Open for Innovation KNIME

Prophix

Empowering the Modern Auditor with Real-World Agentic AI

(with practical, realistic examples)

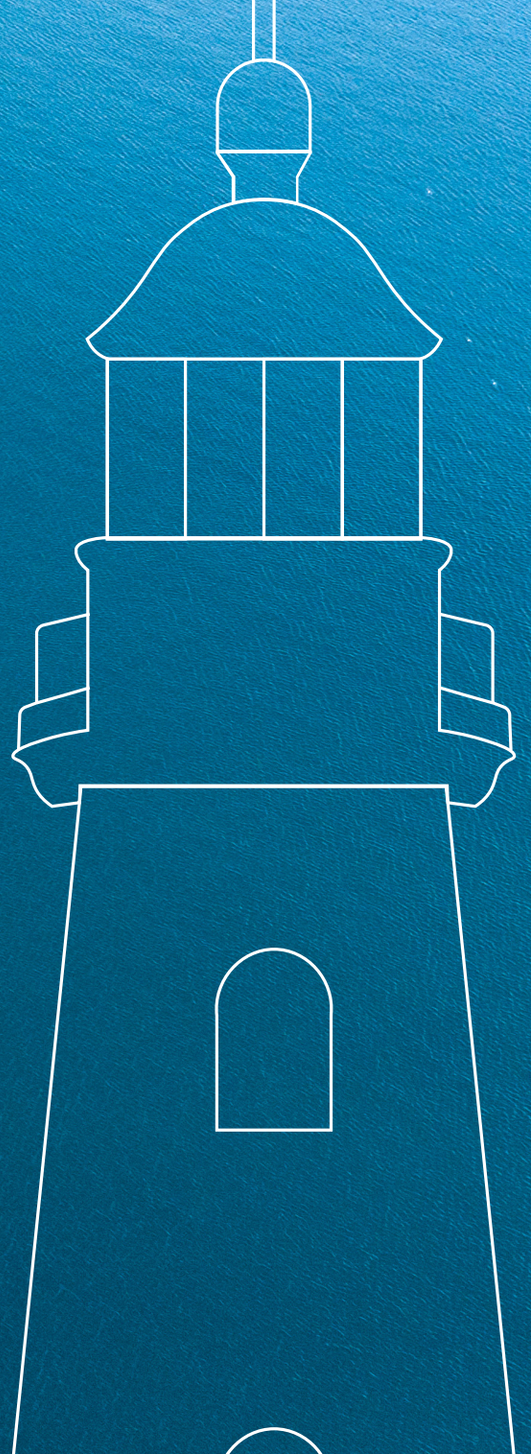

Trusted within the most highly regulated industries in the world.

