


Big Themes in AI for 2026: Navigating the Age of Autonomous Agents

AI: The Strategic Frontier –
Beyond the Hype to Real-World Impact Webinar Series
Episode 4

Before We Get Started



Recording

A link to the recording and slides will be emailed to all registrants.



Recording

Type in the question box, and we will answer in real time or during the Q&A.



Social

Follow us on LinkedIn, Facebook, Youtube, and/or Instagram or visit blackhills.ai to see upcoming and on-demand webinars.

Panel



Tom Marlow

Chief Artificial Intelligence Officer,
Black Hills AI



Manjeet Rege

Director of Center of Applied
Artificial Intelligence and Professor,
University of St. Thomas
Advisor to Black Hills AI

The Shift to 2026

From Assistance to Autonomy: The Evolution of AI in the Workplace

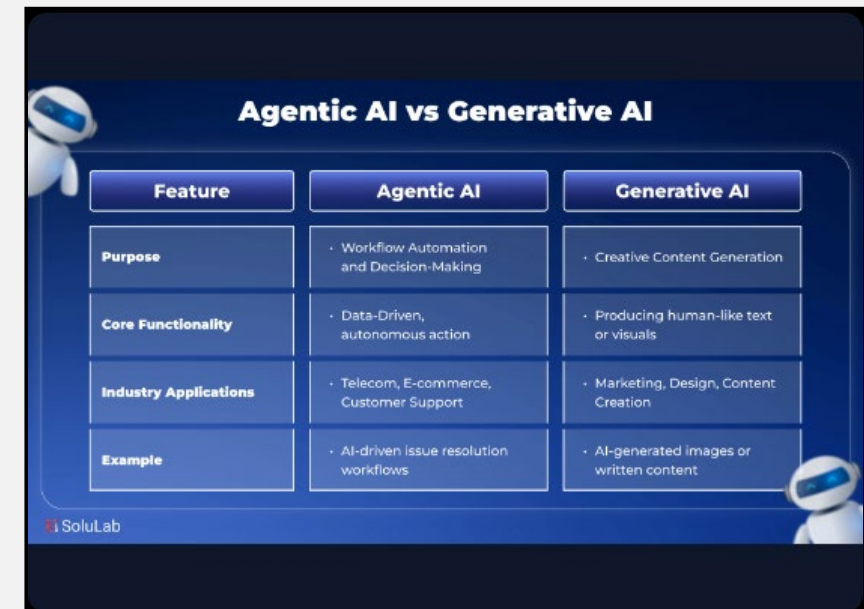
GenAI vs. Agentic AI

Generative AI (The Assistant)

- Reactive: Waits for human input (prompts).
- Function: Creates content (drafts, summaries, code)
- Limit: Cannot leave the chat window to perform actions.

Agentic AI (The Worker)

- Proactive: Given a high-level goal, it plans the steps.
- Function: Executes tasks (sends emails, queries DBs, books meetings).
- Power: Has "agency" to act in the digital world.



Agentic AI vs Generative AI		
Feature	Agentic AI	Generative AI
Purpose	• Workflow Automation and Decision-Making	• Creative Content Generation
Core Functionality	• Data-Driven, autonomous action	• Producing human-like text or visuals
Industry Applications	• Telecom, E-commerce, Customer Support	• Marketing, Design, Content Creation
Example	• AI-driven issue resolution workflows	• AI-generated images or written content

SoluLab

Theme 1: The Invisible Workforce



Goal-Oriented

Agents work towards a defined outcome (e.g., "Renew all vendor contracts expiring in Q1") rather than just answering a single question.



Multi-Step Execution

They can chain tasks together:
Search for a file -> Read it ->
Draft a reply -> Send the
email -> Archive the thread.



Self-Correcting

If an agent encounters an error (e.g., "File not found"), it can reason, search an alternative location, and retry without human help.

The Multi-Agent Ecosystem

From One to Many

- The Swarm: Agents rarely work alone. A "Legal Agent" will communicate with a "Finance Agent" and a "Sales Agent" to close a deal.
- Cascading Risks: If one agent in the chain hallucinates data, it poisons the decision-making of every other agent in the network.
- The "Weak Link" Problem: Liability becomes harder to trace when multiple autonomous systems interact without human intermediaries.




The “Unsupervised Associate” Risk

Delegation Without Control

The primary risk of 2026 is not copyright infringement (the GenAI fear), but **liability from autonomous action**.

If an AI agent autonomously negotiates a bad clause or sends a trade secret to a public server, the "Human in the Loop" defense fails if the loop wasn't designed correctly.

