

GENERATIVE AI & RPA – OPPORTUNITIES FOR FINANCIAL REGULATORY AUTHORITIES

USE CASE 1: AUTOMATING QUALITATIVE
COMPLIANCE ASSESSMENTS

INTRODUCTION

Recent advancements in generative artificial intelligence (AI) have marked a critical milestone in the evolution of technology. While powerful on its own, significant additional opportunities and value can be unlocked by combining generative AI with other technologies, including robotics process automation (RPA).

In the introductory article "Generative AI & RPA: Opportunities for Financial Regulatory Authorities" we examined how generative AI and RPA, if integrated, can have a mutually reinforcing effect that supports their application in more advanced use cases, such as by enabling new design options for the digital delivery of processes that require cognitive abilities.

This first deep dive in our series of use cases extends this discussion to more concrete application opportunities of these technologies by examining how they can streamline the compliance assessment by financial supervisors.

Specifically, drawing on successful pilot work, the deep dive explores the architecture of a generative AI and RPA-powered tool that can fully automatically execute qualitative assessments of licensed entities' internal policies vis-à-vis regulatory requirements, reducing the time to complete an assessment to a fraction relative to a manual review while maintaining a quality and precision in outputs that matches those of supervision staff.

Besides an overview of the tool's core technical components, the article offers an overview of the necessary prerequisites that must be in place for its successful application, summarizes key benefits as well as discusses the implications for the role of supervisors.

01

POLICY COMPLIANCE REVIEWS – AN INTRICATE & LABOUR- INTENSIVE TASK

A central aspect of supervision activities during license application processes and subsequently as part of off-site supervision is the detailed review of licensed entity internal policies and procedures vis-à-vis regulatory requirements with a view to establishing the level of compliance and pinpoint potential gaps or breaches.

At present, the review of such policies often remains a highly manual and time-absorbing process. Carried out fully manually, a thorough review may take hours to complete. The core challenge lies in the qualitative and unstructured nature of these documents, which frequently lack a standardized format and structure, requiring a differentiated review approach for each document and thus exacerbating a systematic assessment at scale. Further complicating this is the dynamic nature of regulatory requirements, necessitating an up-to-date understanding of the latest regulation and re-evaluation of compliance.

This complexity demands a high level of expertise and experience, making the process both labor-intensive and time-consuming, which is at odds with the time constraints that supervisors frequently operate under. Additionally, the need for subjective interpretation introduces an element of variability in assessment outcomes, which can impair the consistency of assessment outputs.

These challenges create a strong case for the application of technology into the compliance assessment process.

02

A PROTOTYPE FOR AN AUTOMATED COMPLIANCE ASSESSMENT TOOL

By blending RPA and generative AI, it is possible to significantly accelerate and standardize the review process. In this section we outline the principal architecture of a tool that integrates both technologies to execute qualitative assessments of licensed entities' internal policies vis-à-vis regulatory expectations fully automatically.

At its core, the tool leverages state-of-the-art generative AI models to perform an assessment of a policy's content against a set of defined and thematically grouped regulatory standards and to draw evidence-based conclusions on the entity's level of compliance and potential compliance gaps as well as formulate recommendations to the supervisor for follow-up actions. Rigorous testing as part of limited-scale pilot work has demonstrated the tool's capabilities to deliver an assessment quality in terms of precision, level of detail and consistency on par with professional supervisors, while significantly surpassing human capabilities in terms of the speed of delivery and the volume it can handle.

The tool operates as an end-to-end automated processes comprising of five consecutive steps, which are summarized in the following and illustrated in Figure 1.

Process initiation

The tool's execution is externally triggered by the supervisor at a point when a new licensed entity policy is available for review. The supervisor securely submits the policy or indicates its location through a dedicated internal user interface and identifies the domain of standards the policy is to be evaluated against.

In an advanced configuration of the tool, the supervisor has the option to select multiple policies from the same entity for analysis against the same standards with a view to obtaining a single integrated assessment.

Data extraction

Following the intake of the file, the tool extracts the complete file contents as is and temporarily stores it for further processing. In its baseline configuration, the tool allows exclusively for the processing of machine-readable files, such as readable PDFs. In an advanced setting, the tool can be configured to handle more complex document designs, such as by leveraging OCR and other techniques to extract scanned content and/or non-readable images and graphs. .

