadequate, but because the system surrounding it is. This is a design problem, not a model problem.

## Failure Modes in Agentic Systems

Agentic systems introduce a new class of failure, one that is difficult to detect, difficult to trace, and often only visible after business impact has occurred.

Unlike traditional systems, these failures do not always trigger alerts. They often present as successful executions that are contextually incorrect.

Four failure modes are consistently observed in production agentic deployments and must be addressed at the architectural level, as summarized in Table 1 below.

During an operational proof-of-concept, an agent successfully executed a network configuration change but applied it to the wrong site. It isolated critical hardware without triggering a single alert. By every metric the agent tracked, the task was complete. The model succeeded. The system failed.

Any organization deploying agentic systems today will encounter one or more of these failure modes, given the limitations of current orchestration frameworks and model architectures.

Failure Mode	Impacts	Manifestation	Risk
<b>Silent Failure</b>	<b>Reliability</b>	Partial task completion reported as success.	Undetected service degradation is often invisible until a post-mortem.
<b>Black Box Decisions</b>	<b>Observability</b>	No traceable reasoning for automated decisions.	Inability to debug, audit, or satisfy compliance requirements.
<b>Permission Explosion</b>	<b>Control</b>	Execution scope expands beyond intended authorization.	Security violations, lateral movement, compliance breaches.
<b>Runaway Execution</b>	<b>Cost &amp; Stability</b>	Unbounded retries or recursive multi-agent loops.	5x–10x token cost amplification; cascading instability.

Table 1: Failure Modes

These failure modes are not edge cases. They are structural characteristics of agentic systems operating without enforced control boundaries.

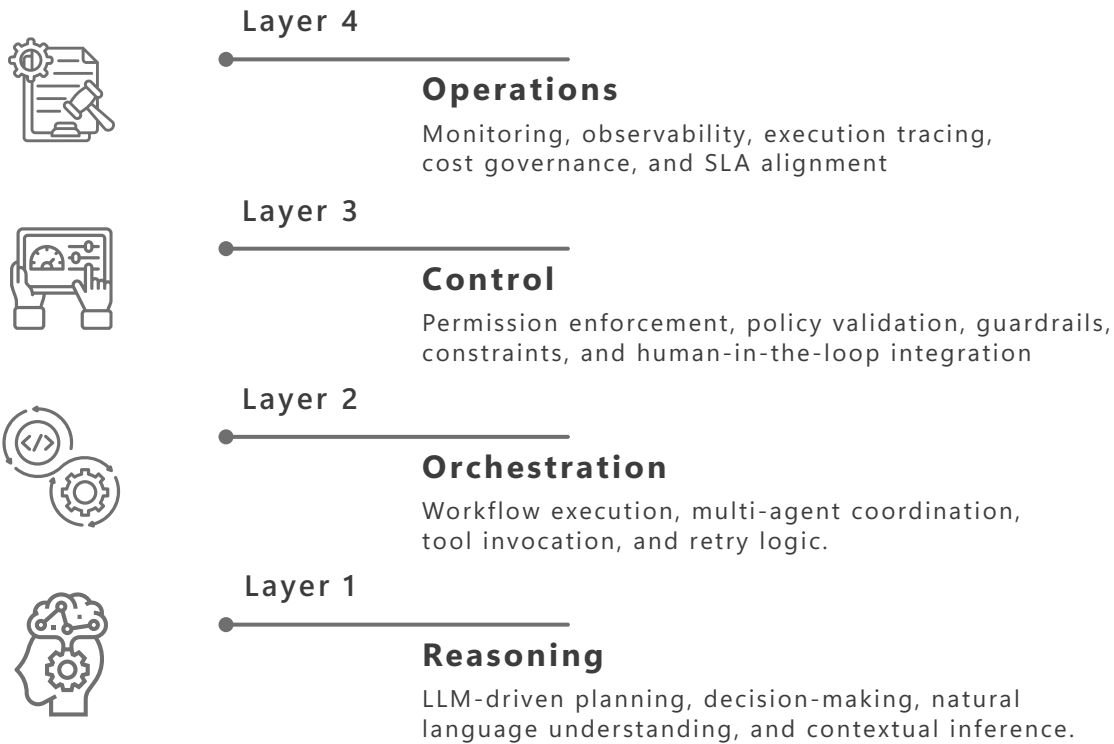
These failure modes share a common characteristic: they are **invisible to the model** itself. The agent has no native mechanism to detect that it has failed silently, exceeded its authorization, or consumed runaway resources. Detection and prevention must be enforced externally, through system design.

