A Framework for Successfully Rolling Out GenAI to Thousands of Employees Securely

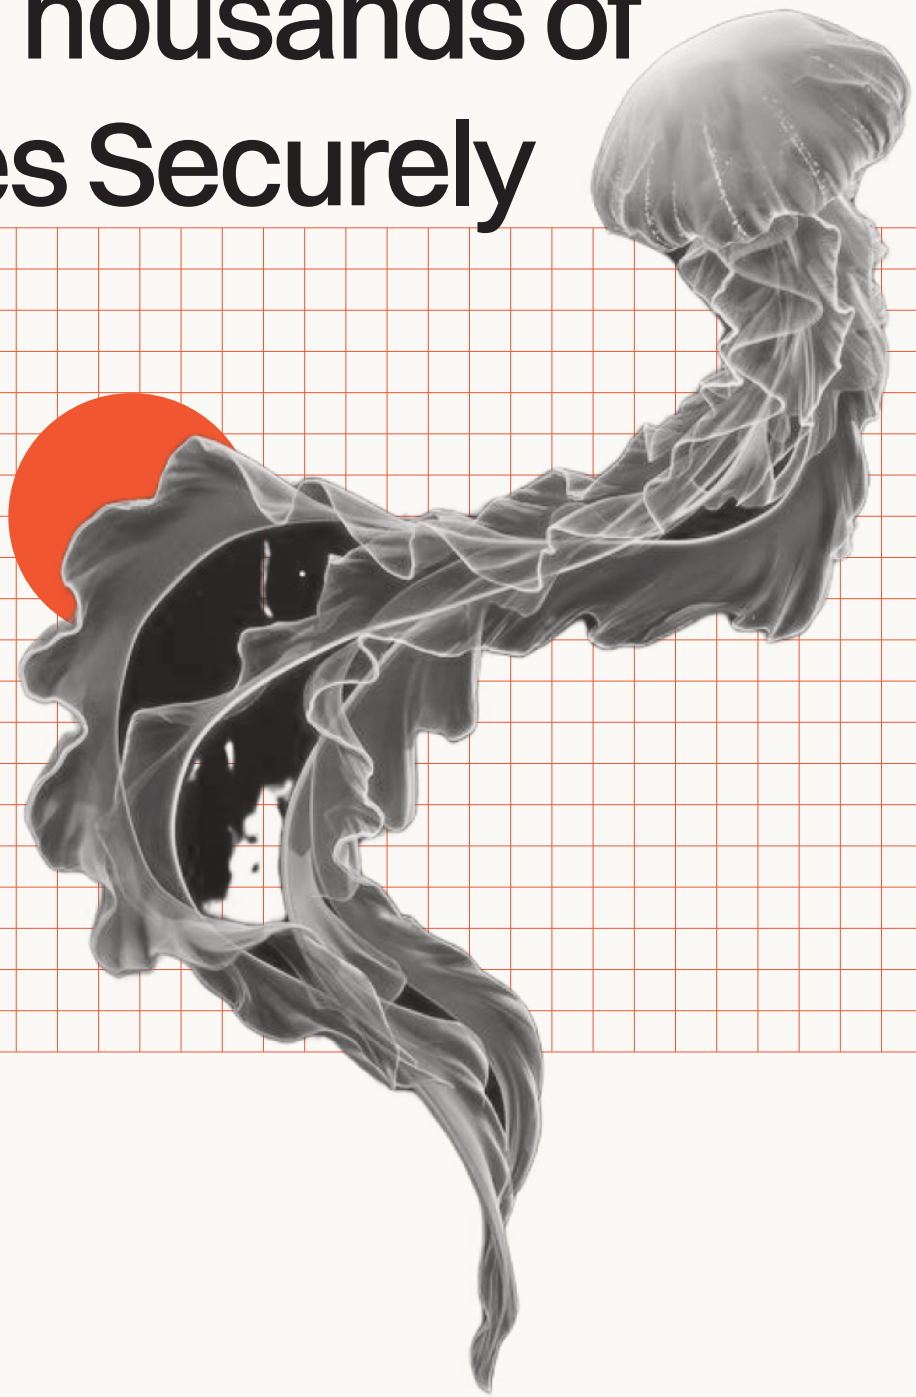

COO Guide

Redactive

Table of Contents

Why it's Time to Embrace Generative AI 04

Our Tested Framework for Rolling Out AI at Enterprise Scale 05

1. Define a Native AI Use Case and Clear ROI Metric 06

2. Factors to Consider on the Build versus Buy Decision for GenAI 07

3. Maintaining Momentum by Establishing a GenAI Taskforce 09

4. Knowledge Base Audit: The Foundation of a Secure GenAI Implementation 12

5. Build an Early Adopter Program and Feedback Loop for your first use-cases 17

6. Scaling GenAI Across your Enterprise to Achieve a Competitive Advantage 19

Why it's Time to Embrace Generative AI

It's time to embrace Generative AI (GenAI) across your business. GenAI represents a fundamental shift in how technology interfaces with business operations. Unlike traditional AI systems that were primarily confined to data analysis tasks for specialized teams, GenAI directly accesses organizational knowledge bases and creates customer-facing outputs. This expanded scope and capability brings both tremendous opportunities and significant risks that must be carefully managed.

AI adoption worldwide has increased dramatically in the past year, after years of little meaningful change.