In parallel to the extraction, the tool retrieves the defined regulatory standards for the domain in question, which are housed permanently in a separate database. To support a focused analysis, regulatory standards are grouped thematically into a subset of categories.

As an optional step pre-analysis, the policy's extracted contents can be thematically segmented in alignment with the defined thematic regulatory categories prior to being ingested for analysis.

Analysis

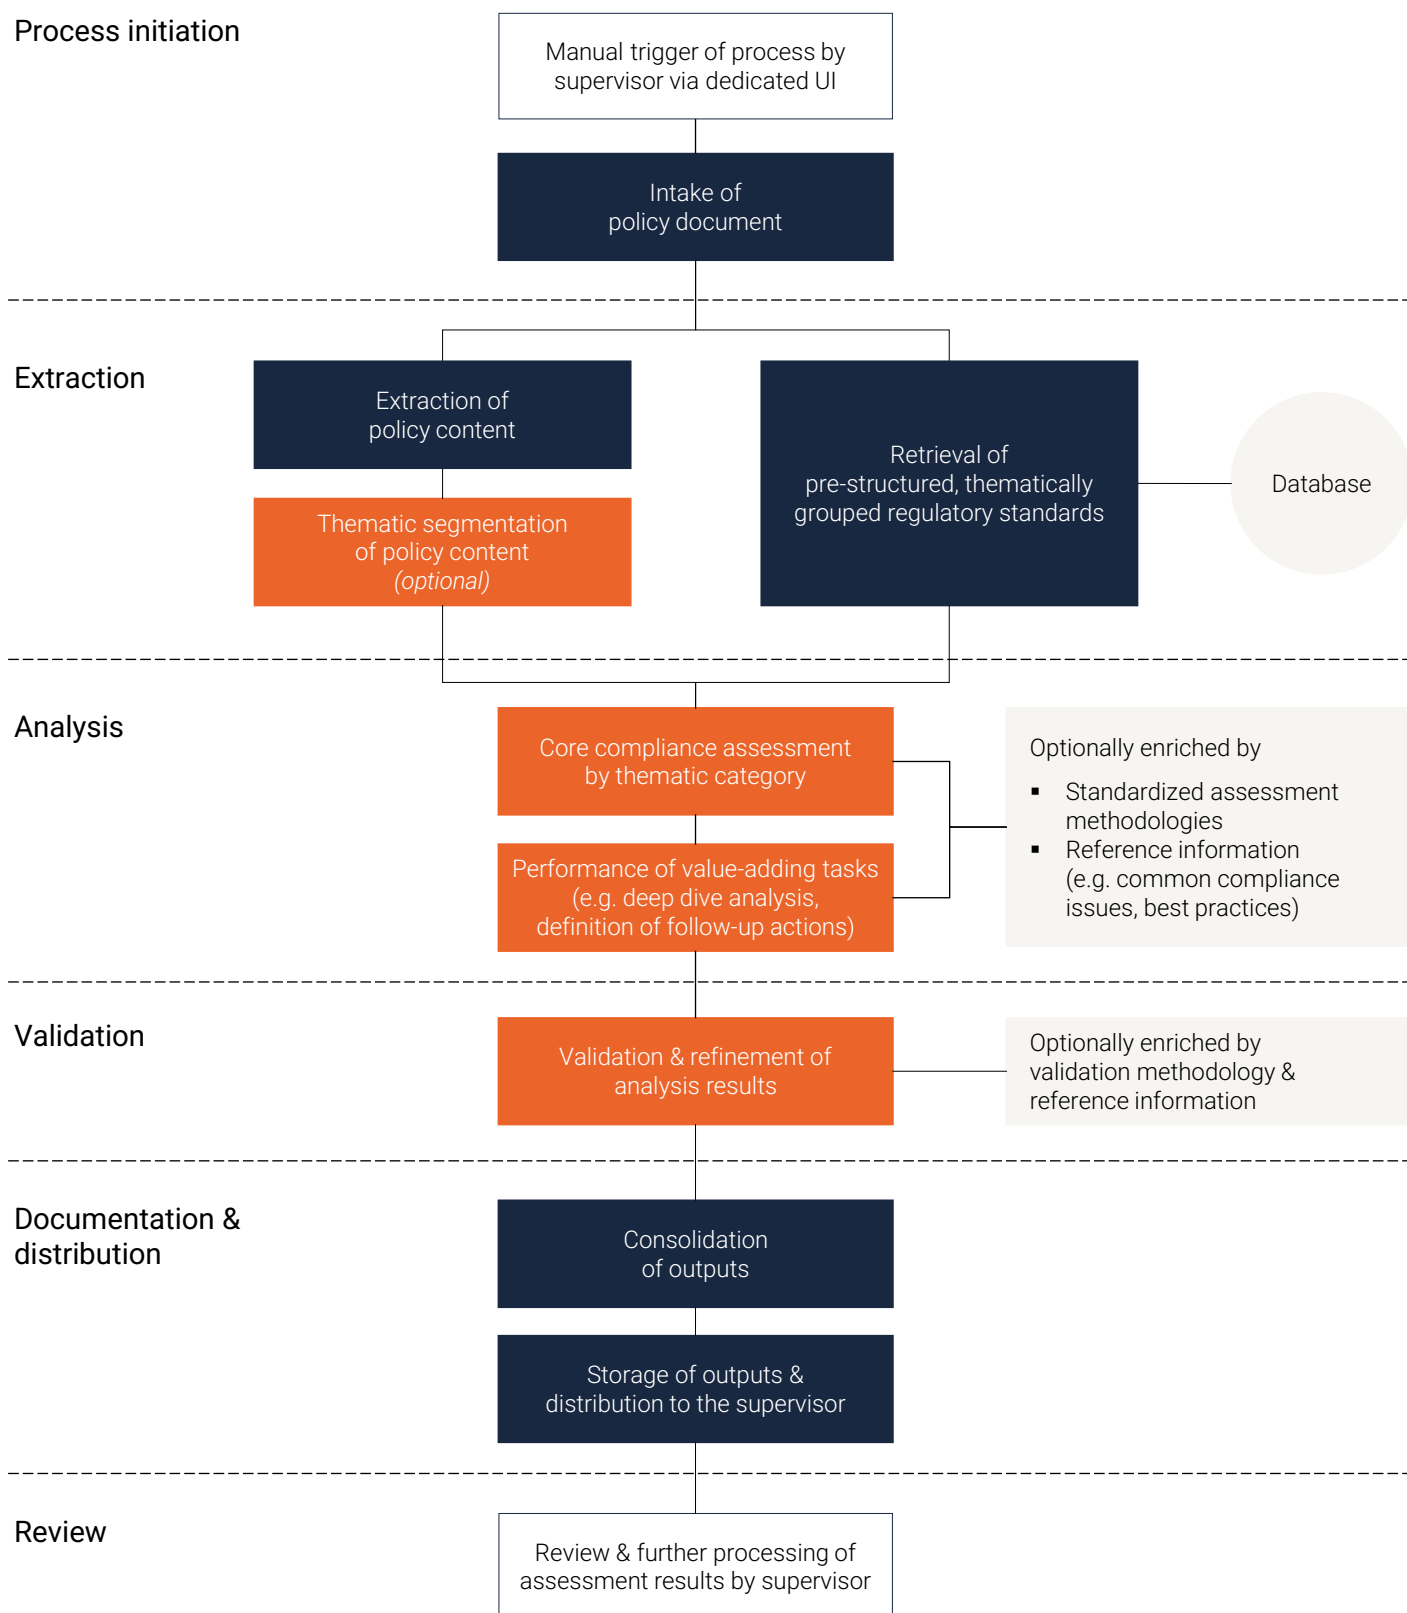
The analysis stage comprises of multiple generative AI-powered steps.

