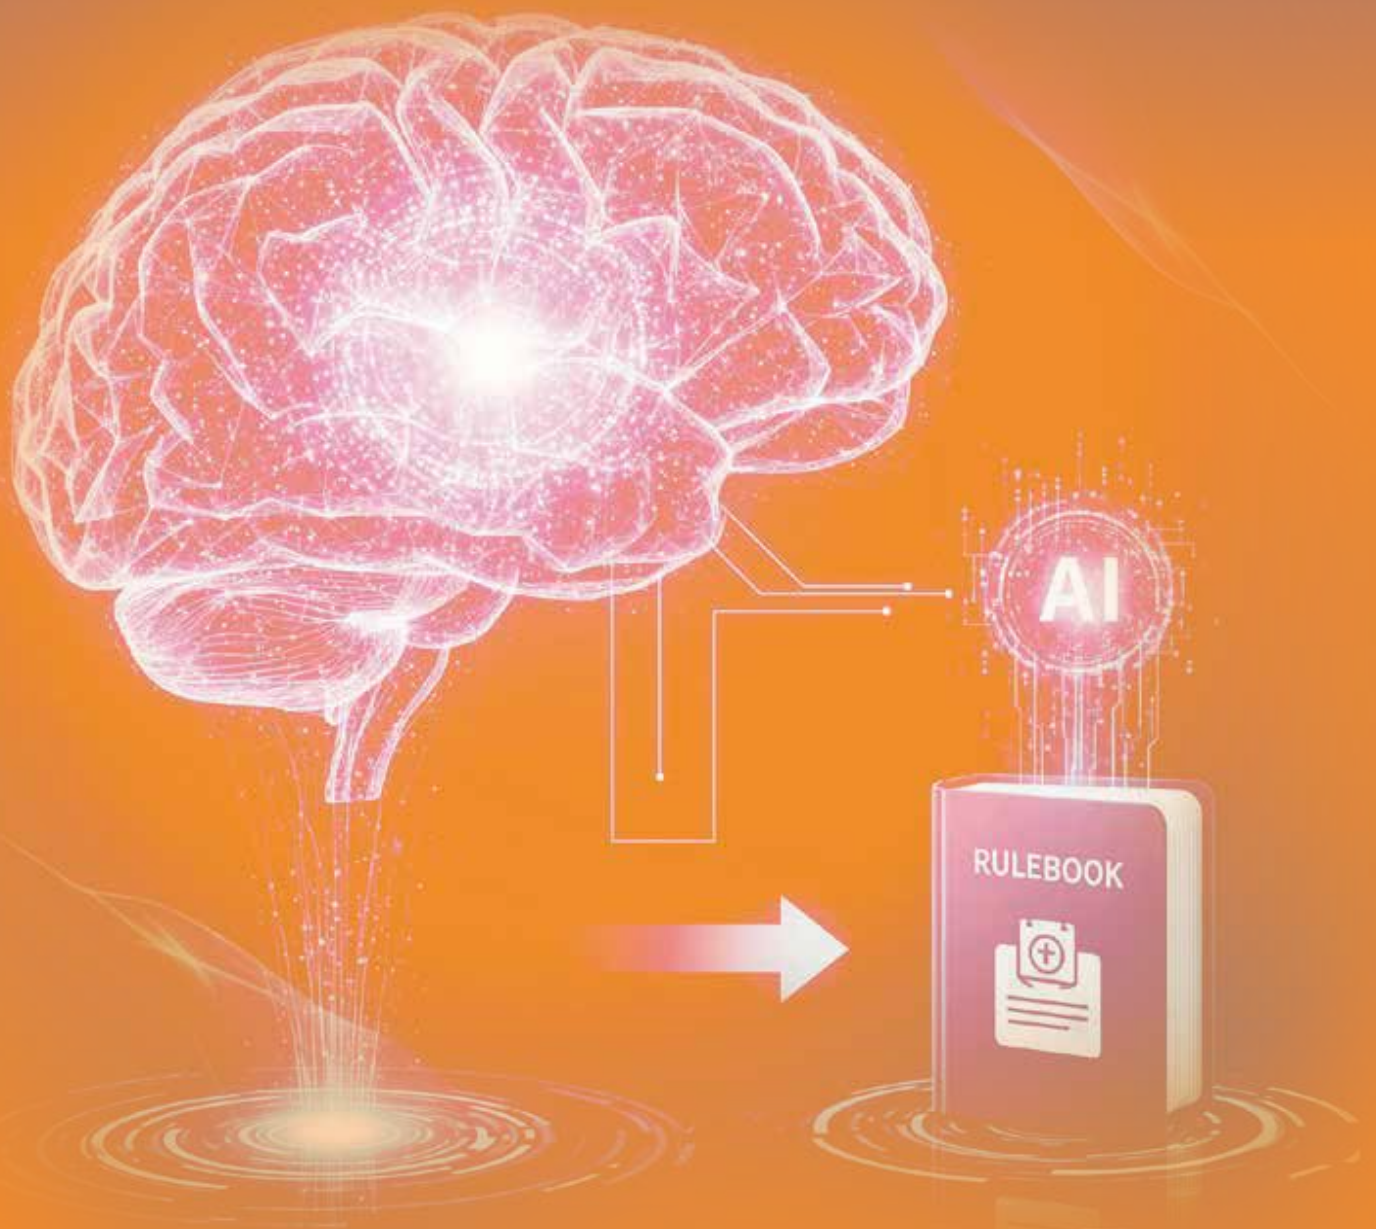




REGULATORY OVERSIGHT, UPGRADED

Agentic Regulatory Management for Transforming
Oversight from Rulebooks to Real Intelligence by
Newgen Software



Guardians at the Gate

Financial regulators have long served as the steady hands guiding banking, insurance, and broader economy. They have carried immense responsibility. Their aim has always been to ensure fairness, mitigate risks, and uphold the rules that safeguard our economy. They work tirelessly behind the scenes, but the tools they place their dependency on do not always match the magnitude of their task.

The Regulator's Burden

It's like a seasoned regulator sits at their desk, surrounded by paperwork, dense legal texts, decades of case files, hastily scribbled meeting notes, and ever-changing compliance reports. Each decision they make carries weight, requiring not just an understanding of black-and-white rules, but also the understanding to navigate gray areas. For instance, 'Was this exception granted before? What were the long-term consequences of that past ruling? Does this new fintech's model introduce unseen risks?'

All the required answers get buried in filing cabinets, long lost emails, and, most fragile of all, **HUMAN MEMORY**. New team members have to start from scratch, relearning lessons the organization already paid dearly to understand



It's an imperfect system, built on dedication, however strained by:

Pain Point	Core Challenge	Impact
Volume	Regulations change faster than teams can manually track them	Critical updates get missed; compliance gaps emerge
Consistency	No institutional memory leads to similar cases being handled differently	Unpredictable outcomes; fairness concerns arise
Transparency	Stakeholders can't see the reasoning behind decisions	Eroded trust in the regulatory process

Let's be real, this has been how things worked. But, what if it didn't have to be this way? Think of a:



A system that instantly recalls every relevant precedent, as easily as asking a trusted colleague



