CORE®
AN EPAM COMPANY

Next-Generation GRC: AI as a Strategic Enabler

Opportunities, Challenges, and Use Cases for AI Deployment in Modern GRC Functions

Dr. Philipp Kleine Jäger
Dr. Carola Bader
Muskaan Multani

September 2025
Whitepaper

Public

Key facts

- **GRC (Governance, Risk, Compliance) is under pressure – and on the verge of transformation:** Compliance departments face increasing regulatory complexity, data overload, and inefficient processes. At the same time, GRC is evolving from a purely control-focused function into a strategic success factor.
- **AI transforms GRC from merely a cost factor into a competitive advantage:** AI enables not only automation but also proactive risk management, faster adaptation to regulations, and strategic steering – particularly relevant given limited resources.
- **AI can be applied with varying degrees of automation:**
    - Rule-based systems automate simple tasks.
    - AI assistants/copilots support analysis and decision-making processes.
    - AI agents operate largely autonomously – although human oversight is essential in complex regulatory matters.

  Rule-based systems are a baseline requirement nowadays, but planning for additional applications (taking into account organizational context, risk appetite, etc.) is key for achieving efficiencies.
- **Concrete use cases can be identified for all major GRC activities** – two examples demonstrate that the potential can be realized not only in theory but also in practice.
- **The AI market in the GRC space is diverse** – spanning from established platform providers to specialized startups – and offers solutions with varying levels of maturity, focus, and use-case coverage. A structured market analysis aids in identifying technology- and company-specific entry points for AI deployment.
- **Transformation potential and outlook:** When implemented correctly, AI can fundamentally transform the GRC function – shifting it from reactive compliance to a strategic enabler. This requires early investments in skills, technology, and governance to ensure that compliance is resilient, efficient, and future-proof in the long term.

1. GRC in transition – three key challenges

The landscape of regulatory requirements is becoming increasingly complex and dynamic. This presents three key challenges for compliance departments.

1.1 Rising regulatory demands

New regulations, such as the EU AI Act, DORA, or FIDA, not only require deep expertise from companies but also continuous adjustments to their processes. This is accompanied by increased resource allocation, such as establishing a dedicated ICT control function in accordance with DORA. According to a recent study, 23% of IT and security experts see the greatest challenge of their department in keeping up with these changes.[1] One reason for this: nearly 70% of companies were required to confirm compliance with six or more regulatory frameworks in 2023 – especially in the areas of information security and data protection.[2]

1.2 Data overload and process inefficiency

At the same time, the amount of data that needs to be processed is growing exponentially. Many GRC processes, from monitoring to reporting to risk analysis, remain manual and are therefore time-consuming and prone to errors. It is also important to note that the tool landscape within organizations is becoming increasingly complex, meaning that compliance with regulations often involves numerous systems. These developments lead to rising costs and a greater demand for specialized GRC professionals. Companies are faced with the challenge of how to balance these increasing demands with limited resources.

1.3 Strategic shift in compliance

The role of compliance teams has changed significantly in recent years: 70% of risk and compliance professionals report that companies are increasingly moving away from a purely "check-the-box" mentality toward a strategic understanding. Compliance is increasingly being seen as a business-critical success factor – not just to minimize risks, but also as a potential competitive advantage. Regulatory bodies are also explicitly mandating this shift. For example, DORA Article 6(8) requires that the ICT risk management framework must include a digital operational resilience strategy, define its implementation, and describe methods for managing ICT risks as well as achieving specific ICT objectives, including supporting the business strategy. This transformation opens up new opportunities but requires more efficient processes and the focused use of technology.[3]

However, a survey of CFOs and compliance leaders reveals that only 16% of companies currently pursue a strategic approach to compliance management. The majority remain in a reactive or administrative mode, focusing on meeting minimum regulatory requirements rather than aligning

[1] https://coalfire.com/insights/resources/reports/securealities-report-2023-compliance
[2] https://coalfire.com/insights/resources/reports/securealities-report-2023-compliance
[3] https://www.thomsonreuters.com/en-us/posts/investigation-fraud-and-risk/risk-compliance-survey-report-2023/

compliance with business objectives. This highlights existing process inefficiencies and the growing pressure to make compliance smarter and more cost-efficient.[4]

Against this backdrop, a central question arises for many companies: How can GRC be transformed using artificial intelligence (AI) from a resource-intensive obligation into a true strategic lever – without compromising on security or flexibility?

2. AI as an enabler for future-proof GRC management

In the face of rising regulatory requirements, growing data volumes, and increasingly complex system landscapes, companies are seeking solutions that not only ensure efficiency and security but also provide flexibility. AI offers transformative approaches that go far beyond mere automation, establishing GRC as an integral part of corporate strategy. The good news is that a wide range of off-the-shelf solutions already exists, which can be tailored to the specific requirements of companies, making the adoption of AI-powered GRC strategies significantly easier. A market overview of GRC providers will be discussed in Chapter 5.