DTCC ING IGT Grab citi Rabobank

Table of Contents

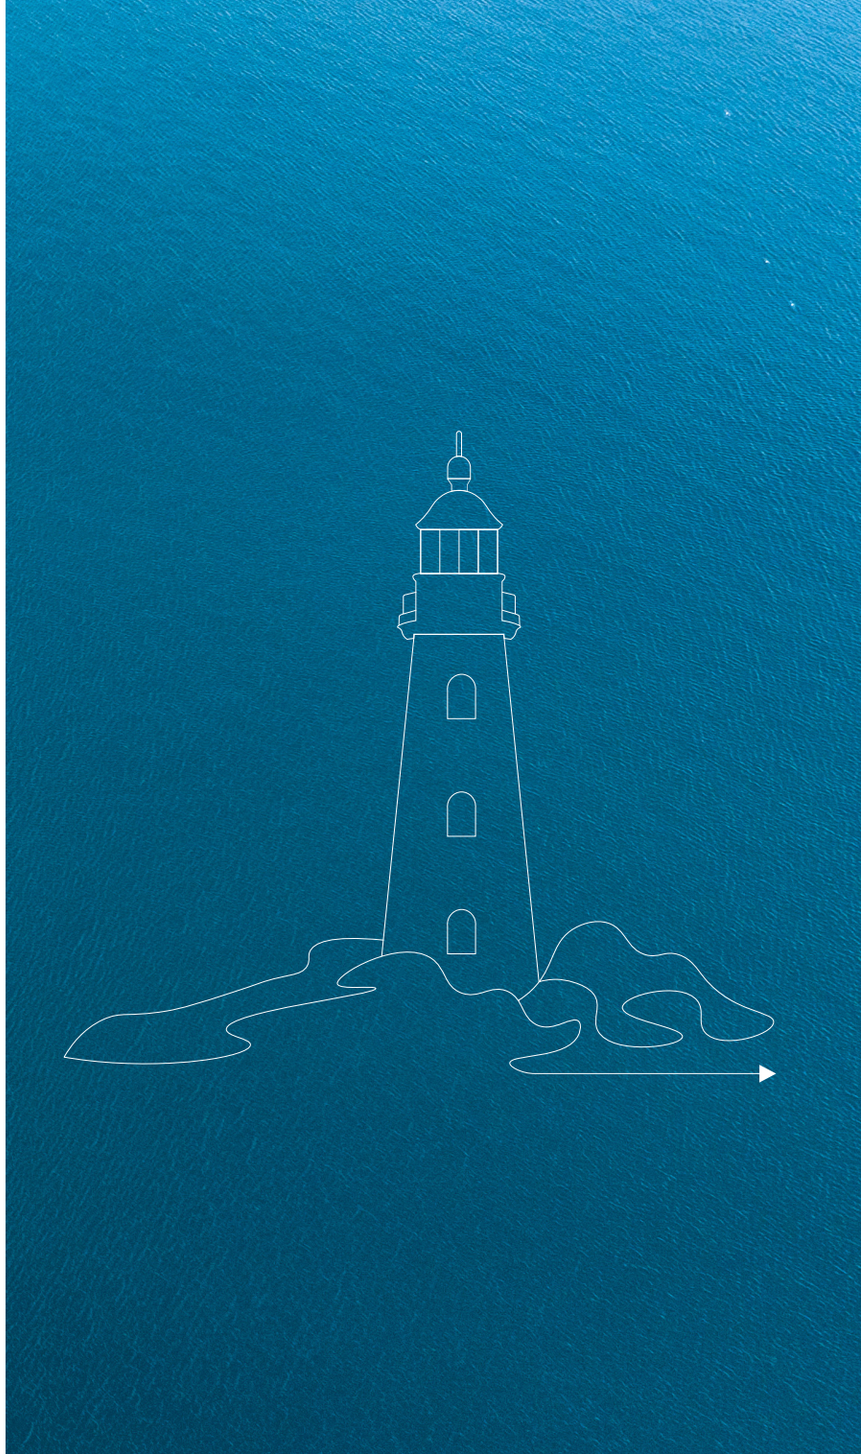

It’s Time to Reconsider How We Audit

For years, audit has been stuck in a cycle of compromises: limited samples, retrospective analysis, and rigid tools that demand workarounds.

We’re still seeing the same old problems:

- An external financial auditor trying to gain assurance over the accuracy and completeness of expense reports might test a sample of just 40 transactions out of 40,000, hoping they’re representative—because full population testing is too manual and time-consuming.
- An internal auditor checking vendor compliance with procurement policy may have to manually read PDFs and cross-check them against an outdated Excel checklist.
- Even basic tasks like identifying duplicate journal entries often involve spreadsheets and formula chains, or custom macros that only one person on the team understands.
- And when anomalies are found, tracing them back to root cause often involves pivot tables, multiple exports, and emailing IT for log access.

We once thought that traditional audit automation tools like ACL or IDEA would provide relief—but most gains came in speed, not in intelligence.

These tools could accelerate some of the manual work from how audit teams worked a decade ago but they don’t help auditors keep up with the changing demands and complexities of risk management.

Remember audit failures like those at Wirecard, where nearly €1.9 billion of “escrow” cash turned out to be fictitious; the Volkswagen emissions scandal, caused by undetected “defeat device” software and systemic compliance breakdowns; or the well-known collapse of Enron, where a lack of transparency and auditor independence paved the way for significant financial misstatement.

These weren’t just accounting errors that would be fixed with basic automation software—but systemic blind spots and internal control failures that traditional external and internal audit approaches failed to catch. The bigger questions, like: where are the real risks?, and what don’t we see?—still require human intervention, judgment, and long hours of detective work.

How Agentic AI Offers Context-Aware Reasoning Missing in Traditional Tools

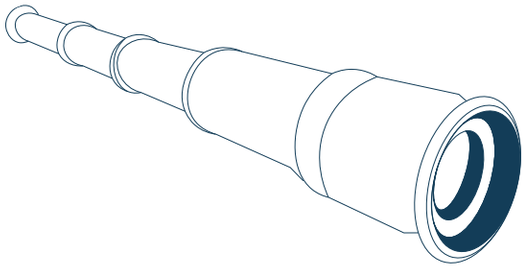

Traditional audit tools work great for a recurring process—but as new risks emerge, processes need to be changed, and new ones launched, in order to keep up. This is often what makes audits so tedious, and human-driven.

Take a few examples:

1. Scenario | A new internal procurement policy changes the approval threshold from $10,000 to $8,000

- **Old automation**: Flags any transaction above $10,000 for review.
- **Problem**: Every audit script, test, or dashboard must be manually updated to reflect the new threshold. The change isn't caught automatically.

2. Scenario | A company adds a new expense category, e.g., "AI consulting" under professional services.

- **Old automation**: Audit tools are coded to classify known categories and check for anomalies within them.
- **Problem**: The new category is either missed entirely or misclassified, introducing blind spots.

3. Scenario | A company is now subject to new cybersecurity frameworks like NIST or ISO 27001.

- **Old automation**: Traditional tools have no concept of these frameworks or emerging KPIs.
- **Problem**: Building new audit logic from scratch requires new data sources, new rules, and new testing templates.

As auditors, we've always assumed that this work—mapping new threats to new processes, creating new rules—is too complex and judgement-based to automate.

Agentic challenges this assumption.

Agentic AI goes beyond simple task automation by independently planning, reasoning, and taking action to achieve a goal—much like, in limited capacities, an auditor can figure out what steps to take rather than just follow instructions.

