Open for Innovation
KNIME

Your roadmap for building a GenAI governance framework

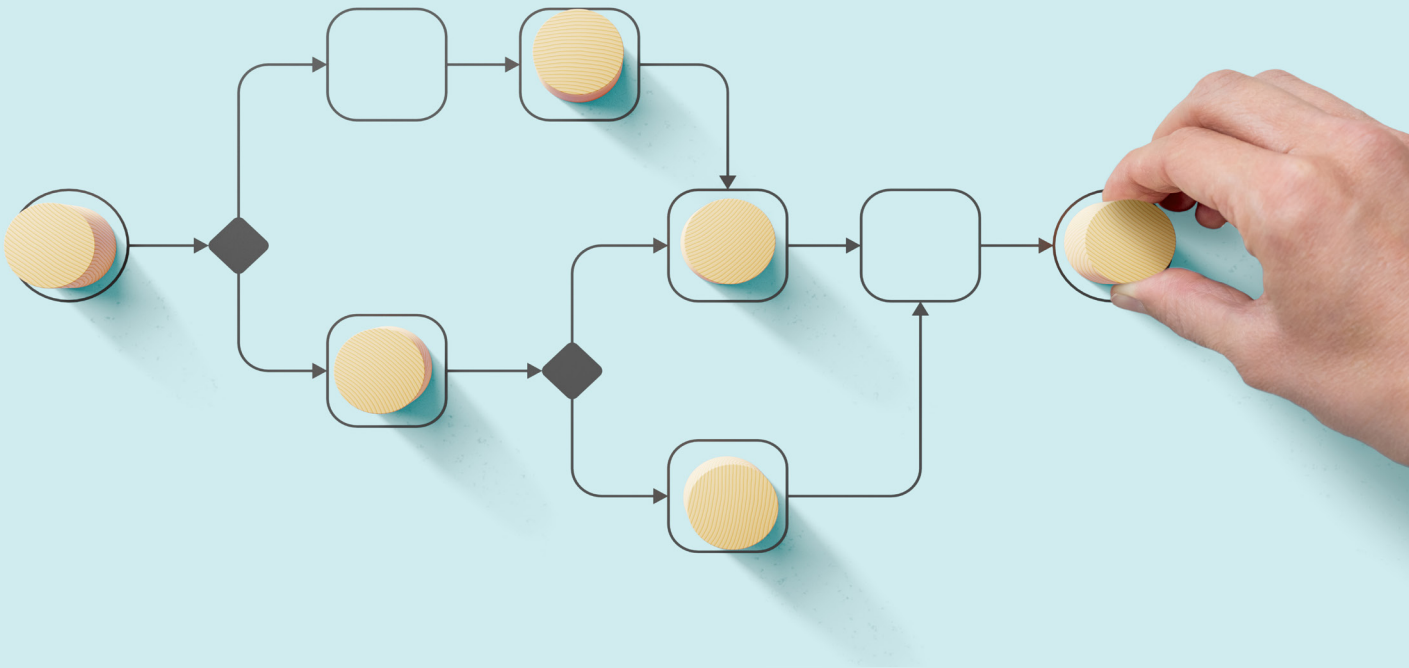

Creating a comprehensive governance framework for GenAI in your organization may seem like a daunting task. But it doesn't need to be. This 8-step roadmap breaks it down into manageable steps so you can easily take your company from GenAI chaos to a secure and structured operating framework.

What's more, you can present this roadmap to stakeholders to transparently communicate what the process of crafting a governance framework looks like – ensuring everyone is aware of the process, knows what needs to be done, and understands their role in it.