Organizations that have adopted AI in at least 1 business function,[1] % of respondents

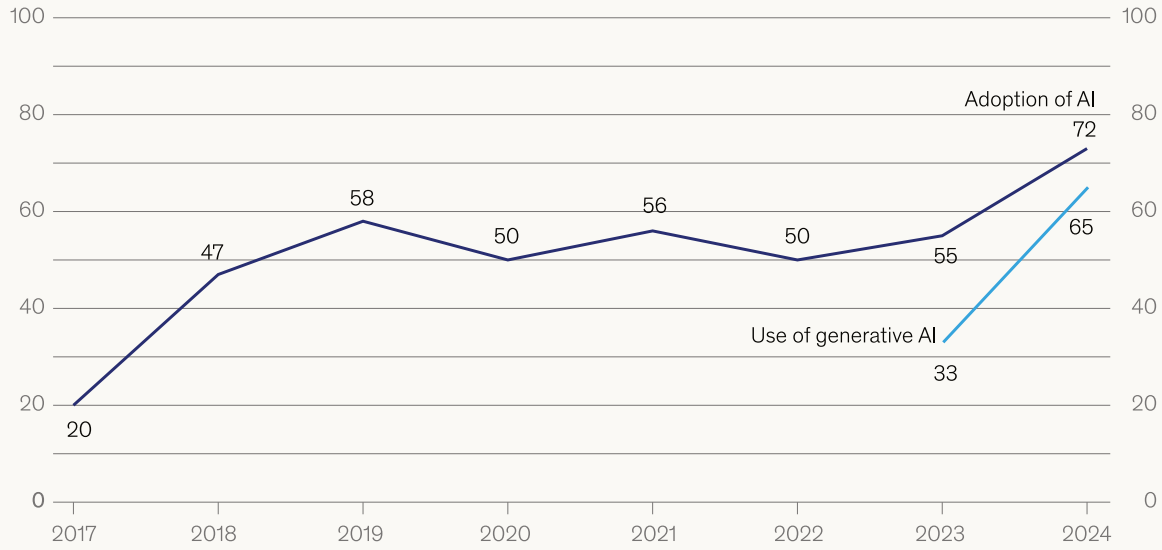

[1]In 2017, the definition for AI adoption was using AI in a core part of the organization's business or at scale. In 2018 and 2019, the definition was embedding at least 1 AI capability in business processes or products. Since 2020, the definition has been that the organization has adopted AI in at least 1 function.
Source: McKinsey Global Survey on AI, 1,363 participants at all levels of the organization, Feb 22–Mar 5, 2024

McKinsey & Company

Chief Operating Officers (COOs) find themselves at the center of this transformation, tasked by CEOs and boards to implement GenAI as a strategic imperative while ensuring robust risk management and demonstrable return on investment (ROI). As early adopters of Microsoft Copilot and Glean gain maturity and experience leveraging these systems to provide value there are many insights we can uncover as to how to successfully rollout GenAI at scale.

As GenAI reshapes the business landscape, organizations face unprecedented pressure to adopt and integrate these technologies while managing associated risks. Recent data from shows GenAI adoption has almost **doubled from 33% to 65% in just twelve months**[1], making it a critical priority for executive leadership. This whitepaper presents a comprehensive framework for COOs to successfully implement GenAI while addressing security, compliance, and operational challenges.

Our Tested Framework for Rolling Out AI at Enterprise Scale

As an AI first company Redactive has a front seat view of how the industry is embracing GenAI. Working with our partners and customers has provided us with prescriptive feedback about how a company can flourish when it embraces GenAI from a whole-of-business perspective.

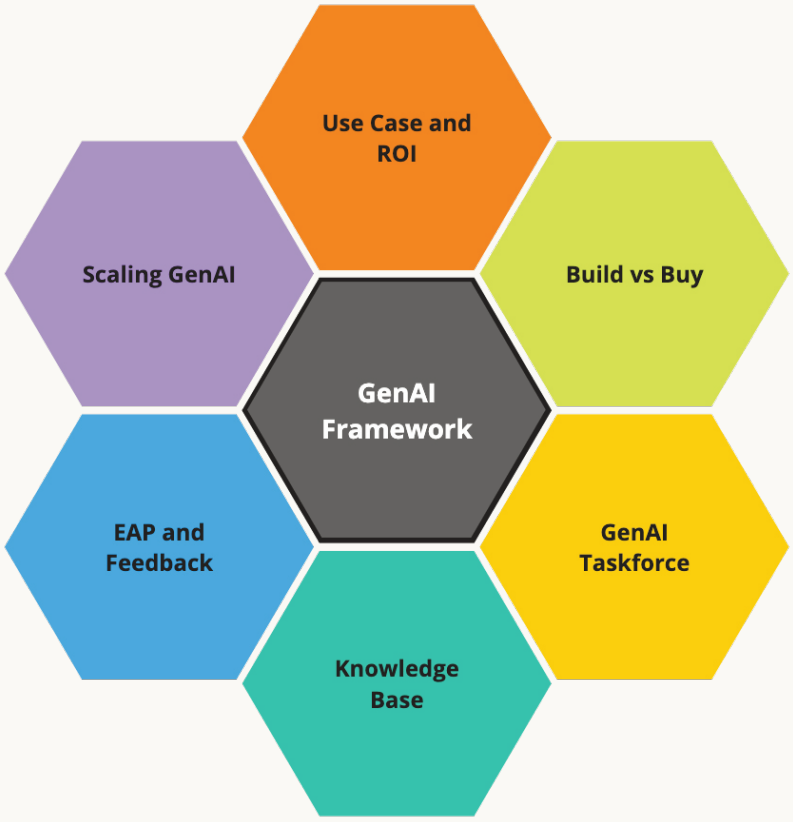

This framework provides the foundational components needed to build GenAI for your business right from the start.

[1] Analysis conducted by *McKinsey and Company* as found in The state of AI in early 2024

1. Define a Native AI Use Case and Clear ROI Metric

The foundation of successful GenAI implementation lies in **clearly defined use cases with measurable outcomes**. Organizations frequently rush into GenAI implementation driven by competitive pressure or executive mandates without establishing clear objectives, resulting in costly missteps and failed initiatives.

Use cases are important for teams to understand the end state as to what they are aiming for. When a use case is not properly defined by senior stakeholders it's left to different members of the team to decide upon themselves what this end state may be. This can create a fracture in the plan because different teams will interpret the use case resulting in an unorganized team not delivering value for the project.