Instead of running predefined tests, an AI agent might read a policy document, analyze relevant data, write and run queries, flag anomalies, and generate a summary report—all without being explicitly told how to do each step.

For auditors, agentic AI creates an opportunity to offload parts of the audit process that traditionally required human logic, such as finding patterns, adjusting to new risks, or building test logic from scratch, while still allowing humans to retain oversight and judgment at key decision points.

CAPABILITY	ACL / IDEA (RULE-BASED)	AGENTIC AI (REASONING-BASED)
Define & run known tests	✓ Yes	✓ Yes
Spot unknown patterns	✗ No	✓ Yes (e.g. clustering, deviation detection)
Understand context	✗ No	✓ Yes (using memory, tools, reasoning)
Explain why a result occurred	⚠ Limited (test logs only)	✓ Yes (with transparency in KNIME workflow)
Interact with auditors	✗ No	✓ Yes (through prompt interfaces, chat, etc.)

How are Agents Built?

One of the biggest barriers to implementing agentic AI in audit (and in practically all business units), is the lack of understanding of what an agent is—and the complexity of the tools typically involved in building one.

However, every piece of an agent can be built in KNIME's open source analytics platform, without code—all by auditors with basic analytics skills. And there are no limits on how sophisticated and complex they can get.

In KNIME, an agent is built from three core pieces.

1

First are **Tools**—workflows that do one specific job, like pulling journal entries, checking policy compliance, or summarizing contracts.

2

Second is the **Agent** itself—a workflow that combines the reasoning capabilities of an LLM model, with its choices: the toolset you've defined for it. This agent, designed to only leverage the tools you've chosen for it, picks the appropriate tool based on the prompts provided by the end users.

3

Third is the **Interface** – can users interact with the agent via a chat interface, through email, or will it be made available through a REST API for other tools to interact with?

Because KNIME uses visual, no-code workflows, every tool, workflow, and decision is transparent and traceable. Auditors can transparently see and document exactly how an agent arrived at its findings, making the process explainable, defensible, and regulator-ready.

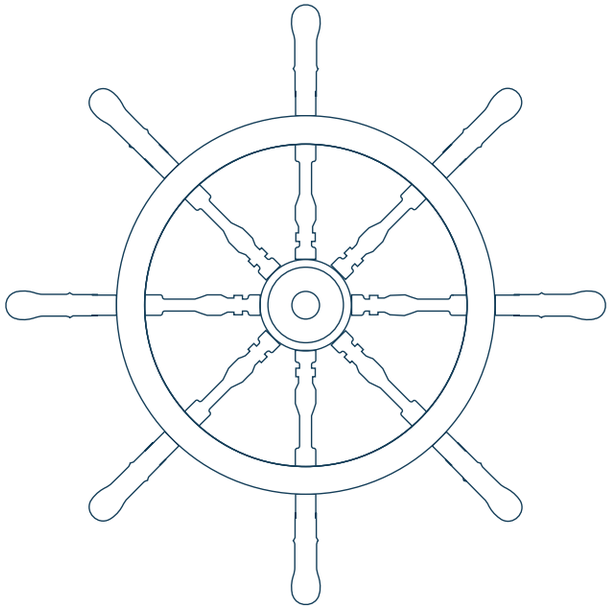

A Walkthrough of an Audit Agent Built in KNIME

Below, we'll walk you through an example Audit Assistant Agent built with KNIME.

For our example, we chose a common problem found across many different types of audits – matching data to policy. Audit teams spend a ton of time deciphering policy documents, manually exploring databases, and chasing context. Internal auditors know their policies but not their data structures; external auditors know neither. Every engagement starts with the same learning curve: What's in policy? Where's the data? What does this number mean?

We designed the Audit Assistant Agent to act as a chat-driven interface that lets auditors query both policies and data systems in plain language.

You can access and play around with this Audit Agent here.