This roadmap has been designed to make sure you achieve all the following things in your GenAI governance framework:

- Risk factors are comprehensively identified and documented
- Identified risk are clearly mapped to internal controls
- Legal requirements are complied with
- Identification and creation of needed data governance policies
- Usage of GenAI tools is aligned with existing internal policies
- GenAI is safely and effectively rolled out to teams

1. Define goals and scope

Identify Goals: What do you want to achieve at the end of this process? Create a concrete set of goal statements to give your process focus. Consider what teams or functions should be governed by your framework and the expected impact the framework will have. Having a clearly defined end state helps you maintain focus and creates clarity for the process. For example:

- Have a clear and up to date data governance approach that is clearly understood across the organization.
- Have a documented rubric for evaluating which use cases or activities are in scope and which are out of scope for GenAI usage.
- Have a documented hierarchy on who has access to which GenAI tooling at which level of the organization and a plan for maintaining it.
- Create a GenAI ethics committee and schedule recurring meetings and an agenda for these oversight sessions.
- Implement internal policies to make GenAI usage safe and secure.

Set the Scope: Define in which departments GenAI will be applied (e.g., data science, customer service, product development, marketing) and whether your governance framework should cover the whole organization or specific teams. This will impact how much risk mapping you need to do.

Document Priority Use Cases (Optional): You may or may not know how you want to implement GenAI in your organization. If you have already identified specific use cases, this can make your risk assessment process more precise and watertight. For example, if I know I want to use GenAI to clean data, then I need to bear in mind risks around managing personally identifiable information.

2. Establish governance stakeholders

Governance Committee: Form a GenAI governance committee comprising stakeholders from IT, legal, compliance, HR, and relevant business units, so you have diversity of thought and a broad perspective on risk factors and ways to mitigate them. When forming a GenAI governance committee, we recommend including at a minimum the following roles or proxy roles:

- Chief Information Officer (CIO)
- Chief Data Officer (CDO)
- Head of IT
- Legal Advisor/Counsel
- Data Security Officer(s)
- Head of Compliance
- Non-Executive Director (or other independent advisor)

Roles and Responsibilities: Define roles and responsibilities on your governance committee. In particular to understand which roles are chiefly responsible for pushing which of your goals forward. For example, based on the goals above, we can assign them to the most relevant members of your committee. This creates oversight and accountability. Here are some examples:

- **Goal:** Implement internal policies to make GenAI usage safe and secure
  - **Owner:** Head of IT

- **Goal:** Have a clear and up to date data governance approach that is clearly understood across the organization.
  - **Owner:** CDO and CIO

Set a Check-in Cadence: Define your cadence of meeting up and discussing upfront and get it booked in the calendar. Having a regular check-in to report on progress, unblock each other, or discuss thorny issues helps move things forward and create accountability.

3. Review your existing data governance framework

Data Governance Framework Review: Ensure robust data governance practices are already in place, including data lineage, metadata management, and access controls. Identifying and solving any weaknesses in your current data governance process is the first step in considering how to manage GenAI risk.

Strengthen Governance for GenAI: Consider what new governance structures you may need to put in place to use GenAI safely. For example, how may your data lineage, metadata management, and access controls need to change in light of this technology.

Data Naming Conventions: Review and ensure you have clear and consistent data naming conventions. Without this, it will be difficult to communicate internally about what data types, data sources, or data processes are being governed by what. This may sound simple, but it's easy to overlook. Miscommunication can cause a lot of problems when we lack a shared data language.

4. Identifying risks in GenAI implementation

Identifying risks is a critical component of a governance framework for GenAI. You want to make sure potential issues are recognized and mitigated proactively.

We recommend listing out all your risks in a spreadsheet as you go along, so you can map internal controls to specific risks.

Data privacy and security

Assess risks related to the exposure of sensitive data, data breaches, and compliance with data protection regulations (e.g., GDPR, CCPA). Guiding questions:

- What sensitive data do we use that needs specific GenAI-related risk management?
- How robust is our existing general data governance framework to prevent GenAI related risk?
- What data is and isn't allowed to interface with GenAI tooling?

Model accuracy and reliability

Identify risks related to the accuracy and reliability of AI models, including overfitting, underfitting, and generalization errors.