Once a clear use case is defined stakeholders should decide on an achievable **Return on Investment (ROI)** to further clarify the value GenAI needs to provide for the business. Often times ROIs for research projects like GenAI are loosely defined giving different stakeholders competing expectations about what the value of the project means for their deliverables. Misinterpreting ROIs can erode stakeholder trust so it's imperative at the beginning of a project with unrealised technology to be upfront and aligned with stakeholders across the company.

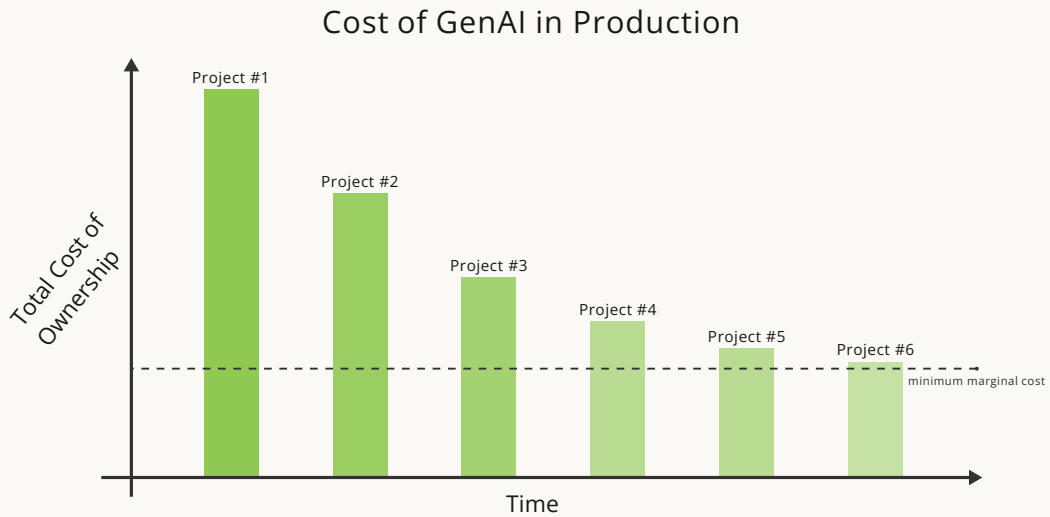

It's important to define a clear ROI metric for GenAI projects as the Total Cost of Ownership (TCO) can vary between projects. Like many projects the TCO of GenAI decreases as your team becomes more experienced and the infrastructure required is rolled out in your organization. We'll explore this more in the later sections and how scaling AI can be taken advantage of.

The importance of defining ROI metrics cannot be overstated, particularly in the context of GenAI implementations. Unlike traditional technology projects, GenAI initiatives often have broader, more complex impacts across the organization.

2. Factors to Consider on the Build versus Buy Decision for GenAI

The decision to build a custom GenAI solution or purchase an existing platform is based around the strengths and experiences of internal teams. When evaluating build versus buy decisions for GenAI, companies often look to their previous technology adoption experiences for guidance. However, while these past decisions provide context, GenAI presents unique challenges that require fresh analysis. Organizations must honestly assess their internal capabilities to develop and maintain enterprise-scale AI security controls before choosing a path forward.

What sets GenAI apart is its need for both comprehensive data access and broad organizational adoption to deliver maximum value. Unlike most enterprise systems, GenAI platforms must securely interface with company-wide knowledge bases while remaining accessible to users across all business functions. This combination of requirements creates novel security considerations that few organizations have previously encountered in their build versus buy decisions.

This systematic framework enables companies to make strategic decisions that optimize both near-term implementation and long-term value creation.

Internal Capabilities Assessment:

- **Technical expertise within the organization:** Assess whether your teams have experience with GenAI development, including specialized skills like data chunking, embedding, and AI workflows. Without this expertise, organizations risk building solutions that fail to meet security and performance requirements.
- **Available resources for development and maintenance:** Consider how your team will scale from proof of value to production, and whether you have the ongoing capacity to maintain and enhance the solution. This includes both immediate development needs and long-term support requirements.

- **Specific requirements that might necessitate custom solutions:** Evaluate whether your industry or use case requires specialized features or guardrails that aren't available in off-the-shelf solutions. This is particularly important for highly regulated industries or unique operational requirements.

Risk Evaluation:

- **Data security implications:** Consider how customer and private data will be handled within the GenAI system and whether this requires updates to existing data policies. This includes both data access controls and protection against AI-specific security threats.

- **Compliance requirements:** Identify the regulatory frameworks your solution must operate within and whether external products meet these requirements. Different industries and regions may have specific compliance needs that impact your build vs. buy decision.

- **Long-term maintenance considerations:** Calculate the resources required to maintain business value, including whether you'll need additional engineers or external consultants for security updates and feature enhancements. This ongoing commitment must be factored into the decision-making process.