2.1 Adaptability to regulatory changes

Studies show that compliance teams currently spend between one and seven hours per week tracking and analyzing regulatory changes. A reduction in this workload across many organizations indicates the growing adoption of RegTech solutions.[5] In particular, AI-based systems offer great potential: they quickly detect regulatory changes (such as updates to sanctions lists), automatically analyze relevant data sources (e.g., legal databases with implications for operational work), and directly integrate new requirements into existing processes, with compliance teams retaining control over the final adjustments. This not only significantly reduces manual effort but also minimizes the risk of regulatory violations – freeing up resources for more strategic GRC work.

2.2 Automating compliance processes as an efficiency driver

A core benefit of AI lies in the intelligent automation of time-consuming routine tasks. Activities such as document reviews, transaction monitoring, or report generation can be significantly accelerated by AI-powered systems. This intelligent automation not only reduces manual effort but also minimizes errors caused by human intervention. According to a study, nearly 38% of companies were able to automate 51–75% of their compliance tasks by 2024, leading to a reduction in time spent by over 50%.[6]

[4] https://www.globenewswire.com/news-release/2024/09/18/2948251/0/en/Survey-Only-16-of-Organizations-are-Approaching-the-Next-Frontier-of-Compliance.html
[5] https://legal.thomsonreuters.com/content/dam/ewp-m/documents/legal/en/pdf/reports/cost-of-compliance-report-final-web.pdf, pp. 7-8
[6] https://www.middesk.com/blog/how-ai-and-automation-are-reshaping-the-compliance-landscape

2.3 Proactive risk management through data-driven AI approaches

The increasing prevalence of Generative AI (GenAI) is giving AI use in the GRC space a new boost, accelerating the adoption of innovative solutions. GenAI shifts the focus of compliance management from purely reactive risk management to proactive and predictive steering. While established techniques like Predictive Analytics and Machine Learning already analyze large data sets in real-time to detect patterns, anomalies, and risks early, GenAI enables deeper and more automated processing of complex data structures, combined with significantly enhanced contextual understanding and adaptability. Natural Language Processing (NLP) is complemented by GenAI's ability to more intuitively and precisely analyze unstructured content such as emails, chat logs, or documents – for example, in fraud prevention or risk detection. At the same time, GenAI accelerates the evolution of Predictive Compliance, utilizing historical data with even greater precision to forecast future risks and regulatory violations and optimize preventive measures. In areas like anti-money laundering, the potential is particularly evident: AI not only significantly reduces the number of false positives but also detects complex fraud patterns, such as nested money flows or networks. Another example of AI's effectiveness is in curbing data breaches: organizations that comprehensively use AI and automation detect and halt such incidents almost 100 days faster on average than those that do not leverage these technologies.[7] This combination of proactivity and precision makes compliance not only more effective but also a strategic competitive advantage.

2.4 Strategic benefits through AI integration

Organizations that successfully implement AI ultimately transform GRC from an obligation into a strategic advantage. Cloud-based solutions provide companies with the ability to efficiently adapt compliance processes across global locations while accurately accounting for local regulatory requirements, which is becoming increasingly important for multinational companies facing diverse and complex regulatory frameworks. Furthermore, proactive compliance strengthens trust among customers and partners, offering a significant competitive edge in heavily regulated industries like financial services. The positive effects of AI adoption are clear: in a survey, nine out of ten users reported that AI significantly improves risk and compliance management while making processes more efficient – particularly through the automation of repetitive tasks and controls, early risk detection (for example, in fraud prevention), and better organization and analysis of large volumes of data. This enables deeper and more comprehensive insights, supporting informed decisions and reducing errors.[8] With its ability to create scalable and agile operations, AI significantly elevates the strategic value of GRC for modern organizations.

[7] https://www.ibm.com/downloads/documents/us-en/107a02e94948f4ec, p. 18
[8] https://www.compliance-manager.net/artikel/ki-koennte-risiko-und-compliance-management-beschleunigen/

3. High-impact use cases: How AI transforms GRC activities

The use of AI in the GRC space can be differentiated by various levels of autonomy. These levels reflect the extent to which a system independently takes on tasks, makes decisions, or supports human users. The maturity of AI application can be categorized into four key stages:

1. Rule-based systems

These systems are based on predefined "if-then" logic and automate clearly structured, repetitive processes. They form the foundation of digitalization in GRC but are neither self-learning nor adaptive. Typical use cases include automated audit mechanisms, structured workflows, or rule-based reporting functions. In many companies, they are considered a baseline requirement for efficiently managing the growing documentation and control workload.

2. Copilots (Out-of-the-box AI tools)

These AI-powered solutions, often referred to as "Copilots" or "AI Assistants," are already widely used in other business areas (e.g., marketing or software development) but have so far been utilized sparingly in the GRC space. They are usually generically trained and not specifically tailored to regulatory contexts. Nevertheless, they hold great potential: they assist in structuring and formulating content, information research, or data visualization. However, it is critical that human users review and assume responsibility for the results – full control remains with the user.