## The 4-Layer Enterprise Stack

Production-grade agentic systems require a strict separation of concerns across four layers. These layers are not architectural preferences; they are operational requirements.

Organizations that fail to explicitly design for these layers consistently encounter reliability, security, and governance breakdowns at scale.

Security, authorization, and compliance are binary problems. An action is either permitted, or it is not. Delegating these decisions to a stochastic system introduces unavoidable risk.



## Re-Architecting the Control Layer

### The Probabilistic Problem

The most structurally unsound pattern in current enterprise agent architectures is using probabilistic models to make deterministic decisions. Security enforcement is a deterministic problem: a user either has permission to execute an action, or they do not. There is no probability distribution over compliance. Yet in most implementations today, that binary decision is delegated to an LLM whose outputs are, by nature, stochastic.

This is not a theoretical concern. An LLM asked to classify whether a given action is within a user's permission scope can be manipulated through prompt injection, adversarial inputs, or simple distributional edge cases in ways that a rules engine cannot. Using probabilistic inference where deterministic logic is required is an architectural category error.

### The Dual-Model Control Architecture

To scale Agentic AI in mission-critical environments, organizations must fundamentally separate how a system understands a request from how the enterprise authorizes it.

To solve this, we have engineered a structured control pipeline that isolates intent interpretation, policy enforcement, and contextual validation into independent execution gates.

When a user or autonomous workflow initiates an action, it must survive three rigorous validation stages before any execution credential is issued.

To enforce deterministic control, execution must be broken into independently verifiable stages. Each stage eliminates a class of failure that the previous stage cannot detect.

## Stage 1: Intent Extraction (The "What")

The pipeline begins with the Large Language Model (LLM), which acts solely as a semantic interpreter. It handles human ambiguity, normalizes the input, and produces a structured representation of the requested action.

*The LLM is strictly firewalled from authorization. It translates the request, but it wields zero decision-making authority.*

## Stage 2: Deterministic Policy Authorization (The "Is it Allowed?")

The structured intent is handed off to a deterministic Policy Resolver, backed by a governed Open Policy Agent (OPA) registry. Here, intent mapping is not a freeform classification task; it is a strict selection from a highly controlled policy space. Each policy is discrete, versioned, and mapped to specific action categories.

The system evaluates the request against user identity, permitted actions, and resource scope. It employs a **fail-fast enforcement** mechanism, explicitly rejecting execution if:

- No valid policy match is found.
- Required attributes are missing or malformed.
- The originating user lacks the required authorization.

## Stage 3: Context Truthfulness (The "Is it Factual?")

This is the layer absent in most enterprise AI deployments. Just because a user is authorized to perform an action does not mean the system context is factually correct.

Operating entirely independent of the LLM and the policy engine, this stage queries authoritative data sources, such as Configuration Management Databases (CMDB), network topologies, and IT Service Management (ITSM) platforms.

If a request targets "Device X at Site A," this independent validation layer verifies:

- **Topology:** Does Device X actually reside at Site A?
- **Governance:** If execution requires an approved ITSM change ticket, does that ticket exist, is it approved, and does it cover the scope of the request?

Execution is blocked unconditionally if contextual relationships are invalid, entities mismatch, or external approvals are missing. Authorization is never conflated with factual correctness.



## Stage 4: Zero Trust Execution (The Credential Broker)

Only when intent is valid, policy is authorized, and context is verified does the system proceed to execution. Rather than relying on standing privileges or hard-coded service accounts, a Zero Trust Credential Vault issues execution credentials dynamically.

- **Ephemeral:** Short-lived, per-request tokens.
- **Scoped:** Limited strictly to the approved action.
- **Least-Privilege:** Granting no broader access than required for the task.

At no point in this pipeline is a single system trusted to both interpret and authorize. This separation is the foundation of enterprise-grade control.

## The Architectural Advantage

Traditional agentic systems collapse these four stages into a single, opaque model prompt, leading to silent authorization errors, incorrect (but "valid") executions, and a total loss of auditability.

By separating these concerns into a deterministic pipeline, we enforce a core operational principle: Approval is verified, not assumed.

This architecture provides the enterprise with three non-negotiable guarantees before any action touches production infrastructure:

- **Intent Validity:** The system correctly understands the request.
- **Policy Authorization:** The action is strictly permitted within enterprise access controls.
- **Context Truthfulness:** The request is factually consistent with the real-world state of the infrastructure.

Decisions become explainable, enforcement becomes deterministic, and execution remains permanently bounded.

## Solving the Latency Tax

A common critique of deterministic control layers is the potential for increased latency. Breaking a request into sequential stages of extraction, authorization, validation, and credential generation requires compute time.