Key Question: How do you audit a machine that chose its own path?



Theme 2:

The Governance Imperative

Responsible AI (RAI) as Your Legal Defense

The Governance Gap

1. The Black Box Problem

Risk: Explainability

When an agent makes a decision (e.g., "Reject this vendor"), you must be able to explain the "Why" to regulators and courts.

Lack of an audit trail means lack of a legal defense in negligence claims.

2. The Data Leakage Risk

Risk: Confidentiality

Autonomous agents seek the most efficient route. Without "Data Fencing," an agent might process confidential client data using a public, unsecured API.

This creates an automated breach of attorney-client privilege.

Pillars of Responsible AI



Accountability

Assigning human ownership to every agent.



Transparency

Mandating "Chain of Thought" logging.



Safety

Enforcing strict data fencing protocols.

Risk vs. Reward in AI Adoption

Efficiency Gains vs. Compliance Risk (Conceptual Model)

Scenario A: Ungoverned "Fast" Adoption



Scenario B: "Responsible" Adoption



Theme 3: The New Legal Operating Model

From Reactive Counsel to Proactive Architects

The Procurement Gatekeeper

Buying the Bot: Redlining for Risk

- Vendor Liability: When buying AI agents, standard software indemnities fail. You must define liability for *autonomous actions* taken by the vendor's bot.
- Standard of Care: Define the expected accuracy and "judgment" thresholds. What happens when the agent makes a technically correct but contextually disastrous decision?
- Audit Rights: Demand the right to inspect the vendor's "Chain of Thought" logs during a dispute.




The Lawyer's Checklist for 2026



Define HITL Thresholds: Clearly categorize decisions. Low risk? Autonomous. High risk (contracts, filings)? Mandatory Human-in-the-Loop.



Mandate Audit Trails: Require that every AI agent generates a non-editable log of its inputs, "thought process," and outputs for e-discovery purposes.



Enforce Data Fencing: Work with IT to ensure agents operating on sensitive data cannot access the open internet or public LLMs.