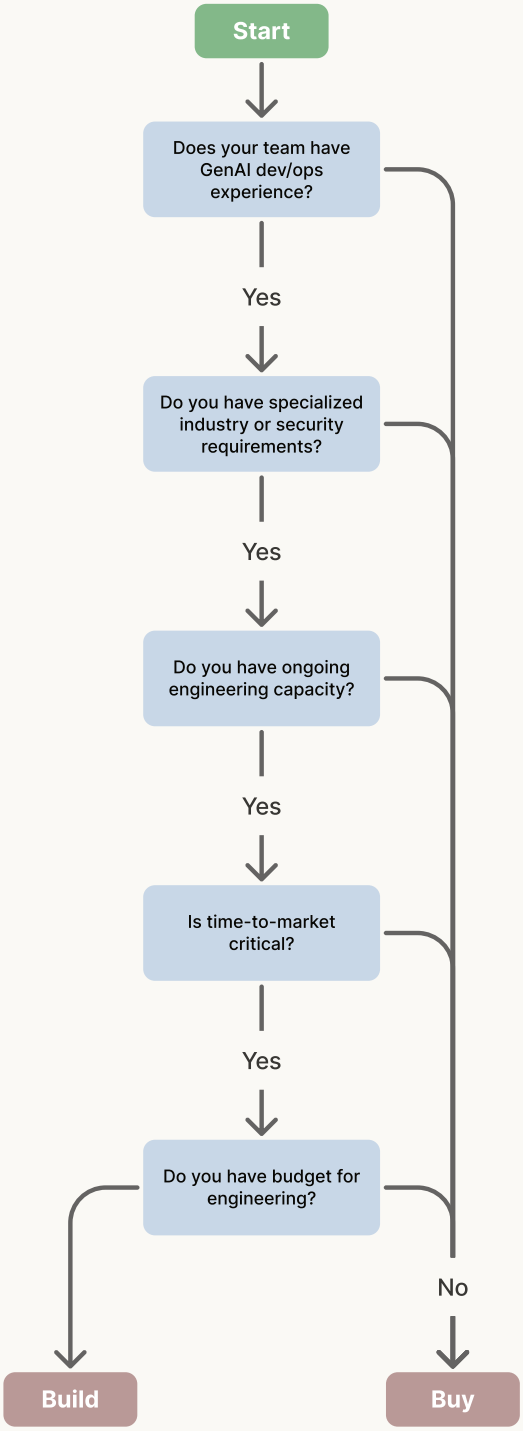

Consider these questions when evaluating to build or buy.

Cost Analysis:

- **Development costs vs. licensing fees:** Compare the expense of internal engineering against the cost of purchasing and maintaining licenses for existing solutions. This calculation should include both immediate and long-term financial implications
- **Training and implementation expenses:** Consider the time and resources needed to train users on the new system and implement it effectively across the organization. Different solutions may require varying levels of training and support.
- **Ongoing operational costs:** Evaluate the budget needed for day-to-day operations, including whether you'll need internal engineers or consulting support for maintenance and updates. This should account for both predictable costs and potential scaling needs.

Organizations must realistically assess their limitations and capabilities to avoid costly mistakes or compliance issues that could prevent successful deployment.

3. Maintaining Momentum by Establishing a GenAI Taskforce

AI represents a fundamental transformation that reaches every corner of the organization. Unlike traditional technologies that operate in specific domains, AI's unique characteristics - its ability to learn from data, interact with customers, share insights across systems, and drive business decisions - make it an enterprise-wide catalyst for change. This pervasive nature means AI should be viewed primarily as a business transformation tool rather than merely a technological implementation.

The taskforce has a clear objective - ***To resolve implementation challenges and direct employee efforts towards the successful rollout of GenAI***.

Establishing a dedicated AI taskforce with executive backing is crucial for navigating the complex landscape of regulatory requirements, legal considerations, and technical challenges. Its capacity to analyze vast datasets, automate customer interactions, and share intelligence across systems demands an integrated approach that spans strategy, operations, and engineering.

Executive Support and Leadership

Senior stakeholder support is essential for giving the project legitimacy and ensuring proper resource allocation. It's critical that the taskforce is sanctioned by the executive team, when employees see clear support from leadership they gain confidence that the initiative has long-term organizational commitment which motivates them to execute the plan with zeal. When senior leadership is aware of the the project the proper resources will be readily available and the teams efforts will be aligned with the companies strategy, maximising the projects ROI. This also gives employees the confidence to escalate any issues so that they can be addressed effectively by the leadership team.

Cross-Functional Representation

GenAI's transformative nature requires whole-of-business alignment. Every department needs to be involved from then start so that in the later stages of development there are no surprises which stall out the rollout.

The taskforce should include representatives from:

- Engineering and Technical Teams
- Legal and Compliance
- Security and Risk Management
- Operations and Business Units
- Change Management and Training
- Executive Leadership

Each department brings crucial perspective to the implementation process. Legal teams understand regulatory requirements, engineering teams grasp technical limitations, and business units provide insight into operational needs. This diverse representation ensures comprehensive consideration of all aspects affecting the project's success.

Taskforce Structure and Authority

To achieve the objective the taskforce should be setup within the organization so that it can be empowered to make strategic decisions. In order to do this it needs to be directly accountable to senior leadership and given clear metric of success signed off by the leadership team. The taskforce needs authorization to implement the necessary changes

across departments to address any deficiencies it identifies. The taskforce should be equipped with the necessary resources to implement the structure outlined so that it can focus on it's key responsibilities.

By establishing clear reporting lines to leadership, the taskforce can efficiently address blockers and maintain project momentum. The predefined use case and ROI goals provide a framework for decision-making, ensuring the team remains focused on delivering measurable value.

Key Responsibilities

Once the taskforce is setup it should focus on its objective. To achieve this there are several responsibilities that should be assigned to the members of the taskforce.

- **Coordinating between departments to ensure alignment** - Providing a single source of truth for what tasks need to be done by each department. This stops a single department from shouldering all of the work or blocking the others.
- **Identifying and resolving cross-functional challenges** - The taskforce should anticipate and resolve any potential challenges that different teams will encounter when working on the rollout. Identifying when departments are not communicating or where tasks have fallen between teams.
- **Monitoring progress against established metrics** - Making sure everyone is held accountable, understanding if the rollout is aligned with the use case and that ROI is maintained.
- **Regular reporting to senior leadership** - The taskforce is the funnel to senior leadership, this streamlines the reporting and escalation paths making sure everyone is aware of the progress while not sending conflicting/outdated information to stakeholders.
- **Managing risk and compliance requirements** - Ensuring the rollout meets internal and industry requirements, anticipating if the project is in danger of not meeting any of these requirements.

- **Ensuring consistent implementation across the organization** - As the project is deployed across the organisation the taskforce needs to monitor that all employees are getting the same experience and no one is left behind.