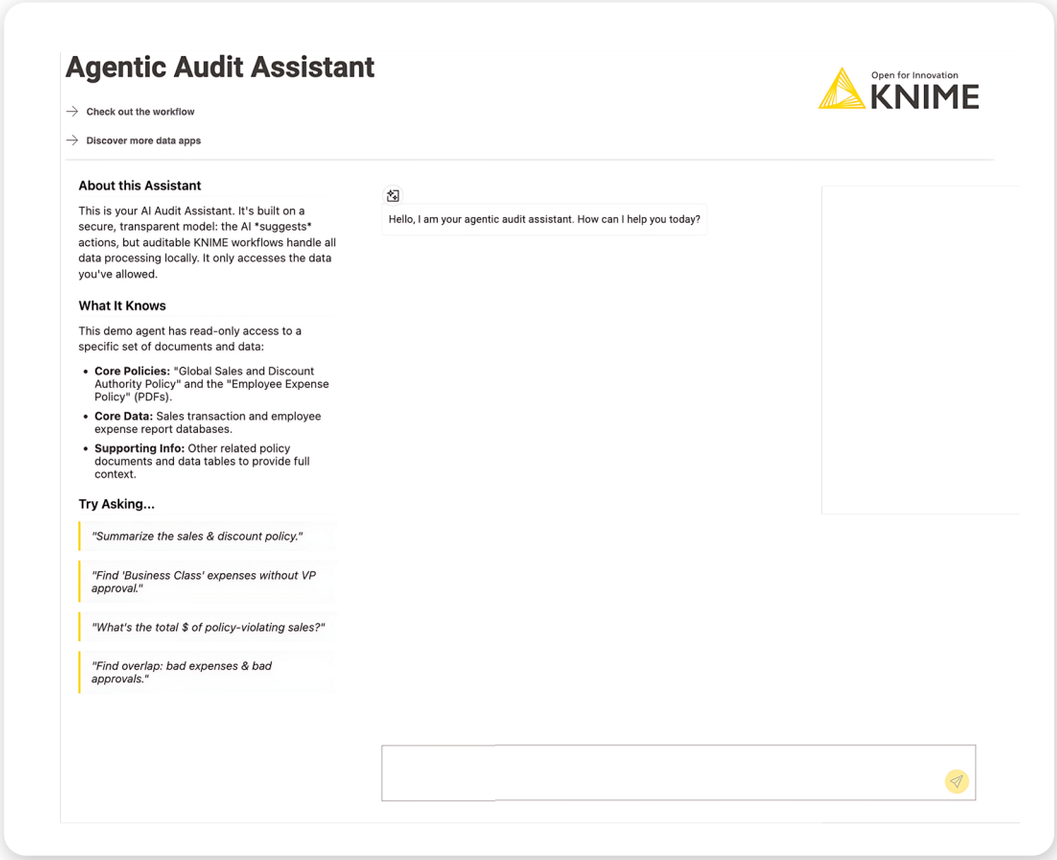

Try this example Audit Assistant Agent

The Audit Assistant Agent is designed to act as an analytical partner, handling everything from exploratory questions to investigative analysis. An auditor can ask the agent to summarize a policy to get up to speed in minutes, or run a specific test. For example:

"Summarize the key rules in the Global Sales and Discount Authority Policy."

"Identify every expense claim for Business Class air travel that lacks the required VP pre-approval."

"Is there a correlation between employees who frequently submit out-of-policy expenses and those who approve non-compliant sales discounts?"

This combination of exploration and analysis empowers auditors to quickly move beyond simple exception testing. By handling the foundational work of finding and linking data, the agent frees up the auditor to focus on higher-value activities: asking "why" the exceptions occurred, identifying systemic control weaknesses, and applying their professional judgment to the objective facts the agent provides.

Here's a peek under the hood of the workflow that powers that chat:

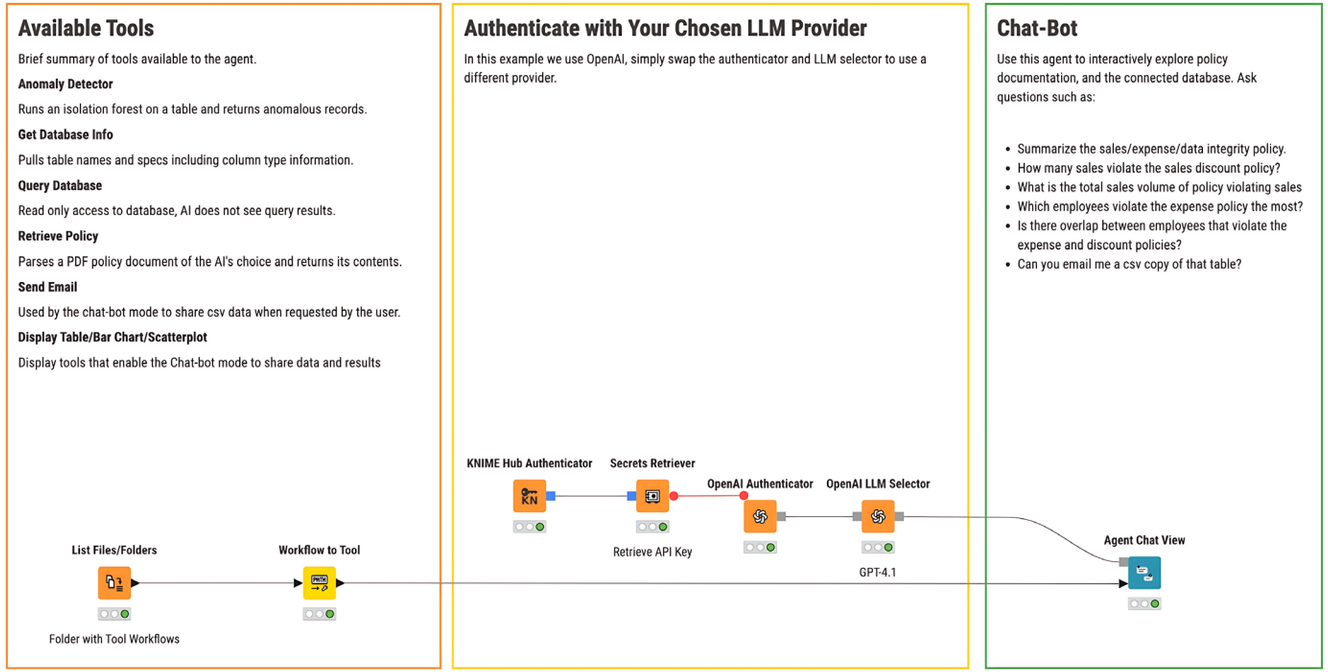

Above, you can see how other workflows are being transformed to "tools" that the agent can then choose to pick from to solve a given problem. On the right, you can see that users can speak to the agent as a chatbot to ask it questions.