3. Humans with agents

At this level, GRC professionals actively collaborate with more specialized AI agents that independently perform tasks, prepare decisions, or initiate processes. The human remains in a leading role but is greatly supported by these agents – for instance, in continuous risk monitoring, internal control systems, or dynamic policy adjustments. The systems learn over time, adapt to changes, and are more seamlessly integrated into existing GRC structures. This reduces manual effort while increasing the speed and quality of execution.

4. Agents with humans

The highest level involves highly autonomous AI agents that independently make decisions within defined guardrails, control other systems, and, when necessary, interact with users or stakeholders. Here, the human takes on an overseeing or supervisory role, intervening only selectively. This form of AI implementation enables extensive automation of complex GRC processes but also requires a high degree of governance, transparency, and risk awareness. Issues like accountability, traceability, and ethical considerations come to the forefront.

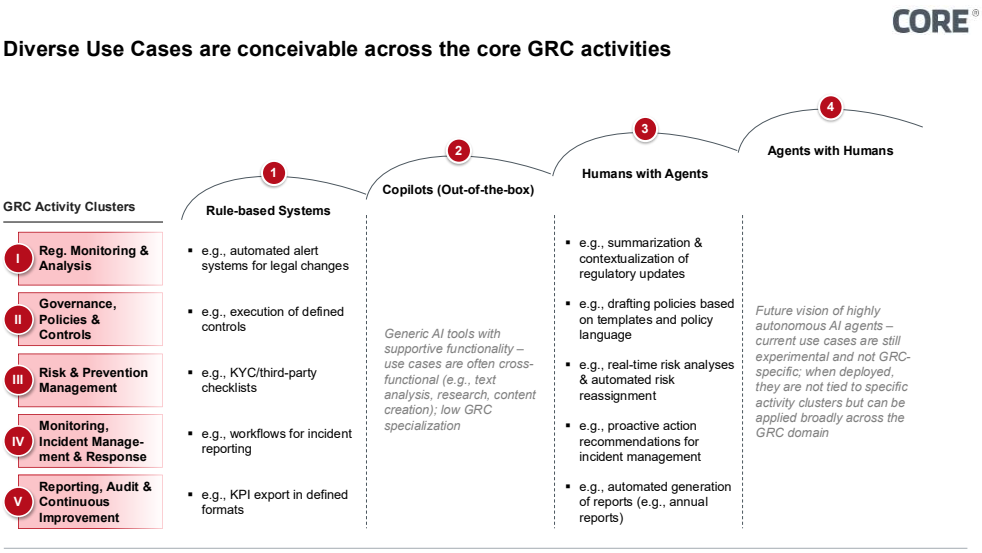

Figure 1: Application areas of AI across GRC activities

This differentiation is essential for accurately assessing the potential applications of AI solutions within the GRC function. Whether rule-based models or adaptive AI agents – the key lies in holistic integration that combines efficiency, scalability, and predictive decision-making, thus establishing GRC as a central component of modern corporate management.

The degree of implementation and the areas of application are strategic considerations that must, in particular, consider the organization's framework conditions (maturity, resources, etc.) and its appetite for AI-related risks (as part of an AI strategy or similar framework).

Rule-based tools – such as automated audit mechanisms or structured workflows – are now considered the minimum requirement for achieving a basic level of efficiency in day-to-day compliance tasks. Without this foundation, increasing documentation and audit workload can hardly be managed. "Humans with agents" models already represent the next leap in efficiency. A growing array of specialized solutions is emerging in this space, which can be integrated into existing GRC processes and provide significant relief to compliance teams. The highest level, "Agents with humans," remains, for now, more of a future vision. Currently, there are no practical solutions available that enable such autonomous management of complex GRC processes. Here, the focus is not only on technological feasibility but also on a thorough examination of risk-based challenges. Who takes responsibility in the event of making a wrong decision? How can transparency be ensured? And how does all this fit into existing governance structures?

4. Two practical examples: AI in reducing regulatory implementation efforts

4.1 Smart monitoring: AI deployment for efficient implementation of regulatory requirements

Currently, monitoring regulatory changes often remains a manual task or relies on simple plugins and APIs. Despite the tools available, the effort is still significant: content must be individually reviewed, categorized, and structured. As a result, the process is time-consuming and prone to

errors. This is particularly challenging in dynamic industries such as financial services, pharmaceuticals, or international law, where laws, rulings, consultations, and standards must be monitored on a daily basis.

a. Functionality of AI-powered news analysis

Artificial intelligence can significantly alleviate these challenges by aggregating, analyzing, and structuring regulatory information automatically – making relevant developments more accessible, faster, and more targeted. Typical functional components within the workflow of such a solution include:

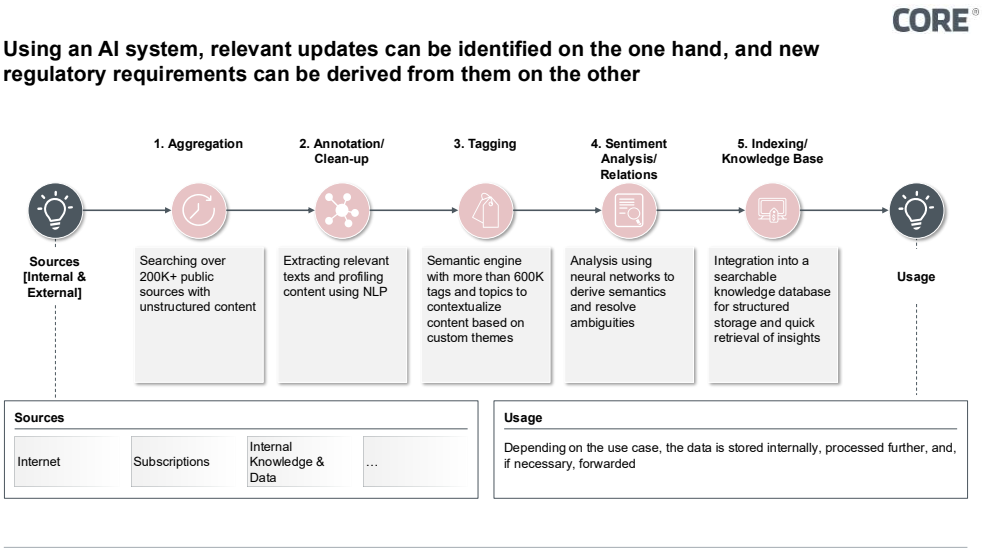

Figure 2: Functionality of AI-powered news analysis

- **Aggregation of news and research sources:**

  Modern systems collect content from a variety of external and internal sources, such as specialist portals, official publications, websites, social media, as well as internal company repositories like SharePoint, databases, or cloud services. Existing subscriptions or internal document archives can often be integrated via interfaces, consolidating them into a unified data pool that provides the foundation for further analysis. This data pool should ideally consist of several (hundreds of) thousands of sources to ensure a robust basis for searches.

- **Text extraction and annotation:**

  After raw data collection, the information is processed: systems extract the actual text from websites, PDFs, or emails and enrich it with metadata such as source, publication date, or document type. Frequently, an automated summary is generated to condense the key points of a document – this is a critical step for efficiently managing large volumes of text.

- **Contextual tagging and categorization:**

  Using NLP and Machine Learning, content can be categorized thematically. Industry-specific taxonomies are often used, systematically mapping terms, entities (like companies or government bodies), and topics. Sentiment analysis can additionally provide insights into the

tone and possible implications of regulatory changes – indicating, for example, whether a measure is likely to be restrictive or relieving.

- **Duplicate detection and cleanup:**

  In large datasets, repetition is inevitable – e.g., when multiple sources publish the same regulatory notice, slightly altered. AI systems detect such redundancies, merging or filtering out duplicate content, thereby reducing information overload without losing relevant insights.

- **Filtering and faceting of content:**

  Users can filter aggregated and structured information based on various criteria, such as topic, region, source, or affected institution. Faceting functions allow users to intuitively search through complex datasets and quickly navigate subsets of information.

- **Access and rights management:**

  In larger organizations, it's important that not everyone can view or edit all internal information. Modern systems offer granular access controls, providing different data visibility and editing permissions based on user roles or team affiliations.

- **Content creation and distribution:**

  Using structured data, curated content can be assembled automatically or semi-automatically for specific audiences – such as regular compliance newsletters, email alerts on specific topics, or individualized dashboards. Content is often delivered via discovery portals that enable personalization and topic favoriting.

- **Entity recognition, entity linking and knowledge graph construction:**

  A key feature is identifying and linking entities (e.g., authorities, laws, or companies). Using knowledge graphs, entities and their interrelations are modeled: for instance, a new regulatory decree is automatically linked to existing laws, affected industries, or companies. This allows complex relationships to be quickly identified, visualized, and operationally interpreted – for example, determining which compliance guidelines need to be updated.

- **Hybrid approach: Combining rule-based and AI-powered methods:**

  Many solutions combine traditional rule-based detection methods (e.g., keyword searches) with adaptive AI techniques (such as machine learning). This hybrid approach strikes a balance by providing reliable, explainable results while enabling flexibility to adapt to new topics and terms without requiring frequent manual adjustments.

b. Results and Impact

The use of AI-powered systems for identifying and analyzing regulatory changes is transforming how compliance teams operate. Instead of combing through vast masses of unstructured information, relevant developments today can be systematically identified, prioritized, and contextualized.