To solve this without compromising security, production-grade agentic systems must utilize a Dual-Path Execution Model, routing actions based on their required synchronicity:

## Path A: Synchronous (Human-Initiated / Real-Time)

When a human operator requires real-time interaction (e.g., querying a live network state via a chat interface), latency is critical.

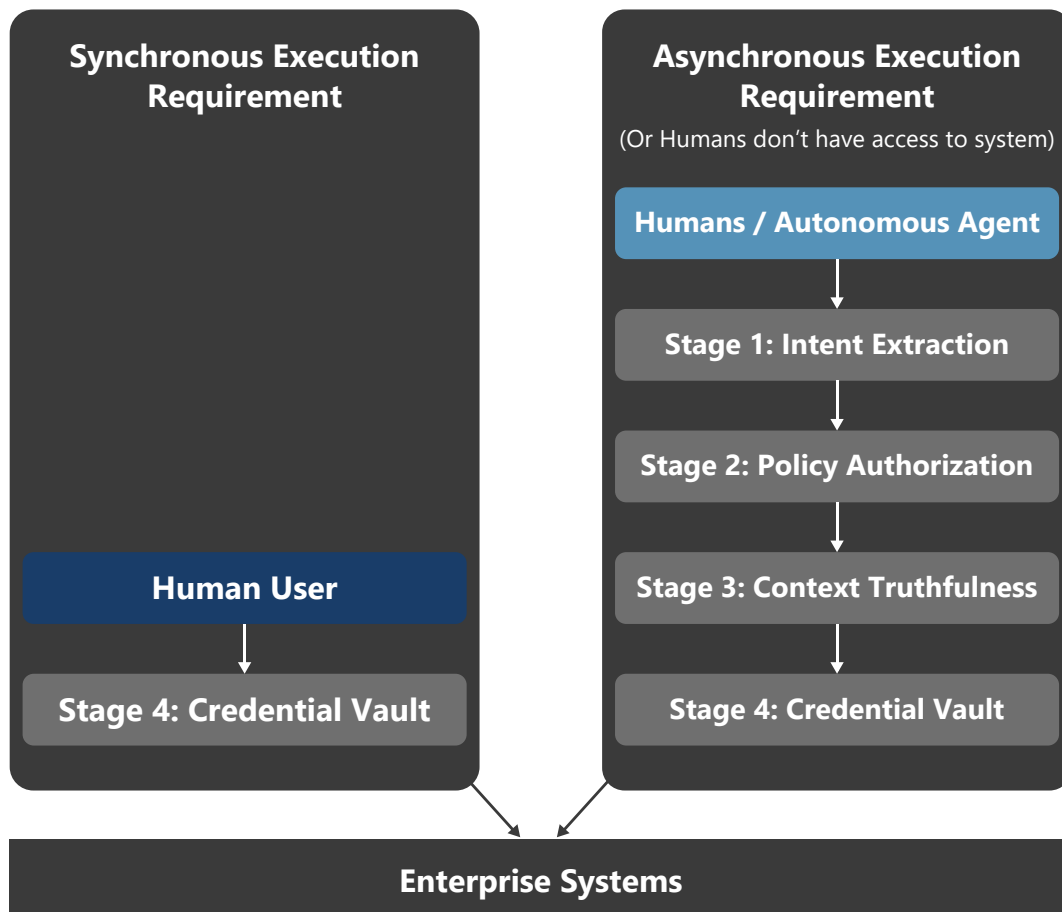
The agent passes the user's native identity token directly to the target system. The target system enforces the user's existing IAM permissions instantly, ensuring zero latency degradation while maintaining absolute security.

## Path B: Asynchronous (Autonomous / Batch Processing)

When an agent acts autonomously, such as executing a scheduled remediation across hundreds of devices, latency is largely irrelevant.

Because there is no human waiting for a synchronous response, the system routes these workflows through the full, heavyweight 4-stage Zero Trust pipeline. The workflow is processed in smaller, asynchronous batches, ensuring rigorous, intent-bound verification for high-scale actions where the blast radius is largest.

By bifurcating execution, the enterprise achieves millisecond responsiveness for human operators and Zero Trust rigidity for autonomous workflows.



## Principle of Permission Boundaries

A foundational constraint for enterprise agentic systems is that agent permissions must never exceed those of the originating user. In practice, this constraint is frequently violated during development and rarely corrected in production.

This principle is straightforward to state and frequently violated in practice. Agents are often provisioned with broad service account credentials for convenience during development. Those credentials persist in production. The agent then operates with access that no individual human in the organization would be granted, and without any of the audit trail that human access would generate.