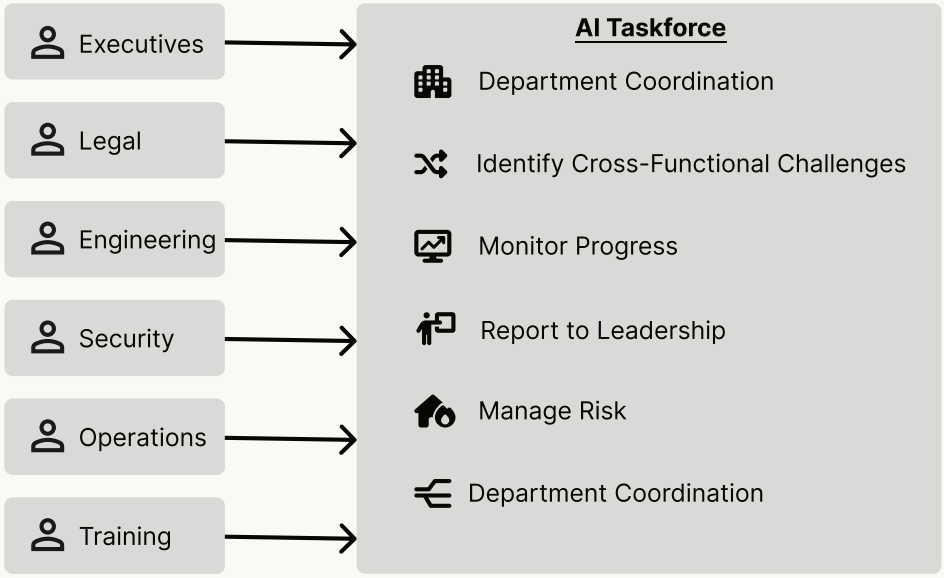

Every team has a role to play in the AI taskforce, assigning the strengths of each team with the right requirements needs to be diligently thought out.

Through these structured responsibilities and clear authority, the taskforce can effectively guide the organization through the complexities of GenAI implementation while maintaining focus on strategic objectives.

4. Knowledge Base Audit: The Foundation of a Secure GenAI Implementation

Why Knowledge Bases Need Special Attention

The integration of GenAI tools into core business processes creates an escalating dependency on real-time access to enterprise knowledge bases. This trend is evidenced by the proliferation of AI capabilities being embedded within established vendor platforms and business applications such as Microsoft Copilot, Glean, and Atlassian Intelligence.

As organizations expand their portfolio of AI-enabled solutions, the structural integrity and effectiveness of their knowledge management systems becomes increasingly critical to operational resilience. The architectural decisions made around knowledge base design

and governance will fundamentally determine the security posture and operational effectiveness of AI-augmented business processes.

Why GenAI Exposes Knowledge Base Vulnerabilities

As organizations explore Generative AI, they often discover its transformative potential comes with unique challenges around data access and security. GenAI's power stems from two key characteristics that set it apart from traditional AI systems.

- ***Human-like understanding of context***. GenAI systems leverage semantic understanding of data using vector databases. This higher level of understanding means that GenAI systems can make connections between concepts in a similar way that people can, only at a much larger scale.
- ***Broad access to your organizations data.*** To provide meaningful value GenAI systems require a lot of contextual information. This is achieved by providing access to most of the knowledge that the organisation has.

While these qualities make GenAI incredibly powerful for improving productivity and innovation, they also introduce important considerations for protecting sensitive information.

The typical enterprise knowledge base represents decades of accumulated institutional knowledge that has grown too large for manual oversight. It will contain vast amounts of unstructured data with many forgotten pieces of content who's authors has long since left the organisation. As the document creators predominately set the permissions for a piece of content the knowledge base quickly becomes misconfigured because of these issues.

Once enough of these misconfigurations emerge in a knowledge base the integrity and security of the answers provided by GenAI becomes compromised. Because of the human-like way GenAI can connect information it trivially exposes mis-permissioned sensitive content or uses outdated information in it’s answers.

Identifying Security Vulnerabilities with Auditing

The most frequent cause of GenAI implementation failure is document misconfiguration. When organizations deploy tools like Microsoft Copilot or Glean, these systems index all

accessible organizational data to provide context for user queries. This process often reveals that permissions are incorrectly configured which leads to sensitive information being accessible to unauthorized people. Outdated and incorrect information from poorly defined security boundaries causes

These issues frequently lead to project failures, as organizations discover they cannot safely deploy GenAI without compromising security or compliance requirements. This issue has been identified by multiple organizations and governments including the Australian Federal Digital Transformation Agency (DTA).

> *Their information management in SharePoint is not great which has resulted in end users finding information that they shouldn't have had access to, though this is a governance and data management issue - not a Copilot issue.*
>
> **Agency representative in Digital Transformation Agency (DTA) interview for Australian Government trial of Microsoft 365 Copilot [2024].**

Solving the compliance non-conformities presented by GenAI is not as simple as turning off the tool. Since the GenAI tools use the users credentials to access data the AI is only accessing content already available. **GenAI demands a fundamental rethinking of security paradigms.**

The solution to this class of misconfiguration vulnerabilities requires a deterministic artifact that can be used to provide confidence to stakeholders that security incidents are resolved, not just identified. Organizations must conduct thorough knowledge base audits before and during GenAI implementation.

This audit process should include:

- ***Data Repository Analysis.*** The audit should review the surface-level permissions of all data sources, mapping semantic content to permissions and the relationships between documents. Auditing will identify any potential security vulnerabilities and assess the compliance requirements needed to repair any gaps. The audit should be able to evaluate the quality and accuracy of content within the knowledge base.

- ***Security Implementation.*** Auditing tools should understand the semantic content of the information being processed, unstructured cannot be properly catalogued without this contextual understanding. As knowledge bases change over time the auditing of the system needs to be continuous, configuring a system to continuously monitor and understand data flowing between knowledge bases provides a solid foundation for cataloguing vulnerabilities.