In a foundational step, the tool dynamically combines a standardized prompt with specific analysis instructions with the contents from the policy and the thematically grouped regulatory requirements to perform the core compliance assessment for a given thematic category of regulatory requirements.

As part of the assessment, the tool compares the policy's contents with regulatory requirements clause by clause, matching it with the relevant contents from the policy and subsequently identifying areas where the policy aligns with requirements and where it falls short and/or exhibits material discrepancies. In doing so, it follows an evidence-based approach and documents the justification and the supporting evidence from the policy content for each finding, thereby making conclusions transparent and verifiable during the subsequent review by the supervisor. Likewise, outputs from the analysis are structured in a standardized format to facilitate the review process.

On the basis of the results of the core compliance analysis, the tool can optionally be instructed to perform additional value-adding tasks. These can range from performing deep dive analyses on specific aspects of interest, prioritizing findings by criticality, or computing an aggregate compliance score. For greater action-orientation, the tool can also serve to define specific points for follow-up by the supervisor with the entity.

Figure 1 Tool architecture (simplified illustration)



For an augmented analysis or complex assessment scenarios, the tool's logical reasoning capabilities and consistency during the analysis can be further enhanced by supplying it with standardized assessment methodologies, best practice examples or lists of common compliance shortcomings as an additional reference point. To achieve this, retrieval augmented generation (RAG) or other advanced prompting techniques can be applied to inject the relevant knowledge.

Validation

To ensure optimal assessment quality, it is recommended to complement the analysis engine by a validation layer.

Akin to a human-led review process, the validation step involves assessing the produced analysis output against certain quality criteria such as coherence, consistency and comprehensiveness of conclusions, flagging conclusions whose supporting evidence is vague, lacks substance or does not align with the policy content, and identifying other potential weaknesses.

This validation layer can have different levels of scope and sophistication and can too be realized through different mechanisms, ranging from finetuned models or agents specifically trained to assume a challenger role to advanced prompting techniques. Similar to the analysis step, additional supervisory intelligence can be injected into the process to enrich the validation.

Based on the validation results, the analysis output can either be further refined and/or validation findings be documented for the supervisor's attention.

Documentation & notification

In a final step, the results from the individual analysis steps are consolidated and documented in a pre-defined, structured format complete with the supporting details and returned directly to the supervisor for review and/or stored in a secure local or cloud environment.

Designed for flexibility, operational efficiency & scalability

The tool's architecture under the proposed design has been optimized for flexibility. Its modular design allows to flexibly tailor the individual steps during both the analysis and validation stage as well as the structure and depth of its outputs in accordance with an authority's specific needs and preferences.

Furthermore, all steps are designed to achieve high execution speed without compromising on the quality of outputs. The average processing time for a single policy page is approximately 30 seconds, with the potential for scaling efficiencies in longer policy documents. For instance, an end-to-end analysis of a 30-page policy document at high precision, complete with the detailed documentation of results and supervisory recommendations, can be achieved in as little as 12 minutes.

Moreover, across all steps, the tool has built-in error management controls, designed to minimize manual intervention required if a process step fails.

Depending on the desired scope of application, the tool can be enhanced to accommodate multi-document processing. This can entail both integrating the assessment of multiple documents from a single entity against the same standards into a single, consolidated compliance assessment as well as the ability to execute multiple analyses in parallel.

“

An end-to-end analysis of a 30-page policy document at high precision, complete with the detailed documentation of results and supervisory recommendations, can be achieved in as little as **12 minutes.**

”

03

BENEFITS BEYOND RESOURCE SAVINGS

Implemented with the right prerequisites in place, the tool can offer a broad range of benefits to supervisory departments. Besides multiple operational benefits, the tool leads to improved assessment outputs in multiple respects.

By uniformly applying predefined criteria and methodologies, the tool eliminates the subjective interpretations and biases that often arise in human analysis, ensuring impartial and equitable treatment of all entities, free from any influences.

Moreover, it significantly reduces the risk of inconsistencies that are inevitably prevalent in human-led assessments due to variations in individual supervisors' profile, understanding and interpretation of policies and requirements, contributing to more reliable and uniform compliance evaluations. Similarly, it reduces the risk of errors and oversight, especially in the case of complex and detailed policies that require a systematic analysis approach and attention to detail that is difficult to achieve under time constraints.

Figure 2 - Summary of core benefits

CONSISTENCY & OBJECTIVITY

- Provides for the uniform application of criteria and methodologies, minimizing human biases and external influences
- Results in more objective assessments and reduced inconsistencies in evaluating compliance

ERROR REDUCTION

- Minimizes the likelihood of oversight and errors, particularly in complex and lengthy document analysis
- Enhances accuracy in assessments, mitigating risks associated with manual human analysis

DOCUMENTATION QUALITY & AUDIT TRAIL

- Generates comprehensive, evidence-based assessments with specific references to regulatory requirements, supporting an improved audit trail
- Offers a level of granularity and transparency that surpasses the capabilities of human analysts within the same timeframe

OPERATIONAL EFFICIENCY

- Reduces the time for assessment completion compared to manual processing, reducing burden on and freeing up resources
- Provides continuous operational capability, adding flexibility and efficiency in the execution of assessments

In addition to these core benefits, the design of the tool supports seamless integration and compatibility with other existing systems and, importantly, flexible integration into existing workflows.