Practical examples show that companies can significantly reduce the time required for gathering and preparing information – especially in highly regulated sectors such as finance, healthcare,

and law. In one specific case in the healthcare industry, over 8,000 users were supplied with a personalized news and alerting portal, delivering over 25,000 relevant notifications daily – while halving the time spent on research compared to traditional solutions. In another example involving a global telecommunications company, over 80% time savings in the entire process of data extraction, annotation, and integration of research results into the customer's ecosystem were achieved.[9] The impact can be observed across multiple levels:

- *Faster response:* Regulatory developments are identified early, enabling companies to respond within hours instead of days.
- *Precise prioritization:* AI filters relevant information from the flood of news, enhancing efficiency and work quality.
- *Deeper understanding:* Knowledge graphs reveal systemic relationships and connections.
- *Scalability:* The more data sources and users are integrated, the more efficiently these systems operate – ideal for global organizations.
- *Greater consistency:* Automated classification reduces interpretative ambiguity and improves traceability

This also initiates a strategic shift in resource allocation: tasks that previously required time-intensive manual research and documentation are now automated. Compliance teams can focus more on analyzing implications, designing measures, and engaging with clients – thus playing a more active role in business development.

c. Challenges

Building a system for automated identification of regulatory changes and relevant news comes with significant technical and domain-specific challenges, including:

- *Real-time indexing of sensitive data without central storage:* Processing must occur at the data's point of origin without centralized storage, which imposes high demands on architecture, interfaces, and security.
- *Broad and targeted connectivity to relevant regulatory sources:* General news sources are not sufficient; domain-specific regulatory sources must be comprehensively covered.
- *Definition of precise semantic filters and taxonomies*
- *Correct integration and enrichment of information:* Logical linkage of internal and external datasets is critical.
- *Assessment of relevance and impact:* Beyond identification, systems must evaluate regulatory changes' implications on internal processes, policies, and risks.

[9] Data is based on project references from EPAM Systems

d. Success factors

Successful implementations address these challenges deliberately and focus on the following success factors:

- *Real-time, privacy-compliant architecture:* Systems enabling indexing without central storage preserve data sovereignty and security.
- *Targeted coverage of relevant sources:* Comprehensive domain-specific coverage of regulatory information sources is a key success factor.
- *High-quality semantic modeling:* Custom semantic filters and taxonomies provide the foundation for precise classification and relevance assessments.
- *Smart information linking:* Knowledge graph construction improves contextualization and traceability of regulatory documents and relationships.
- *Automated impact assessments:* Systems that not only detect regulatory changes but also assess their specific impacts on the business deliver real added value for decision-making.
- *Clean and structured data lineage:* Ensuring traceability of data sources and transformations maintains transparency, trust, and compliance requirements.

In summary, the success of an AI-powered system for regulatory and news analysis depends not only on technological excellence but also on the ability to

1. access relevant sources effectively,
2. precisely filter and enrich content,
3. intelligently assess impacts, and
4. effectively support users in their decision-making processes.

4.2 AI-powered reporting: From manual reports to intelligent reporting workflows

Companies face the challenge of creating compliance reports that are accurate, timely, and comprehensive. Traditionally, this process has been highly manual: data is pulled from various systems such as Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), or Document Management Systems (DMS), processed in spreadsheets, and transformed into reports. This approach is not only time-consuming – according to recent industry surveys, compliance teams spend up to 50% of their working hours on manual administrative tasks such as data collection and report preparation[10] – but also prone to errors. Manual data entry increases the risk of inconsistencies and non-compliance, while regulatory changes are often identified too late. Additionally, real-time transparency into the current compliance status is frequently lacking, which makes management even more challenging.

a. Functionality of AI-powered reporting

By integrating AI-powered tools into a modular system landscape, data from various sources can be automatically consolidated and analyzed. This significantly reduces manual effort, improves

[10] https://hyperproof.io/it-compliance-benchmarks/

efficiency, and ensures seamless compliance with regulatory requirements. AI-based solutions enable continuous monitoring of compliance processes, early detection of risks, and real-time report generation, allowing companies to respond more quickly and accurately to new requirements. Typical components of such a solution include:

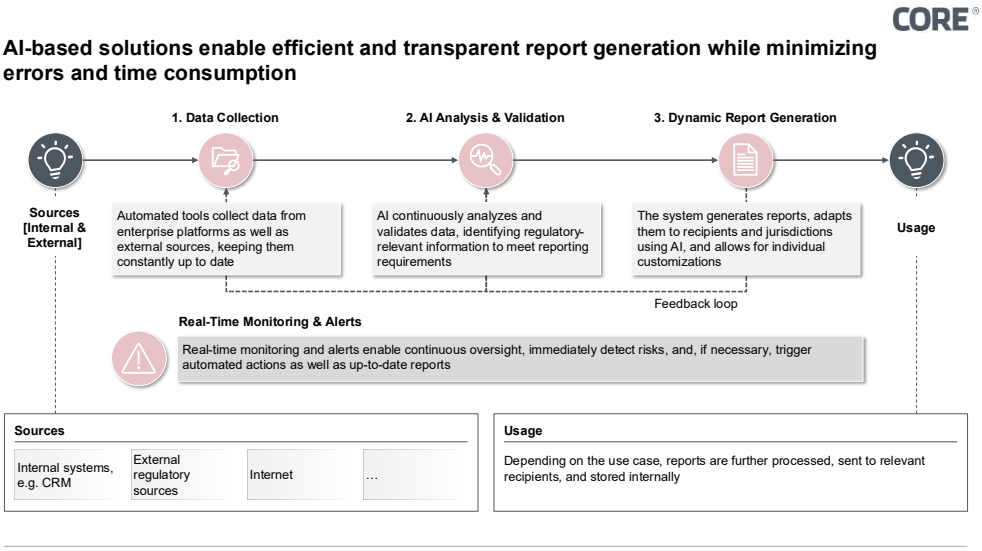

Figure 3: Functionality of AI-powered reporting

- **Automated data collection:**

  Automated systems integrate with corporate platforms like ERP, CRM, and DMS, continuously collecting data and establishing a centralized, up-to-date repository. External sources such as regulatory updates or industry benchmarks are also incorporated to ensure that the data remains aligned with current requirements.

- **Continuous AI analysis and validation:**

  AI assistants analyze and validate collected data, including unstructured formats like emails or PDFs. They identify regulatory-relevant information, process it, and ensure it meets reporting requirements. Using NLP, the AI interprets regulatory texts, identifies adjustments, and monitors updates to automatically adapt reporting templates and compliance controls.

- **Dynamic report creation:**

  Based on the analyzed data, the system generates compliance reports – whether in the form of dashboards, interactive reports, or traditional text documents. GenAI adds the ability to automatically tailor reports to meet the specific detail needs of recipients, such as regulators or the board of directors, while simultaneously creating versions for different jurisdictions, such as the EU and USA. Additionally, a low-code/no-code user interface allows employees to customize report structures as they prefer.

- **Real-time monitoring and alerts:**

  Automated compliance tools continuously monitor all relevant systems, identifying breaches or risks in real-time. For critical events such as rule violations or approaching reporting

deadlines, notifications are automatically triggered, along with immediate recommendations or corrective actions if needed. The data gathered feeds directly into dynamic reports, providing a current picture of the status quo, making them more up-to-date than traditional retrospective reports. Additionally, this information can be used for ad-hoc reporting.

b. Results and impact

AI-driven automation is transforming compliance reporting by replacing manual processes and error-prone workflows with dynamic, smart systems.

Case studies from highly regulated industries such as financial services and technology highlight measurable impacts: a Fintech company reduced the turnaround time for internal approvals and reporting by 60% through the use of AI agents, while significantly improving audit readiness. Another global financial institution automated risk and compliance reporting using AI tools that structured data across platforms, minimized manual intervention, and made processes more efficient, scalable, and less error-prone.[11] At the same time, AI systems improve the consistency of reporting by enabling standardized formatting and thorough traceability. The impact is visible across four strategic dimensions:

- *Efficiency gains:* Automated reporting processes significantly shorten the time required for data collection, validation, and report creation. This relieves compliance teams and allows them to focus more on strategic analyses and the development of preventive measures.
- *Error reduction and consistency:* AI-powered analyses minimize errors by largely eliminating human input mistakes and inconsistencies. Automated validation checks and real-time alerts ensure that discrepancies are identified and addressed immediately. This leads to higher quality and timeliness of reports, while the consistency across reporting periods, tailored to the intended audience, is improved.
- *Early risk detection:* Continuous evaluation and monitoring of compliance data allow potential risks to be identified early. AI systems analyze large data sets in real-time, recognize patterns, and proactively generate alerts before violations occur. This enables companies to respond quickly to new risks and implement targeted countermeasures.
- *Audit readiness and traceability:* The seamless, automated documentation of all compliance-related activities simplifies internal and external audits significantly. Digital audit trails and centrally available reports ensure fulfillment of audit requirements at any time. Additionally, this strengthens the trust of customers and regulators.

c. Challenges

The implementation of AI-powered compliance solutions offers numerous benefits but also introduces specific challenges that companies should address early to ensure long-term success.

- *Integration of heterogeneous systems:* The wide variety of different data formats, interfaces, and legacy systems complicates the smooth consolidation and processing of information. In

[11] https://springsapps.com/knowledge/top-compliance-problems-solved-by-ai-agents-real-world-case-studies

complex corporate landscapes, this results in fragmented workflows and higher maintenance and coordination efforts.

- *Data quality and timeliness:* The effectiveness of automated compliance reports is highly dependent on the quality and timeliness of the underlying data. Incomplete, inaccurate, or outdated information can lead to faulty analyses. Moreover, there is a risk that AI systems may generate false alarms or ambiguous results, requiring additional manual verification.
- *Data protection and security:* Processing sensitive and personal data within GRC workflows places high demands on privacy and IT security. It must be ensured that all data is protected in accordance with applicable privacy regulations (e.g., GDPR) and that AI systems adhere to principles such as Privacy-by-Design through anonymized data aggregation and localized processing of sensitive information.
- *Explainability and fairness:* AI models must be transparent, explainable, and free of bias or discrimination. This is particularly critical in creating reports that must remain impartial – such as whistleblower reports that should not favor the employer's position. In regulated industries, this introduces significant legal and reputational risks, requiring a high level of transparency and accountability for the deployed models.
- *Regulatory dynamics:* Regulatory requirements for AI are changing rapidly and becoming increasingly complex. Companies must continuously update their compliance processes and AI systems to align with new laws and standards, requiring additional resources and ongoing monitoring.

d. Success factors

To fully realize the potential of automating GRC processes, certain success factors are critical. These ensure not only technical functionality but also acceptance, security, and the future-proof nature of the solution.