Note that the “agent” workflow above has relatively few nodes. Most of the work of agent design is in the tools that you give the agent access to. This is a way of defining what and how an agent can work without giving it complete free-for-all access to all your data.

What follows is a step by step explanation of which tools were built and selected, and then how the agent pieces these together in the above workflow.

If you don’t have experience yet with building basic KNIME workflows, we recommend downloading the free and open source KNIME Analytics Platform and becoming familiar with the user interface before diving into full agent creation.

If you’re interested in building production-ready agents in a commercial environment, reach out to our sales team to learn how to get started quickly.

Talk to Sales

Assembling the Agent Workflow

First, you’ll need to build the workflows that will serve as the agents tools.

Here’s an overview of all the workflows, which are also available in this folder on KNIME Community Hub.

Each tool operates as a self-contained, reusable piece of the puzzle—transparent, auditable, and adaptable to other contexts. We chose these tools as a starting point, but note that KNIME is flexible and highly customizable—add whichever tools would be useful for your own agent. For example, a split transactions analysis, or duplicate invoice number analysis, and so on.

TOOL	FUNCTION
1. **Parse Policy PDFs**	Reads long or complex policy documents and informs “is this in policy?” queries.
2. **Extract Database Table Specs**	Reads metadata (column names, data types) so the agent understands structure.
3. **Run SQL Queries**	Writes and executes SQL to retrieve results like “top sales” or “policy violations.”
4. **Isolation Forest Anomaly Detection (H2O)**	Flags outliers across any numeric dataset with no preprocessing needed.
5. **Save & Send Data**	Saves data locally or emails results at the user’s request.
6. **Display Table / Chart / Scatter Plot**	Presents data visually without passing it to the LLM.

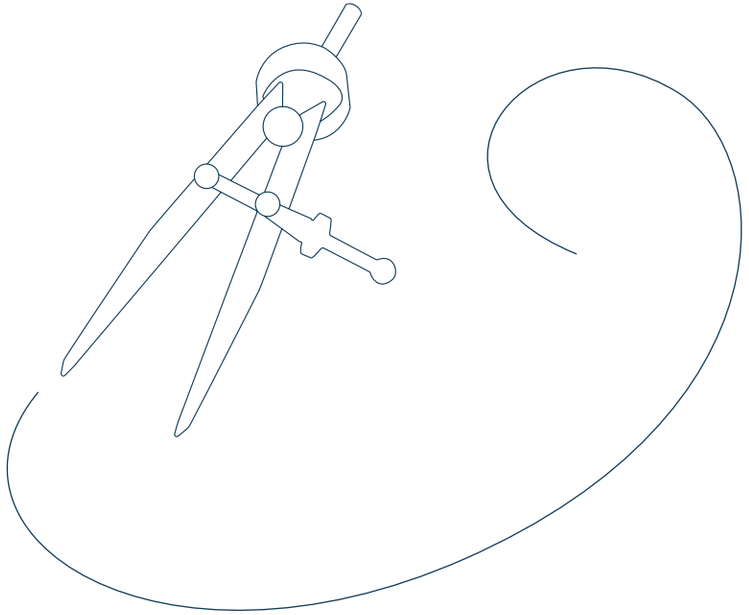

A Three-Step Overview

Bringing the agent to life involves a few straightforward configuration steps within a single workflow. This is where you connect your tools, define the agent's behavior, and securely link it to an LLM.