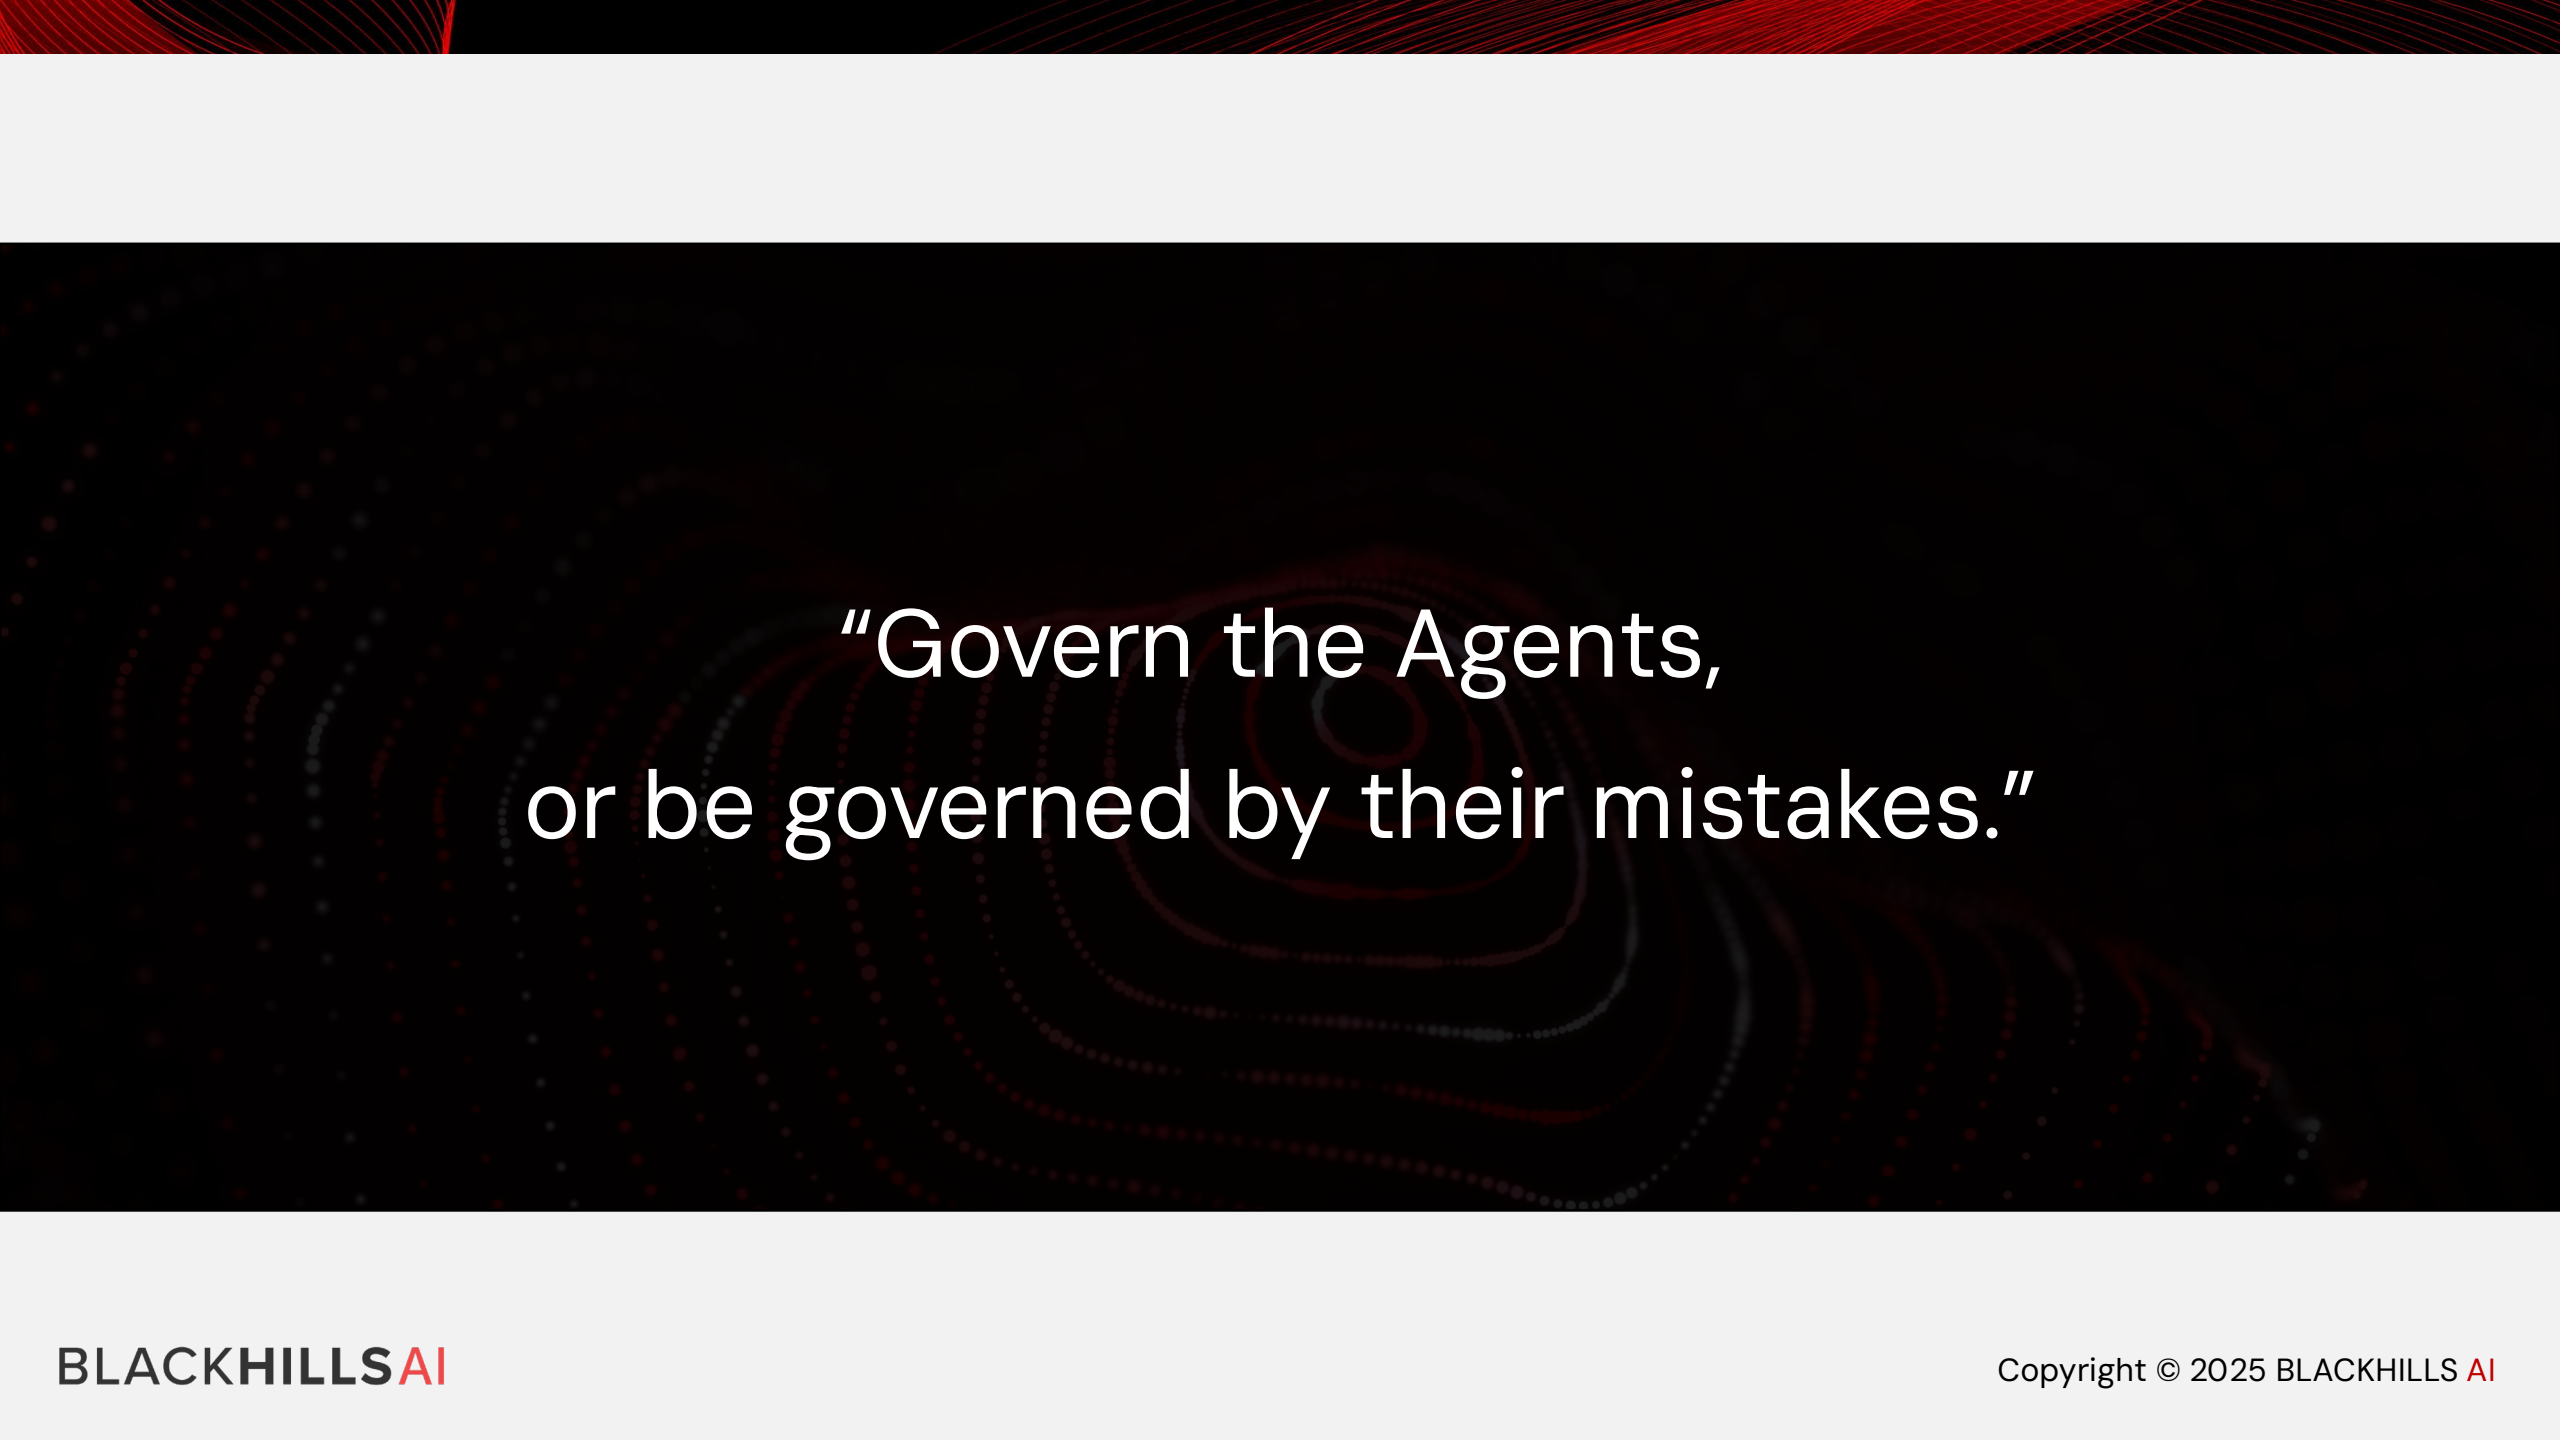


Create the Agent Policy: Move beyond a "GenAI Policy" to an "Autonomous Systems Policy" that defines liability and oversight roles.

The Human Element

"The law is more than a system of rules; it is a system of judgement.
AI can process the rules, but only humans can provide the judgement."