Guiding questions:

- What are the risks of GenAI inaccuracy and how big are they?
- What's an acceptable range of accuracy that we are happy with?
- Where is the accuracy and reliability risk too high or pervasive to be managed?
- How do we verify accuracy and reliability for model output?
- What testing (if any) do we require to ensure a model or workflow can deliver accurate results.

Bias and fairness in algorithms

Assess risks of inherent biases in algorithms that could result in unfair treatment of certain groups.

Guiding questions:

- What particular bias or fairness risks do we need to control against given our likely GenAI use cases?
- What testing (if any) do we require to ensure our GenA) usage is not reinforcing harmful biases?
- Which specific use cases are most risky when it comes to bias and fairness?

- How do we identify signs of bias or unfairness in our GenAI usage and how can we effectively scale that process?

Copyright and intellectual property

Are the models you plan to use in any active legal battle around copyright infringement that could pass risks onto your company.

Guiding questions:

- What are our organization's criteria for picking GenAI tooling regarding copyright?
- How can we make sure we avoid GenAI tooling that is under legal contention for copyright infringement? **(see this helpful tracker)**
- What is our stance on intellectual property of work created by or with GenAI tooling?
- How can we reduce the risk of copyright infringement that could backfire on us?
- How can we reduce the risk of internal intellectual property leaking to the public via GenAI prompting?

Cost control

Identify risks related to budget overruns in the development and deployment of GenAI systems.

Guiding questions:

- How do we mitigate the risk of uncontrolled costs?

- What warnings or systems can we put in place to alert us early to high expenditure/request volumes?
- How do we mitigate the risk of bad actors flooding requests to cause financial damage to us?
- What is our AI budget?

System integration

Identify risks associated with integrating GenAI systems with existing IT infrastructure and business processes.

Guiding questions:

- Which systems or databases need to interface with GenAI tooling to fulfill our use cases?
- Which systems or databases should GenAI be prevented from interfacing with?
- How deeply embedded in systems should GenAI be?
- What are the potential downstream effects in terms of accuracy, bias, or cost control of integrating GenAI into systems?

Performance and Dependence on AI Systems

Assess risks related to over-reliance on GenAI systems, including potential failures and downtime.

Guiding questions:

- What are our criteria for ruling in/out GenAI for a given use case?
- For processes run with GenAI, what's the business or reputational risk of something not working?
- What performance-related obligations could be threatened by GenAI usage (for example, up-time guarantees with paying customers or any core product functionality that begins to malfunction).
- What level of dependence on AI systems do we want to have?

5. Develop controls, policies and guidelines

Build an Internal Controls Matrix: Match identified risks to internal controls based on the questions above and ensure this is documented. Many of the risks you have identified may already be covered by your existing data governance framework. But there will be others where you need to add specific additional controls.

Develop Controls for GenAI-Specific Risks: Document these in the same place, as these will form the basis of the specific controls you need to effectively govern GenAI usage at scale.

Note Unmanaged Risks: Make sure to mark up risks that you are not sure how to manage. Knowing which risks cannot be effectively managed at scale is important, because this will put some important limits on GenAI use cases.

6. Implement governance initiatives

Select and Roll Out Your Models: Based on your risk analysis and review of tooling, choose and set your approved internal tools, which may be a mixture of external and internally-built tools. Document this and put in place access limitations for unapproved tooling.

Implement New Controls: During the controls brainstorm you will have come up with a number of new controls to implement. Now's the time to put them into action. Ideally a few people in your governance committee can take responsibility for this – for example the CIO, CTO or CDO with colleagues.

Modify Access Controls: Depending on your level of system integration, you may need to review access controls detailed in your general data governance framework.