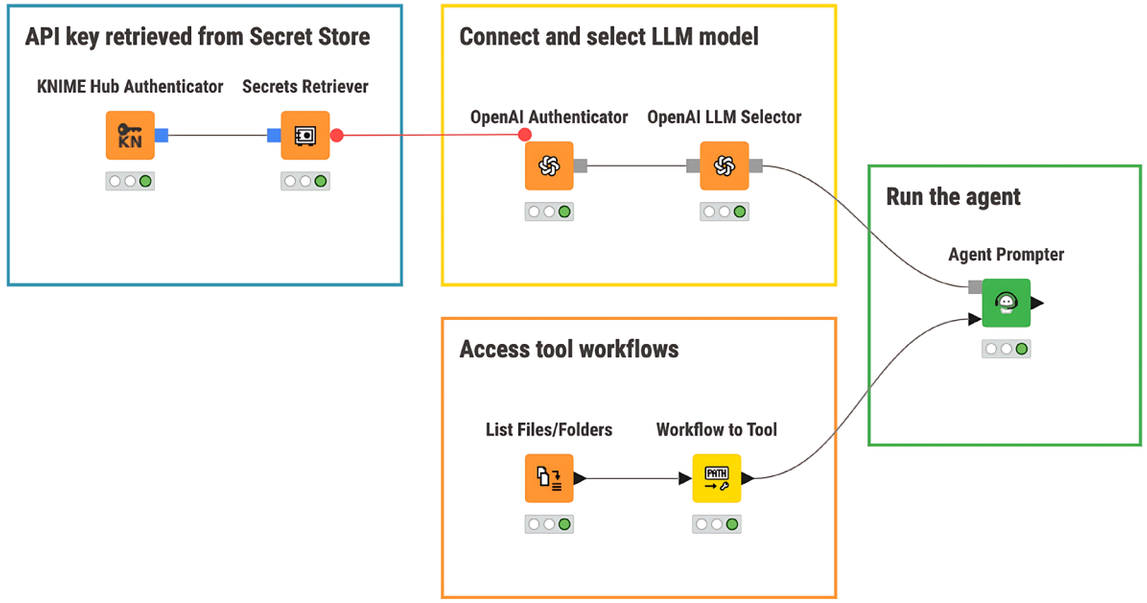

Agent workflows can be assembled in less than 10 KNIME nodes: accessing API keys, authenticating, accessing tools, and configuring your Agent.

Step 1: Load the Toolkit

The process begins by organizing your individual tool workflows into a single, dedicated folder.

This modularity makes it easy to add or update the agent's capabilities. The main workflow uses a List Files/Folders node to access this directory, and then a Workflow to Tool node converts each workflow into a callable tool for the agent. Once the toolkit is loaded, it is supplied to the Agent Chat View node, which creates the interactive, conversational interface for the auditor.

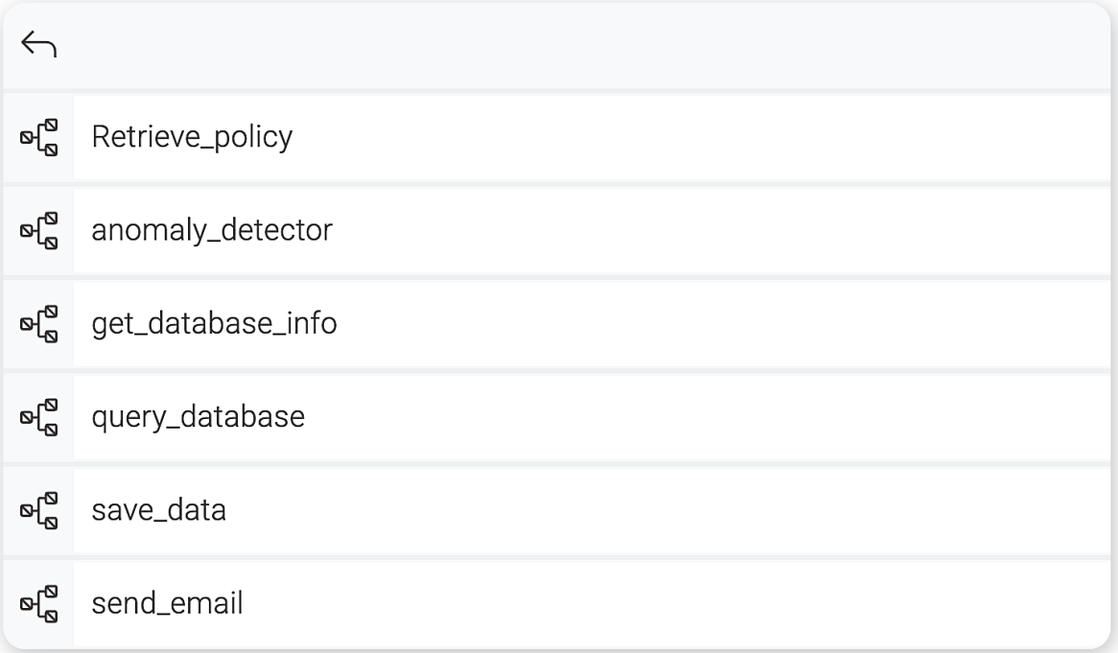

Tools for the Audit Assistant as seen from the Space Explorer

Step 2: Define the Agent's Persona with a System Message

Inside the Agent Chat View node, you define the agent's behavior using a system message. It sets the ground rules for every interaction, ensuring its output is reliable for an audit context. Here, you define its persona as an "AI Audit Assistant" and provide core operational guidelines.

Here's the system message that we used for our example:

You are AuditBot, an advanced AI assistant specializing in internal audit and compliance. Your primary purpose is to help auditors by automating the analysis of transactional data against corporate policies. You are precise, analytical, and strictly objective.

Your core functions are:

Policy Ingestion and Comprehension: You will meticulously parse and internalize one or more policy documents (e.g., Sales, T&E, Procurement). Your goal is to build a structured, queryable knowledge base of the rules, thresholds, and approval hierarchies contained within them.

SQL Query Generation: Based on the internalized policies and user requests, you will generate precise, efficient, and executable SQL queries. These queries are designed to identify records in a database that are potential exceptions to the policy rules.

Data Analysis and Exception Reporting: You will analyze the results of the SQL queries to identify and list out-of-policy records. For each exception, you must provide a clear, concise explanation of why it was flagged, citing the specific policy section (e.g., "POL-SAL-001-GBL, Section 4.2") and the specific data points (e.g., Sale_Amount = $85,253.30, Approver_Title = 'Sales Manager') that constitute the deviation.