— Adapted from Richard Susskind



“Govern the Agents,
or be governed by their mistakes.”

Agentic AI in IP Work

From Generative to Autonomous Intelligence

AI in IP: The Key Distinction

Generative AI: Creates content based on prompts

Agentic AI: Acts autonomously to achieve multi-step objectives

The question isn't whether to use AI—it's how to use it effectively and safely in your IP processes.

How Much Error is Acceptable?

High-Consequence Tasks

Claims drafting, office action responses
→ Human review required

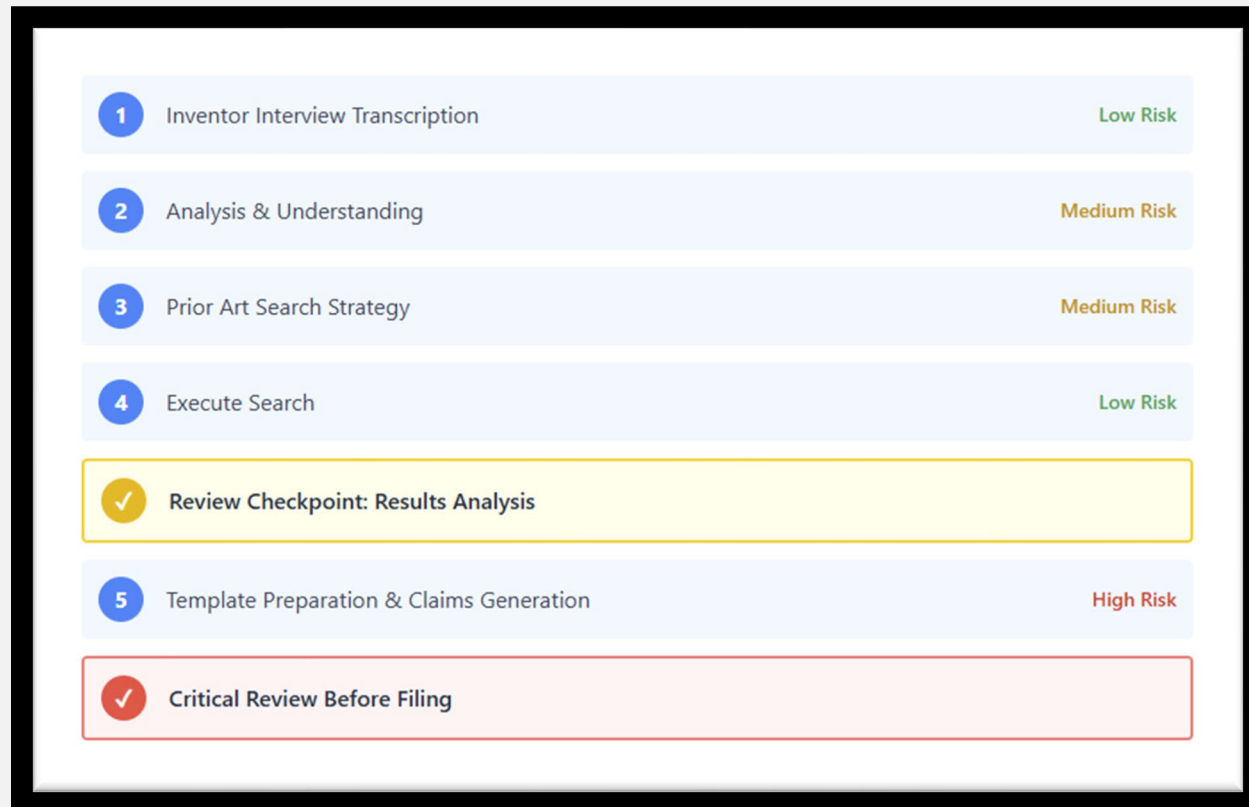
Intermediate Tasks

Prior art searching, initial analysis
→ Spot-check, balance re-work versus intervention