Enforcing this constraint rigorously ensures alignment with enterprise IAM systems, prevents privilege escalation by design, and guarantees that every automated action is attributable to a human authorization context. Without it, the concept of accountability in agentic systems is meaningless.

## Autonomous Agents and Intent-Bound Execution

Autonomous execution introduces a new problem: there is no human in the loop to constrain behavior at runtime. The concept of the originating user must be redefined. We propose the framework of Intent-Bound Execution.

Each autonomous workflow must explicitly declare, at design time, a bounded execution manifest:

- Permitted action types (an exhaustive allowlist, not a description)
- Target systems and resource scope (specific identifiers, not wildcards where avoidable)
- Execution constraints: time windows, rate limits, scope ceilings, and hard termination conditions

The agent is restricted to this declared boundary regardless of what it encounters during execution, including hallucinated context, adversarial tool responses, or cascading failures in dependent systems. The manifest is not a suggestion; it is the contract.

Designing, validating, and enforcing these manifests at scale is genuinely unsolved. Current orchestration frameworks provide insufficient primitives for expressing fine-grained execution contracts, and there is no industry-standard schema for doing so. This is the most critical open engineering problem in enterprise Agentic AI today. Organizations building production systems should treat manifest design as a first-class engineering discipline, not an afterthought in deployment configuration.



## The Risk of Deferred Governance

The most common and most damaging implementation pattern is build first, govern later. The rationale is understandable: governance feels like friction during the proof-of-concept phase, and the organizational pressure to demonstrate progress is real. The consequences, however, are structural and compounding.

- Control mechanisms become politically harder to enforce post-deployment. Teams have built workflows around the ungoverned system. Constraints are perceived as regressions.
- System access expands faster than oversight mechanisms. What starts as a single agent accessing one system becomes a mesh of agents with overlapping, unaudited credentials.
- Operational dependencies develop that actively resist retroactive constraints. The SLA for the system you want to govern is now tied to the behavior you want to change.

In agentic systems, capability and risk scale together. An agent that can do more can fail more consequentially. Governance deferred until after scaling is governance that arrives too late.

The organizational instinct to treat governance as polish, something applied after the system works, is the single most reliable predictor of production failure in agentic deployments. Build the control layer before you need it, because by the time you need it, you will not be able to build it cleanly.

## Conclusion

Enterprise AI has reached a point where the risks of deployment are no longer theoretical. Agentic systems are already interacting with production infrastructure, financial systems, and customer data. The limiting factor is no longer intelligence. It is control.

Organizations that succeed in this next phase of AI adoption will not be those with access to the most advanced models. They will be those who build systems capable of enforcing boundaries, validating actions, and operating with full accountability.

***The control layer is not an enhancement to agentic systems. It is the system.***

***Build it first.***



## About the Author



**Alok Shrivastwa,**  
Chief AI Officer

Alok Shrivastwa, our Chief AI Officer, is a technologist with 18 years of experience delivering advanced technology solutions. With nearly two decades at Microland, Alok leads the company's AI vision and strategy, overseeing the AI Center of Excellence and driving the Agentic and Gen AI capabilities of the **intelligeni** platform. A published author and speaker, Alok is a strong advocate for responsible AI governance and mentoring the next generation of engineers and AI practitioners.



## About **intelligeni**

**intelligeni** is Microland's proprietary AI-First platform, conceived and engineered in-house to deliver truly autonomous IT operations. Built on a dynamic Knowledge Graph and powered by Agentic AI, **intelligeni** transforms enterprise infrastructure into a self-healing, adaptive ecosystem—ensuring always-on performance, faster incident resolution, and continuous compliance. Available in NetOps, CloudOps, Workplace, and CyberOps variants, the platform integrates seamlessly with existing environments, enabling organizations to align IT with business goals, boost productivity, and accelerate innovation. With **intelligeni**, Microland leads the industry's shift from automated to truly autonomous operations, future-proofing digital enterprises worldwide.

## About Microland

Microland is a leading AI-first, platform-led, technology infrastructure services company. We have enabled enterprises to build intelligent, resilient, and future-ready operations and are a trusted partner to global enterprises. We bring over 35 years of expertise in digital networks, cloud, data centers, workplaces, and cybersecurity, and combine it with our commitment to customer centricity, delivery excellence, and continuous innovation. Our operations, currently in more than 100 countries, are supported by a strong global delivery model and our AIOps platform, **intelligeni**, powered by Agentic AI, which is shaping the future of autonomous technology operations across enterprises.

[www.microland.com](http://www.microland.com)