Furthermore, its capability of handling large volumes of documents and assessments simultaneously, enable it to scale effectively and to adapt flexibly to varying scales of workload at any given point in time.

04

ESSENTIAL PREREQUISITES

For the effective implementation and operationalization of the tool in its proposed design, several prerequisites are essential. These not only ensure that the tool delivers optimal outputs but also meets stringent information security requirements.

Information infrastructure

The foundation of the tool's design and effectiveness as presented herein lies in having documented a structured set of regulatory expectations, articulated in clear and specific language at sufficient granularity. The tool therefore caters particularly well to predominantly rules-based regulatory frameworks, which inherently meet these structural requirements. Breaking down regulatory expectations into thematic groups is critical to ensuring that the generative AI model can execute the compliance assessment with sufficient focus and accuracy, avoiding key information to be left out from the analysis.

Alongside this, robust internal mechanism must be established for managing updates to regulatory expectations and ensure they remain current and comprehensive and the tool's assessment approach remains reflective of evolving practices and insights from ongoing supervisory work. These requirements extend to the management of additional contextual information, which may be injected into the tool's analysis and validation layer.

Information security

Given the sensitive nature of the information processed, the tool must meet high information security standards. The implementation of advanced encryption protocols to protect data both in transit and at rest is a critical starting point. This includes ensuring that the tool is operated in a secure local or cloud environment where data is sufficiently isolated and that all communications between the tool and generative AI models are conducted over HTTPs with SSL/TLS encryption with clearly established boundaries that prevent information processed through the tool to be transmitted directly to the generative AI model provider.

Moreover, stringent access controls must be established for the operation and maintenance of the tool, permitting only authorized personnel to view and process sensitive information. This must be supported by audit trails that transparently log every access and action in the system.

User training


Comprehensive training and consistent support for users are crucial to ensure the tool's effective adoption. Supervisors must be thoroughly trained in operating the tool, interpreting its outputs, and integrating it into their existing workflows, while also understanding its limitations.

The design of the proposed tool seeks to facilitate that by reducing the technical burden on supervisors and facilitating a seamless interaction experience. Additionally, both the structure of the tool's assessment outputs, and the additional controls implemented at the validation support easy identification of potential errors and weaknesses in logical reasoning that necessitate the supervisor's validation.

Performance monitoring & refinement

Once operational, the tool requires continuous performance monitoring and a systematic schedule for updates and maintenance. This demands effective collaboration between regulatory experts who provide the necessary subject matter expertise for the tool's generative AI modules, and technical professionals responsible for maintaining the technical components of the generative AI and RPA systems.

Particularly critical is a regular validation of the tool's assessment quality through sample testing and direct supervisor feedback and the incorporation of the insights in the form of refinements to the techniques applied at the analysis and validation stage and amendments to the underlying information infrastructure.





CONCLUSION

The integration of generative AI with RPA holds significant potential in expanding financial regulatory authority's digitalization opportunities by extending into areas that require higher cognitive analysis, such as qualitative compliance reviews at the licensing and supervision stage. The tool presented in this inaugural use case deep dive demonstrates the opportunities for streamlining and materially enhancing the operational efficiency of core supervision processes without comprising quality. However, realizing its full potential requires careful attention to the underlying information infrastructure, data security, user training and continuous performance monitoring and refinement.

For supervisors specifically, the advancement necessitates a redefinition of the role of supervisors. It marks a shift towards a more strategic oversight role and creates additional bandwidth for more value-adding tasks including productive interactions with entities over potential compliance deficiencies and their timely resolution.

Achieving this transition will require supervisors to develop skills in interpreting and leveraging the outputs of such technology tools with an intimate understanding of their limitations as well as to rethink their approach to supervisory decision-making.

ABOUT REGXELERATOR

Regxelerator is a technology-powered intelligence platform and innovation solution partner to the financial regulatory community.

As an intelligence platform, Regxelerator focuses on the development of novel, AI-enabled solutions that enhance access to and smart extraction of insights from financial regulatory developments. Its flagship platform Regxplora integrates regulatory updates from over 700 sources globally including financial regulators, central banks, standard setters and other key actors in the financial regulatory ecosystem.

As a solution provider, it assists financial regulatory authorities in accelerating their organizational performance and digital transformation agenda through novel technology solutions for the delivery of regulatory and supervisory processes, combining generative AI and RPA-powered solutions and tools with ancillary consulting services.

For more information visit regxelerator.com or send an inquiry to contact@regxelerator.com.