- *Modular architecture and tool compatibility:* Flexibility within the system landscape allows for the exchange or expansion of individual modules without major system adjustments. Selecting tools with existing regulatory certifications (e.g., SOC 2, ISO 27001) simplifies implementation and builds additional trust.
- *Transparent AI algorithms:* Traceable and auditable AI decisions foster trust in automated reporting. Using Explainable AI ensures that decision-making pathways are documented and can be transparently reviewed by regulators.
- *User-friendliness:* Intuitive low-code/no-code interfaces empower departments to make individual adjustments quickly and independently. This reduces reliance on the IT department, avoids bottlenecks, and lowers IT costs by enabling purely configurative adjustments. Accompanying training sessions further enhance adoption and minimize application errors.
- *Continuous monitoring:* Real-time dashboards and automated audits ensure continuous compliance and quick responses to changes. Alert systems assist in the early detection and escalation of deviations.

- *Hybrid workflows and human oversight:* AI provides recommendations, while the final decisions remain with humans. This reduces risks of errors and ensures regulatory legitimacy – especially in critical areas such as financial reporting or patient data processing.

5. The question of the right solution – The AI market in Compliance management

There are already numerous AI-powered solutions available for the diverse use cases in the GRC domain, with varying levels of maturity and focus.

A structured market overview reveals that the landscape of vendors is highly diverse – both in terms of technological development and target audiences. While established GRC providers such as OneTrust, MetricStream, or NAVEX are increasingly embedding AI into their comprehensive platforms – which cover and integrate almost the entire value chain of GRC activities (whether regulatory monitoring, risk management, governance, or reporting) – specialized startups like Compliance.ai or Drata focus specifically on making AI the core of their GRC products. These focus areas might include policy monitoring, third-party risk, or continuous control monitoring. Another key differentiator is the breadth of GRC use cases covered. While many providers focus on individual use cases, such as risk management or reporting, some players like Compliance.ai, Vanta, or Archer go much further, aiming to cover as many use cases as possible across the GRC value chain. Overall, the market analysis demonstrates that all relevant GRC providers are integrating AI into their solutions, establishing this as a market standard:

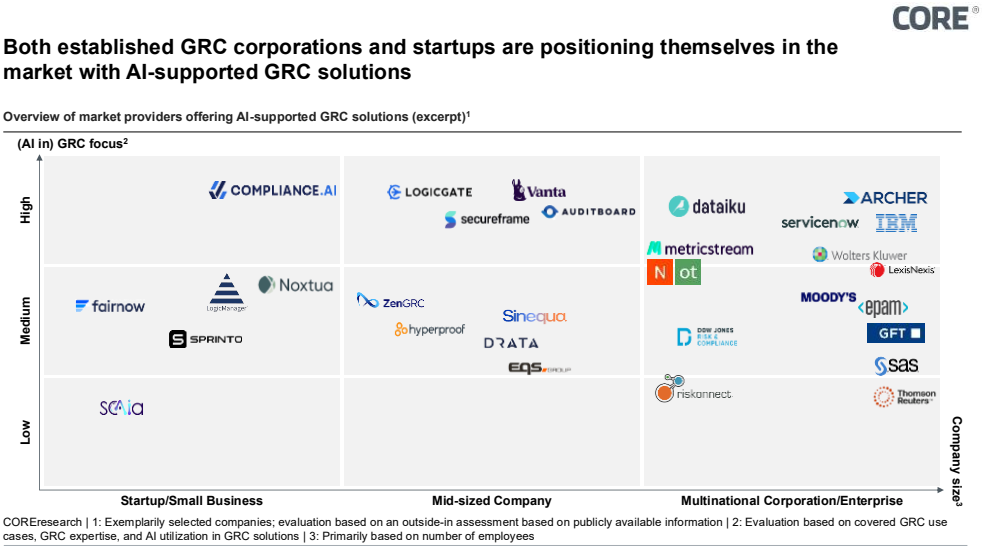

Figure 4: Providers of AI-supported GRC solutions

A market overview provides the foundation for identifying realistic entry points for leveraging AI in GRC processes – depending on individual requirements, internal data maturity, and existing IT infrastructure.

6. From vision to implementation: Success factors for strategic AI deployment in the GRC function

For AI to create long-term value in GRC practice rather than just having a limited impact, a robust strategic foundation is essential. Before deploying AI, companies need to clarify key framework conditions: What are the strategic goals and business priorities for leveraging AI? How can this be implemented in a legally compliant manner, particularly considering the risk strategy? How much automation makes sense – and how can control remain in human hands?