Choosing GenAI providers: Your choice of GenAI providers also has governance implications. Not only do you need to implement who can access what models, for what, and at what level of authority. But it is also crucial that your GenAI providers are chosen wisely and monitored carefully. Do some GenAI providers pose an increased risk around data privacy and intellectual property compared to others? Implementing governance initiatives runs alongside implementing tools. GenAI tooling needs to support your governance initiatives.

Now your GenAI governance framework is documented and implemented. The following sections focus on important steps to take after this work is done, including fine-tuning your approach, ensuring oversight, and educating your employees.

7. Training and awareness

Employee Training: Develop and implement training programs for employees on approved GenAI use cases, risks, and governance. Also be sure to train staff on how to report breaches or leaks and press upon them the importance of doing so.

Approved Models: Make sure to communicate which GenAI tooling or models are allowed and which are prohibited.

Use Case Manual: Provide a list of approved use cases that teams can refer back to and create a process for people to suggest and get new use cases approved.

8. Ethical review and oversight

Ethics Board: Establish an AI ethics board to review and provide guidance on GenAI projects on an ongoing basis. This could be the same governance committee you put together at the start of this project. For added accountability, consider adding independent members to this board.

Ethical Audits: Conduct regular ethical audits of GenAI applications. You may want to do this on a more frequent basis during your first few months rolling out GenAI.

Ombudsperson: Designate a person or team of people to manage information breaches or issues that may come up, or to support whistleblowers who have seen unethical/out of policy GenAI use in your organization. Make sure you have a process for documenting GenAI related leaks or issues.

Feedback Mechanism: Implement feedback loops to gather insights from users and stakeholders about how scaling GenAI is going, and what's working and not working.

Periodic Reviews: Conduct periodic reviews of the governance framework to ensure it remains effective and up-to-date. During these reviews, make sure to review the ombudsperson's list of GenAI related leaks or issues and use this feedback to fine tune your GenAI governance framework.

By following this detailed checklist, an organization can establish a robust and effective governance framework for the responsible and strategic use of GenAI.

How KNIME's AI Gateway supports your GenAI governance journey

Identifying risks is the easy part of any governance process. Finding effective and practical controls that can be implemented is the difficult part. That's why we've built a comprehensive set of GenAI-specific governance tools into KNIME, so you can scale up GenAI in your data practice while driving down risk.

In case you don't know, KNIME is a data science tool that makes it easy to access, blend, transform, model and visualize data. Whether you're a data analyst, data scientist, or sitting in the c-suite, KNIME makes all your data processes clear and auditable, and creates a central hub for all your data-related workflows and processes.

KNIME's commercial software KNIME Business Hub gives you access to a suite of customizable controls to help you gatekeep, guardrail, and audit GenAI functionalities. In case you're at a loss for how to implement controls we've discussed in this roadmap, here are some of the features on offer within KNIME Business Hub's AI Gateway:

GenAI routing control: Ensure all GenAI requests are routed through your approved provider – whether those are internally built or external providers like OpenAI, Azure OpenAI Service, Hugging Face, and GPT4ALL.

Access controls: Admins can set who gets access to GenAI, including which models can be accessed at which seniority level or by team membership.

PII anonymization: Even when we're careful with personally identifiable information, we are still subject to human error. You can set up an internal control in KNIME so workflows anonymize or remove PII before any data is sent to GenAI tooling.

Detect and mitigate bias or inaccuracies: Get access to many KNIME extensions to check for and resolve bias and hallucination. For example, with KNIME's **Giskard.ai** integration, you can check the robustness, accuracy, or any bias in your end-to-end machine learning models with just a few clicks.

Guard-railing workflows: Prevent any other sensitive or competitive information being accessed by GenAI tools with KNIME "guard-railing" workflows that block requests including sensitive data.

Monitor GenAI cost and activity: Benefit from a KNIME data app to analyze AI activity and usage history, so you can stay on top of usage and the cost of your GenAI solutions.

Learn more about all of **KNIME's GenAI capabilities** here.

Interested in finding out how KNIME can support your GenAI roll-out? **Get in touch here.**