PITHY NOTES PUBLICATIONS

PithySignal

PITHY SIGNAL: HYPER TPRM

SIGNAL BRIEF | ISSUE 002

Agent Sprawl Is the New Shadow IT

Why enterprises need governance models for autonomous AI agents operating across vendor ecosystems.

Organizations are rapidly deploying AI agents across enterprise workflows and vendor ecosystems faster than governance models can adapt, creating agent sprawl as the AI-era equivalent of shadow IT.

Created by Pithy Notes Publications | Published under Pithy Signal

KEY SECTIONS

Category: Hyper TPRM
Publication Date: 2026-05-21
Issue: 002
Tags: Hyper TPRM | Agent Governance | Shadow IT | Runtime Governance

This publication was created by Pithy Notes Publications and published as part of the Pithy Signal governance intelligence series.

Executive Summary

Traditional third-party risk management programs were designed for relatively stable software environments operating within predictable review cycles.

AI agents challenge these assumptions.

Organizations are increasingly deploying autonomous and semi-autonomous AI agents across procurement operations, productivity systems, customer workflows, analytics environments, and third-party ecosystems faster than governance models can adapt.

Some agents are embedded within vendor platforms. Others are developed internally by business units attempting to accelerate operational efficiency and automation.

Many operate with:

- unclear ownership
- fragmented oversight
- inconsistent governance controls
- evolving runtime behavior
- uncertain operational boundaries

This introduces a new governance problem: agent sprawl.

Like shadow IT before it, unmanaged agents create visibility gaps and operational uncertainty. But unlike traditional software, agentic systems may autonomously execute actions, interact across systems, chain decisions dynamically, and evolve behavior over time.

This changes the nature of oversight itself.

Static review models designed around periodic assessment cycles are increasingly misaligned with continuously operating AI-enabled environments.

Hyper TPRM emerges as a response to this shift by combining:

- runtime governance
- continuous evidence assurance
- AI oversight
- operational telemetry

- delegated authority controls
- human accountability mechanisms

The issue is not simply technological.

It is operational, organizational, and strategic.

1. Shadow IT Has Evolved

For more than a decade, organizations struggled to govern shadow IT: unsanctioned software, unmanaged cloud services, and decentralized technology adoption occurring outside centralized oversight models.

AI agents represent the next evolution of this governance problem.

The barrier to deploying autonomous operational systems is rapidly decreasing. Employees can now create or deploy agents capable of:

- orchestrating workflows
- accessing enterprise systems
- generating outputs
- interacting with external tools
- triggering downstream actions
- automating decision pathways

Many organizations are adopting these capabilities faster than governance structures can mature.

This creates a widening gap between: AI operational deployment and institutional oversight readiness.

2. Agent Sprawl Is an Operational Governance Problem

Agent sprawl is often framed as a technology management issue.

In reality, it is an operational governance problem.

The core challenge is not merely that agents exist. The challenge is that organizations increasingly lack clear visibility into:

- who deployed them
- what systems they access

- what permissions they possess
- what decisions they influence
- how runtime behavior is monitored
- how accountability is established

Traditional governance models were largely built around:

- identifiable applications
- static permissions
- predictable execution paths
- periodic review cycles

Agentic systems destabilize these assumptions.

Governance increasingly shifts from: static assessment to continuous operational oversight.

3. Why Agentic Systems Change Oversight

Traditional software systems generally operate within relatively deterministic boundaries.

Agentic systems introduce:

- probabilistic outputs
- delegated autonomy
- adaptive execution
- chained decision-making
- dynamic tool interaction
- evolving contextual behavior

This creates a form of operational uncertainty that static questionnaires and periodic evidence snapshots cannot fully capture.

As AI systems become increasingly embedded across enterprise and third-party ecosystems, governance models must evolve toward:

- runtime observability
- continuous evidence validation

- behavioral monitoring
- identity-aware controls
- escalation-aware oversight

The shift is analogous to the evolution from: point-in-time cybersecurity to continuous security operations.

4. The Governance Gap Is Widening

Many organizations currently lack:

- centralized AI agent inventories
- runtime telemetry standards
- delegated authority frameworks
- operational escalation models
- AI identity governance
- cross-functional ownership structures

This creates governance debt.

The issue is not necessarily that organizations lack AI capabilities. The issue is that governance maturity is lagging behind operational deployment velocity.

This creates increasing exposure across:

- procurement
- compliance
- legal
- cybersecurity
- operational resilience
- vendor ecosystems

As organizations scale AI adoption, governance fragmentation itself may become a systemic operational risk.

5. Hyper TPRM Emerges as a Response

Hyper TPRM represents an emerging governance operating model designed for AI-enabled ecosystems operating beyond traditional review assumptions.

The model extends beyond: vendor questionnaires and periodic assessments.

Instead, Hyper TPRM increasingly emphasizes:

- continuous evidence assurance
- runtime governance
- operational telemetry
- AI-enabled vendor oversight
- agent accountability
- delegated authority controls
- identity-aware governance
- human oversight mechanisms

This represents a structural shift in how organizations govern operational trust.

6. Strategic Implications

Organizations should expect:

- increased pressure for AI runtime governance
- convergence between IAM and AI oversight
- expansion of continuous monitoring models
- procurement and governance integration
- increased demand for operational evidence
- greater focus on AI accountability structures

Over time, organizations may increasingly govern AI agents as operational actors rather than passive software components.

This distinction fundamentally changes governance expectations.

7. What Organizations Should Do Next

Organizations should begin establishing:

- centralized AI agent inventories
- AI governance registration processes
- runtime monitoring requirements
- delegated authority thresholds
- operational escalation pathways
- human oversight controls
- AI identity management standards
- cross-functional governance ownership

The organizations that mature these capabilities early may gain significant operational resilience advantages as AI-enabled ecosystems continue expanding.

Closing Signal

Shadow IT introduced unmanaged software risk.

Agent sprawl introduces unmanaged operational autonomy.

This is one of the defining governance shifts emerging behind Hyper TPRM.

Signal Sources

- Gartner - "Six Steps to Manage Artificial Intelligence Agent Sprawl"
- Cloud Security Alliance - "Agentic AI NIST AI RMF Profile"
- Google Cloud Office of the CISO - "AI Governance Tips to Counter Shadow Agents"
- Palo Alto Networks - "What Is Agentic AI Governance?"
- NIST AI Risk Management Framework
- Microsoft Security - "Observability, Governance and Security Shape the New Frontier"
- Reuters - "CrowdStrike to Buy SGNL to Tackle AI Identity Threats"

Created by Pithy Notes Publications Published under Pithy Signal

Governance intelligence publications focused on AI governance, risk management, compliance, and Hyper TPRM.