The first step is to develop a risk-based approach that defines the desired level of AI deployment. A structured decision tree can help make informed "Buy vs. Make vs. Reuse" decisions. The selection and prioritization of suitable use cases should also not be left to chance: companies should deliberately decide whether to start with a quick win or focus on areas where the strategic impact is greatest. Some solutions can be implemented quickly, while others require deeper integration into existing systems and processes. Beyond the time factor, cost considerations such as change-management expenses versus ongoing licensing costs also play a critical role.

The assessment of organizational and technological maturity is also decisive for success – evaluated across dimensions such as technology, processes, data, infrastructure, and governance. This maturity significantly determines how well AI integrates into existing architectures. Ideally, a centralized platform serves as the foundation, connecting all AI modules together to create synergies between compliance, IT security, and data management.

Another success factor lies in the collaboration between humans and machines. While AI handles repetitive tasks like risk classification or monitoring, the final assessment of complex regulatory issues remains in human hands. This hybrid model delivers efficiency without sacrificing the necessary accountability.

There is also the fact that many providers already have AI systems firmly embedded in their product portfolios. However, companies are often only in the early stages of actively using these solutions, and their acceptance must be strengthened through targeted communication and real-world examples. In most cases, there is already collaboration between vendors and organizations – but to fully realize the benefits, it is often necessary to expand existing licenses and more deeply integrate the systems into day-to-day operations.

7. Conclusion: How AI is transforming GRC for the long term

The regulatory framework surrounding AI itself will continue to evolve – from AI-assisted tasks to autonomous compliance systems. At the same time, increased oversight by regulatory bodies is to be expected. Companies are therefore well-advised to establish an AI-driven compliance strategy early on, one that not only focuses on efficiency gains but also builds long-term structures.

This includes addressing the implications for talent, technology, and operating models at an early stage. New role profiles, shifting responsibilities, and the development of AI skills in compliance are key levers for sustainable success. Companies that invest today in the right capabilities, architectures, and processes will lay the foundation for a future-proof compliance function.

Organizations should not hesitate to seek external support to systematically evaluate the complex questions surrounding AI-powered compliance solutions and identify the optimal strategic direction. As current market conditions demonstrate, the breadth of available solutions is no limiting factor – the range of robust AI systems is extensive and readily available. The challenge lies instead in selecting the right solutions for an organization's specific needs and integrating them seamlessly into existing structures.

When implemented correctly, AI can fundamentally transform compliance and IT security – from a reactive obligation to proactive, highly automated, efficient, and adaptive GRC processes. This enables companies not only to meet regulatory requirements efficiently and automatically but also to become more resilient and competitive.

Sources

1. Figure 1: Application areas of AI across GRC activities
CORE

2. Figure 2: Functionality of AI-powered news analysis
CORE

3. Figure 3: Functionality of AI-powered reporting
CORE

4. Figure 4: Providers of AI-supported GRC solutions
CORE

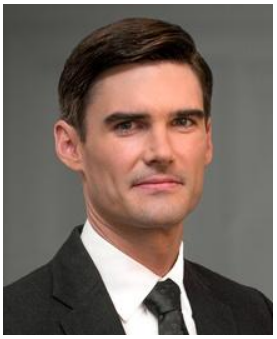

Dr. Philipp Kleine Jäger is responsible for Technology, Strategy & Management Consulting in Europe within EPAM's consulting unit CORE. With over 20 years of experience in the financial industry, including roles such as Lead Enterprise Architect at an international commercial bank, CTO of FinTech companies, and more than 10 years of consulting expertise in Digital, IT, Product & Operations, he advises clients at the executive level, focusing particularly on IT, digital, and transformation strategies.

Mail: philipp.kleine.jaeger@core.se

Dr. Carola Bader is an Expert Manager at CORE and has extensive experience in compliance and risk management within the financial industry. She advises financial service providers on developing and implementing effective compliance strategies to meet regulatory requirements and combat financial crime. With her strategic expertise, she helps organizations successfully implement regulations, minimize risks, and build resilient structures.

Mail: carola.bader@core.se

Muskaan Multani is a Transformation Manager at CORE. With her expertise in regulatory compliance, she primarily supports clients in the financial industry—from strategic planning to execution. Her academic background in Business Law as well as Strategy & Consulting provides her with a unique perspective on regulatory requirements and their organizational implications. Her focus is on supporting technology transformations critical to business success.

Mail: Muskaan.multani@core.se

COREtransform GmbH
Kurfürstendamm 194
10707 Berlin | Germany
https://core.se/
Phone: +49 30 263 440 20
sharedcoreoffice@epam.com

COREtransform GmbH
Limmatquai 1
8001 Zurich | Switzerland
https://core.se/
Phone: +41 44 261 0143
sharedcoreoffice@epam.com

COREtransform Ltd.
9 Devonshire Square, 5th Floor
London EC2 4YF
United Kingdom
https://core.se/
Phone: +44 20 328 563 61
sharedcoreoffice@epam.com