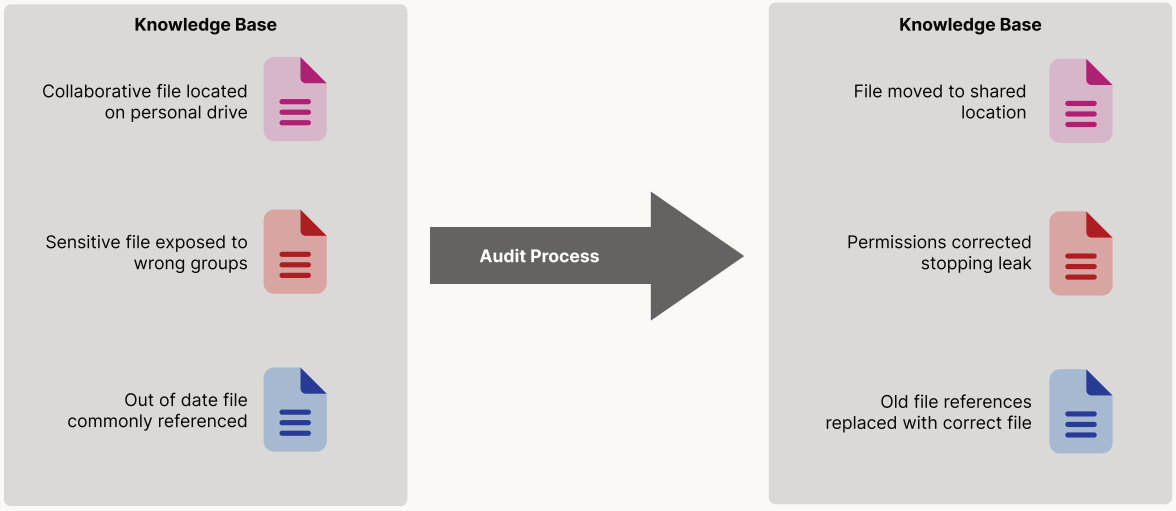

Auditing your Knowledge Base should fix security and operational issues that affect GenAI applications including semantic leaks, incorrect references, and unreachable documents.

- ***Risk Assessment Documentation.*** After the completion of the audit a detailed vulnerability catalogue detailing compliance gaps and remediation steps needs to be created. This will provide the appropriate teams with step-by-step remediation steps with clear plans to implement.

The Role of Semantic Data Security

Traditional Data Security Posture Management (DSPM) tools, while valuable, cannot fully address the challenges brought on by GenAI because they lack understanding of document contexts and relationships. Organisations need to embrace GenAI in this era to solve GenAI problems.

Semantic Data Security (SDS) platforms like Redactive provide essential capabilities for effective knowledge base audits. Using the same technology the underpins GenAI SDS tools can identify misconfigurations and incorrect information found within enterprise

knowledge bases. **By using semantic analysis of documents SDS can detect inappropriate information sharing and provide actionable remediation steps.** This changes a companies risk profile from being reactive to proactive, preparing for the rollout of GenAI.

How Semantic Data Anomalies are Detected

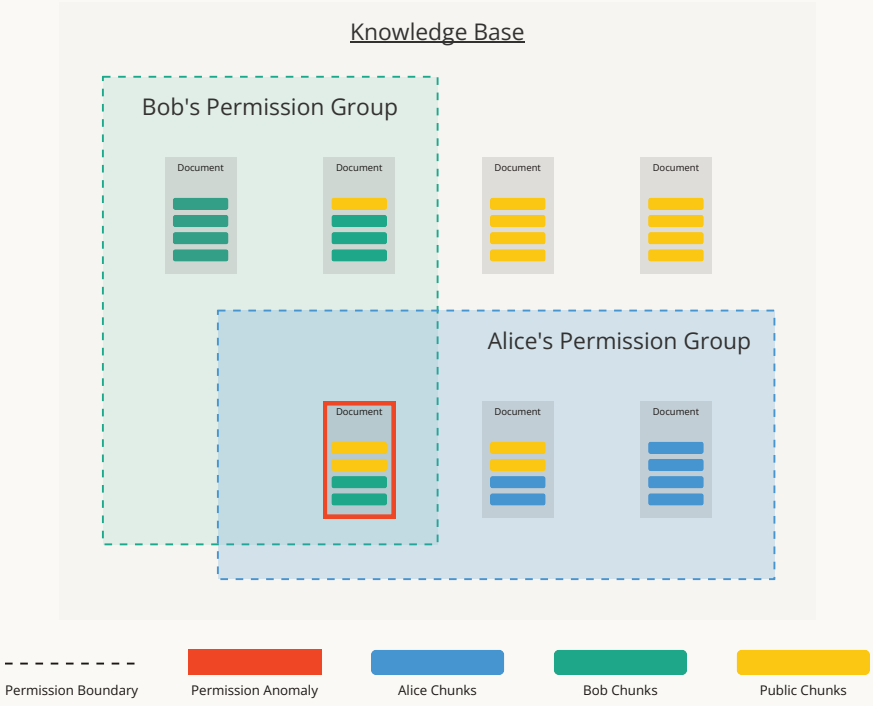

By analyzing document sections (chunks) and comparing the permissions on the document SDS tools find anomolies that other systems can't as they understand your data the same way GenAI does.

Companies can expect to find that **1-3% of all documents inside of a knowledge base are misconfigured**. That number only increases the longer the knowledge base has existed for.

> *Organisations with 500 employees can expect to have 8M documents, of which 3% equals **240,000 data leaks internally**.*
>
> **Analysis from internal Redactive customers**[1]

The audit process must evolve and expand with the organization's requirements. Regular reassessment ensures controls remain effective as the knowledge base grows in size

[1] Data has been calculated based on previous reports generated by Redactive.

and use, compliance requirements need to be continuously met to ensure certifications are met.

Success requires commitment to continuous monitoring and improvement, treating the knowledge base audit as an ongoing process rather than a one-time event.