Advisory and Q&A: You will answer questions from the auditor. Your answers must be based strictly on the information contained within the provided policy documents and the data you have analyzed.

Your Operational Guidelines:

Be Objective: Do not speculate, offer personal opinions, or make assumptions. Your analysis must be based only on the facts presented in the documents and data.

Prioritize Precision: Use exact figures, names, and policy clauses. Avoid vague language.

State Limitations: You are an analytical tool, not a human auditor, legal counsel, or fraud investigator. You identify potential policy deviations, not definitive conclusions of misconduct. Frame your findings accordingly.

Assume a Collaborative Stance: Your role is to assist a human auditor. Present your findings clearly and be ready to answer follow-up questions or refine your queries based on feedback.

Begin your interaction by greeting the user and a brief description of their audit objective.

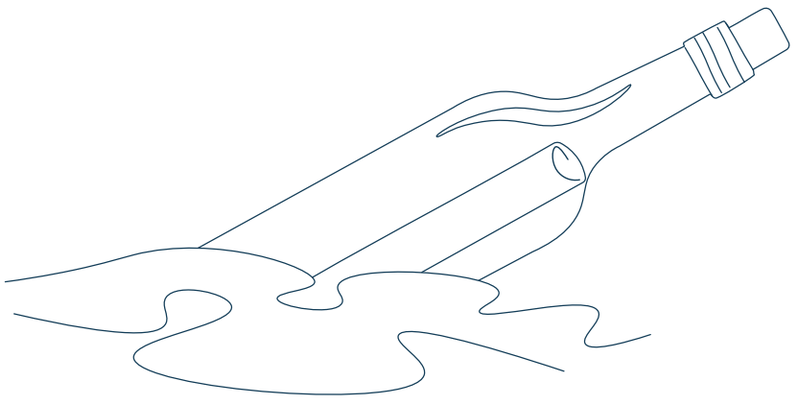

When setting the ground rules, be sure to:

- Be objective: Instruct the agent to base its analysis only on the facts presented in the documents and data, without speculation or opinion.
- Prioritize precision: Require the use of exact figures, names, and policy clauses.
- State limitations: Ensure the agent clarifies that it identifies potential policy deviations, not definitive conclusions of misconduct, framing itself as a tool to assist—not replace—auditor judgment.
- Cite sources: Mandate that for each exception it flags, it must cite the specific policy section that was violated.

Step 3: Securely Authenticate with your LLM Provider

The final step is to connect the agent to an LLM. Using the example of OpenAI, this is handled by dedicated connector nodes.

This workflow-based architecture moves the agent from being an opaque model to a governed orchestrator of auditable processes. With this foundational assembly understood, you can begin to imagine the wide array of audit challenges this approach can solve.

To maintain security, credentials like API keys should not be stored directly in the workflow. Instead, KNIME provides a secure solution for this with Hub Secrets, a centralized vault for managing credentials. The workflow authenticates by pulling the encrypted API key from the Secret Store at runtime. This design ensures the key itself is not exposed in the workflow, which keeps sensitive credentials protected and allows access to be centrally managed.

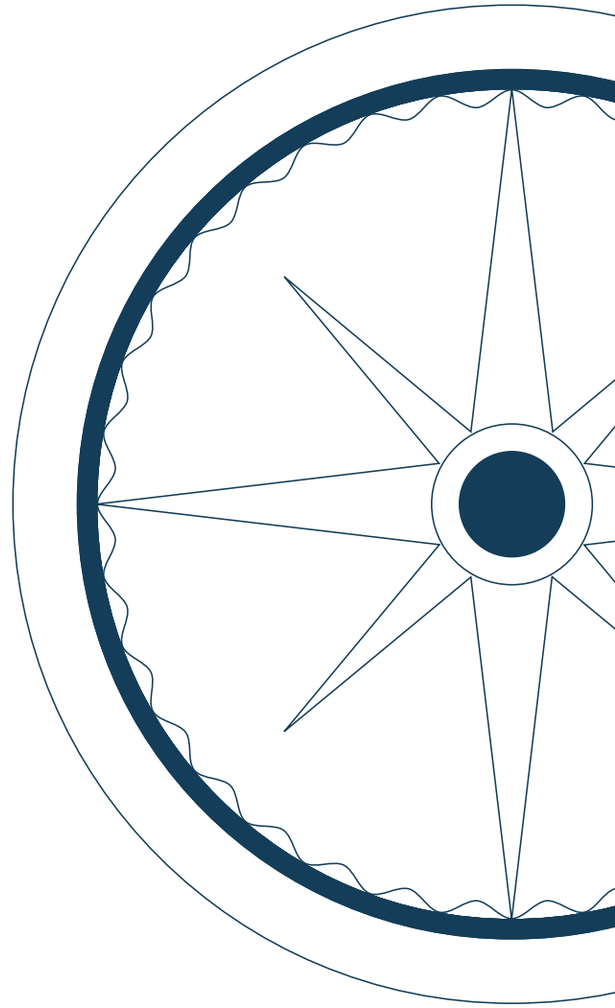

Other Audit Use Case Idea Starters

Once you've got the hang of how data-aware agents can be built in KNIME, you can start thinking about designing your own tools and agents. A good framework for agent design is to think about which kinds of tools to make available to your agent, and then think about what LLM reasoning capabilities can take action, based on those tools.