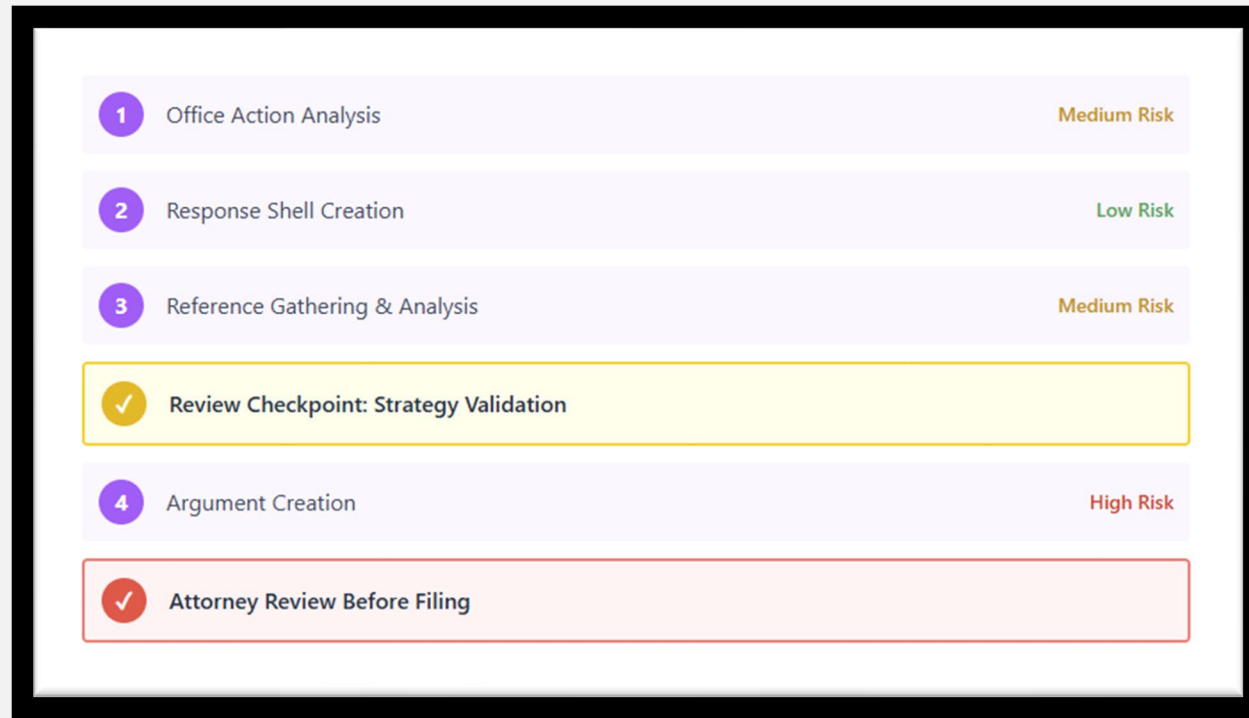
Lower-Risk Tasks

Transcription, formatting, organization
→ Automation ready

Example: Patent Application Development



Example: Office Action Response



Where Do You Want to Step In?

Full Trust Mode

Minimal human intervention on routine tasks

Spot-Check Mode

Random or targeted human review of AI outputs

Checkpoint Mode

Human approval at every critical step before proceeding

Why Transparency Matters

Compliance Reality: Transparency in AI use is moving from best practice to legal requirement.

- EU AI Act (effective Aug 2026): mandatory transparency disclosures
- California AI Transparency Act
- Client expectations: full disclosure of AI involvement
- Audit trails: document every AI-assisted step

Building Transparency: Black Hills AI Model

Guided Skills

Pre-built workflows designed specifically for IP professionals

Transparent Prompts

All prompts visible, editable, and customizable

Key Takeaways

- Agentic AI enables autonomous workflows; use strategically
- Determine your comfort level at each workflow step
- Transparency is no longer optional—it's required
- Maintain human control at high-consequence checkpoints
- Document everything for compliance and IP protection

The Future Belongs to Human + AI

Smart IP professionals won't choose between doing it themselves or using AI.

They'll master the integration—keeping humans in control of judgment,
clients in the loop, and compliance at the center.

Thank you for your interest.

Questions?

Tom Marlow
tmarlow@blackhills.ai