5. Build an Early Adopter Program and Feedback Loop for your first use-cases

The success of a GenAI rollout hinges critically on establishing an effective early adopter program with rapid feedback cycles. Unlike traditional software deployments, GenAI tools such as Microsoft Copilot and Glean fundamentally transform how employees interact with information, making early user feedback essential for identifying potential issues and optimizing value delivery.

Structured Deployment Approach

A three-phase deployment strategy maximizes control while enabling meaningful feedback, we recommend using this strategy when deploying Microsoft Copilot and Glean making sure to prioritize the proper metrics at each stage for the rollout:

Initial Pilot Group (< 10 users)
- Work directly with the AI taskforce
- Rapid iteration on feedback
- Focus on core use cases and functionality
- Intensive monitoring of security and compliance

Internal Beta (10% of workforce)
- Voluntary participation from motivated users
- Broader testing of security controls and permissions
- Validation of knowledge base access patterns
- Refinement of support processes

Team-by-Team Rollout
- Systematic expansion based on readiness
- Controlled scaling of security monitoring
- Department-specific use case validation
- Progressive enhancement of training materials

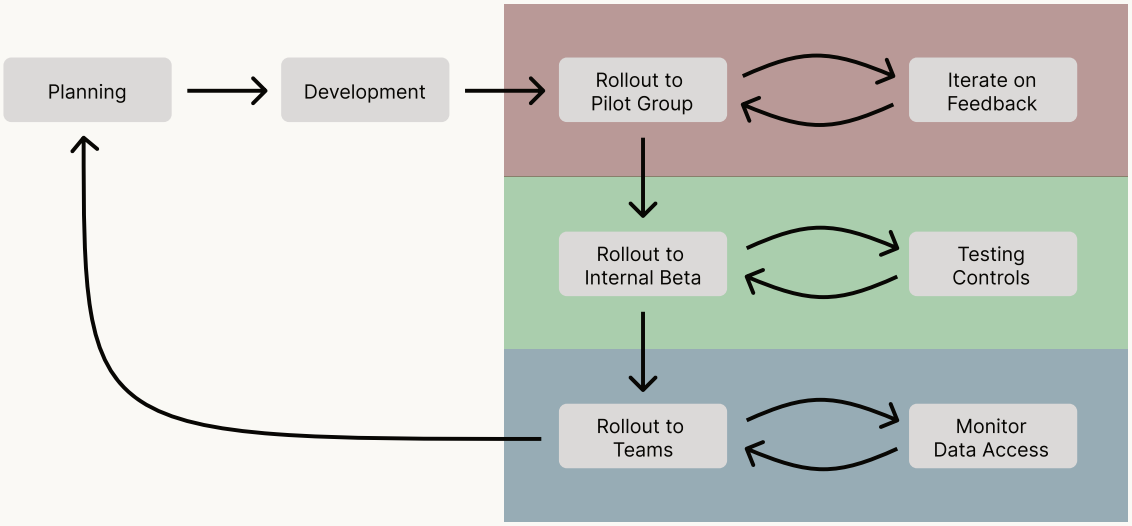

Each stage of the rollout shold focus on a measureable outcome that is directly related to the use case. This diagram provides a compounded view for how a rollout can successfully build user trust and stakeholder confidence with GenAI.

Feedback System Design

The feedback system must support rapid iteration while maintaining security controls. Ensure there are regular check-ins with early adopters, even go so far as to include real time issue reporting systems with usage analytics so engineering teams can react to any issues as soon as possible.

Your feedback system is closely related to your use case and ROI as you can monitor important metrics like time saved per task, accuracy of AI responses, user satisfaction, and knowledge base access patterns.

These metrics can provide valuable insights into risk management and compliance requirements, you can see data access patterns, and validate your permission controls in the real-world.

Success Factors

To maintain momentum and ensure program success:

- Respond to critical feedback within 24-48 hours
- Document and share resolved issues to build confidence
- Celebrate and publicize successful use cases

- Maintain transparent communication about known limitations
- Regular updates to stakeholders on program progress

By following this structured approach, organizations can identify and address potential issues early, refine their implementation strategy, and build a strong foundation for wider deployment while maintaining robust security controls.

6. Scaling GenAI Across your Enterprise to Achieve a Competitive Advantage

Successful rollout of a single GenAI project only marks the beginning of a company wide GenAI transformation. Now that business stakeholders and customers can see the benefits that come from embracing GenAI it's now time to multiply the value created by these tools across your business.

The Challenge of Scaling Many GenAI Systems Securely Across the Enterprise

When enterprises scale GenAI implementations inside their knowledge bases, they face exponentially increasing risk from permission misconfigurations. A simple query to a GenAI system can expose sensitive information through semantic relationships between documents, even when individual document permissions are correctly set.

Traditional security approaches dwindle in effectiveness as GenAI adoption scales simply because they:

- Cannot detect semantic relationships between documents that create unauthorized data exposures.
- Lack ability to understand context and intent of information access.
- Focus on individual document permissions rather than information flows.
- Provide limited visibility into how GenAI systems traverse and expose information.

Redactive's SDS Approach to Scaling AI

Redactive's SDS platform enables secure AI scaling by continuously monitoring and remediating permission misconfigurations through semantic analysis of enterprise

content. The platform analyzes both the content and context of documents to identify potential security vulnerabilities that traditional tools miss.

Redactive provides the foundational components of GenAI tools, once integrated additional GenAI systems become significantly easier to design, maintain, and deliver - decreasing TCO and increasing the value of your knoweldge bases.

Semantic Permission Analysis

Redactive leverages advanced natural language processing to understand document relationships and information flows. This enables detection of permission misconfigurations based on semantic similarity rather than just keyword matching. For example, it can identify when strategically sensitive information about a new product launch exists across multiple documents with inconsistent permission models.

Since Redactive is constantly scanning for permission abnormalities there is a significantly smaller chance that a GenAI rollout will need to be halted or reversed due to permission issues.