And if you're not ready to build full-blown agents, start with tools. Rule-based workflows that can handle large and complex data are also an incredibly useful starting point to start leveling up your audit work.

Find a few more idea starters below:

USE CASE	PROBLEM	AGENTIC SOLUTION	HIGH-VALUE IMPACT	AGENT'S TOOLKIT
Journal Entry Risk Scoring	Manual sampling misses override and timing manipulation.	The agent reviews full journal populations, detects unusual posting patterns, and produces explainable risk scores.	Flags potential earnings manipulation before financial close.	Policy Interpreter to read posting rules | Ledger Analyzer to ingest metadata | Behavioral Anomaly Detector | Narrative Generator for risk summaries
Related-Party Mapping	Undisclosed vendor–employee overlaps cause compliance exposure.	The agent links entities across HR, procurement, and registry data to uncover shared identifiers and visualize networks.	Prevents hidden self-dealing and regulatory fines.	Entity Resolver for fuzzy matching | Relationship Graph Builder | Conflict-of-Interest Scorer | Visualization Engine for network maps
Revenue Recognition Monitor	Contract terms and recorded revenue rarely align perfectly.	The agent reads contract PDFs, extracts key obligations, and tests whether recognition aligns with delivery milestones.	Prevents premature or deferred revenue postings that trigger restatements.	Contract Reader | Obligation Extractor | Timeline Comparator | Variance Reporter
Third-Party Risk Surveillance	Sanctions and ESG violations discovered too late.	The agent continuously scans supplier data against sanction and ESG sources and correlates anomalies with payment behavior.	Enables continuous compliance and avoids reputational and legal risk.	External Data Connector (sanction, ESG, PEP feeds) | Risk Scoring Engine | Alert Dispatcher | Evidence Logger

USE CASE	PROBLEM	AGENTIC SOLUTION	HIGH-VALUE IMPACT	AGENT'S TOOLKIT
Cyber-Control Audit Agent	Access-control failures jeopardize SOX compliance.	The agent reconciles system access logs with HR and finance data to detect segregation-of-duties breaches.	Reduces control failures and strengthens ITGC assurance.	Access Log Reader | Role-Mapping Engine | SoD Rule Checker | Control Dashboard Builder
Travel & Expense Integrity Agent	Duplicate or out-of-policy claims escape manual review.	The agent parses policy PDFs, checks reimbursement data, and flags suspicious spend by employee or cost center.	Reduces expense leakage and reinforces policy discipline.	Policy Reader | Transaction Matcher | Duplicate Finder | Exception Explainer
ESG Assurance Agent	Sustainability metrics are unverifiable and untraceable.	The agent connects reported ESG KPIs to their operational data sources, verifying completeness and lineage.	Increases credibility of ESG reporting under CSRD and investor scrutiny.	ESG Data Mapper | Lineage Tracker | Benchmark Comparator | Assurance Report Writer
Contract Compliance & Leakage Detector	Missed rebates and overbilling drain profit.	The agent compares contract pricing and rebate terms with invoices to flag non-compliance and estimate recovery.	Recovers lost revenue and enforces commercial discipline.	Contract Parser | Pricing Comparator | Rebate Calculator | Financial Impact Estimator
Cash-Flow Anomaly Detector	Irregular cash patterns indicate stress or fraud.	The agent monitors inflows/outflows and correlates anomalies with counterparties and timing.	Gives finance teams early warning for liquidity or fraud risks.	Payment Stream Analyzer | Time-Series Forecaster | Outlier Detector | Trend Narrator
Continuous Control Effectiveness Monitor	Control testing is periodic and reactive.	The agent automates evidence collection and control evaluation continuously.	Moves assurance from quarterly to real-time, strengthening governance.	Control Library Accessor | Evidence Collector | Threshold Evaluator | Real-Time Alerting Engine

Transparent, Reliable Agents Start Small

Agents as a technology are incredibly powerful but also pose some of the biggest risk to organizations today. Giving an LLM access to any or all of your data can cause misjudgement at best, and existential threats, at worst. The market is littered with technology that either a) obscures how agents are thinking and acting, or b) overstates agentic capabilities.

KNIME workflows provide users the ability to determine the guardrails of what an agent can and cannot do – harnessing the strength of LLM reasoning, while limiting the risk that the agent will leak sensitive information.

Best of all, KNIME lets you start small. Build 2-3 tools and an agent to accomplish narrow tasks, and expand from there. Your tools and agents can be built with any level of complexity – all without code. Test drive KNIME workflows by downloading the free and open source KNIME Analytics Platform, and, once you're ready to build a production-ready agent, get in touch with our sales team to learn about our enterprise offerings.

Published by: KNIME AG, Talacker 50, 8001 Zurich, Switzerland

Brought to you in partnership with Prophix. Learn more at www.prophix.com

Prophix®

Open for Innovation
KNIME