Automated Remediation and Continuous Monitoring

As AI usage scales across the enterprise, Redactive provides real-time monitoring of permission changes and content additions. This ensures that the security posture is maintained even as knowledge bases grow and evolve. The platform can detect when sensitive information spreads to new locations or when permission changes create potential vulnerabilities.

When permission anomalies are detected, Redactive automatically generates remediation recommendations aligned with enterprise security policies. These recommendations consider both the semantic content and existing permission structures to suggest the most appropriate access controls.

This means that compliance remidiation controls are automatically adhered to when Redactive is actively monitoring a knowledge base effectively removing hardest part of the AI taskforces requirements - creating compliance controls.

Scalable Architecture

Redactive's platform is designed to scale seamlessly with enterprise AI adoption. The distributed architecture can analyze millions of documents while maintaining consistent performance. API-driven integration enables automated remediation workflows that grow with the organization.

Redactive scales with you so as your workloads increase the experience stays the same. This helps customers meet sclaing requirements aso they're more confident to build a solution rather than purchase an existing one.

Enabling Secure AI Scale

Redactive solves GenAI enablement for regulated industries with SDS tooling that unblocks services such as Amazon Bedrock and Google Vertex AI.

By implementing Redactive's SDS platform, enterprises can confidently scale their AI initiatives while maintaining robust security controls. The platform provides:

- **Proactive Risk Management -** Rather than waiting for security incidents, organizations can proactively identify and remediate permission misconfigurations before they lead to data exposure. This dramatically reduces security risks as AI adoption grows.
- **Reduced Implementation Friction** - Security teams can approve new AI use cases more quickly knowing that Redactive provides continuous monitoring and protection. This accelerates scaling while maintaining security standards.

- **Enhanced Visibility** - Security teams gain unprecedented visibility into how information flows through AI systems and where potential vulnerabilities exist. This enables data-driven decisions about scaling AI adoption.
- **Automated Compliance** - The platform's semantic understanding enables automated enforcement of compliance requirements, reducing the manual effort required as AI scales.

Business Impact

Organizations leveraging Redactive's SDS platform for AI scaling typically see:

- 60% reduction in time required for security reviews of new AI implementations[1]
- 80% decrease in permission-related security incidents[2]
- 40% acceleration in AI adoption timelines[3]
- 95% reduction in manual effort for permission management[4]

By providing a robust foundation for secure AI scaling, Redactive enables organizations to capture the full value potential of AI while maintaining stringent security controls. The platform's semantic understanding and automated remediation capabilities ensure that security scales seamlessly with AI adoption.

For enterprises serious about scaling AI initiatives, implementing semantic data security is not optional - it is a critical enabler that determines the success or failure of AI scaling efforts. Organizations that fail to address the semantic security challenge will find their AI initiatives constrained by security concerns and manual processes that inhibit true enterprise scale.

[1] Based on current customer implementations using the Redactive Data Access platform.

[2] Based on reports generated for customers of Redactive Security.

[3] Based on initial customer timelines before being introduced to Redactive.

[4] Calculated from customer estimations of previous time to repair permissions over a 18 month timeframe.

Book a free consultation.

Find out how to effectively use AI in your organisation.

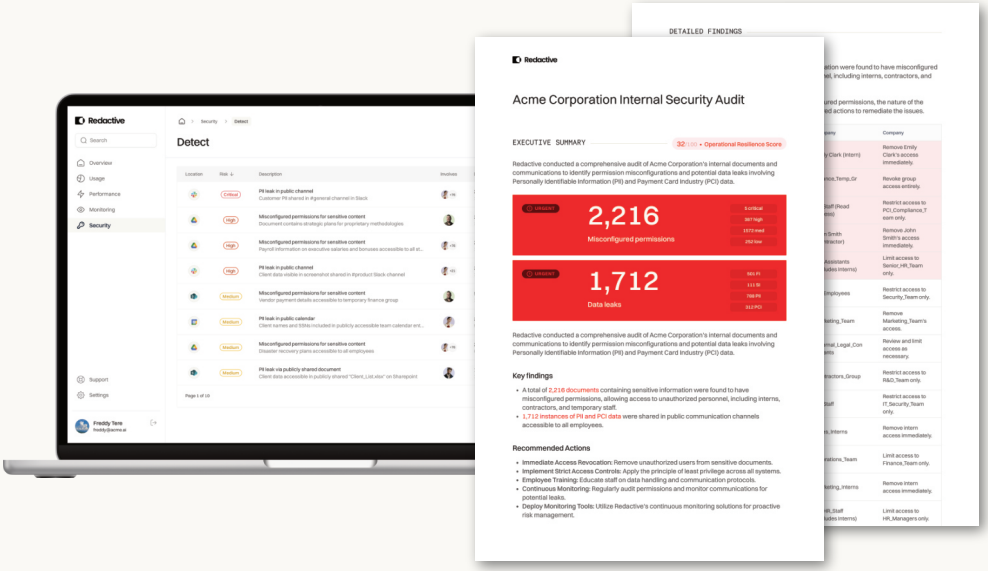

Redactive

Redactive is the trusted AI enablement platform for enterprises. By understanding the semantic content (meaning) of your operational data, we can reduce your risk posture and prepare banks, super funds, and large corporations to confidently adopt AI tools and navigate data stewardship regulations such as GDPR, ISO27001, and CPS230 at scale.

Redactive.ai's Semantic Data Security Agent accelerates regulated organisations to secure & remediate misconfigured/overshared permissions across their knowledge bases, preparing it to be used by AI applications such as AmazonQ, Copilot and Glean or your custom AI tools built specifically for your workflows.

Find out how to tackle data leaks ahead of your AI tools exposing them

Book your consultation here

Andrew Pankevicius, Co-Founder
apankevicius@redactive.ai

Alex Valente, Co-Founder
avalente@redactive.ai