Decisions grounded not in gut calls, but in decades of institutional wisdom



Always-on supervision that spots and addresses issues in their earliest stages

This isn't a dream, it's the promise of intelligent, AI-backed governance. Where once there were piles of paper and uncertainty, now there's a scope to build CLARITY, CONSISTENCY, AND CONFIDENCE.

Transformation in Regulatory Oversight

It's remarkable how much is changing in the world of regulation, not with fanfare, but with practical, meaningful shifts. Generative AI and large language models are simplifying how oversight works, giving regulators the right tools to carry out their increasingly complex jobs in a smooth manner.

What makes this different isn't just the technology, but the way it aligns with what regulators struggle with every day. Enlisted here are the three core capabilities that stand out because they address real, persistent challenges:

Content Intelligence that Actually Understands Context

- Not just finding documents, but grasping why they matter
- Connecting past decisions to current cases in ways that make human sense
- Preserving institutional knowledge that used to walk out the door with retirees

Decision Support that Explains Itself

- Not replacing judgment, but making past rulings easier to understand
- Showing the 'why' behind historical rulings, not just the 'what'
- Making consistency achievable without rigid checklists

Oversight that Keeps Up

- Moving from periodic check-ins to always-on awareness
- Spotting deviations while there's still time to course-correct
- Turning reaction into prevention

This isn't about flashy tech solutions. It's about having ways to handle the overwhelming complexity that regulators face daily, with more clarity, less guesswork, and greater confidence that the right call is being made.

GenAI and LLMs are fundamentally reimagining regulatory oversight, empowering regulators with deeper contextual intelligence, evidence-based decision support, and end-to-end transparency. GenAI empowers regulators to manage high-stakes applications with confidence and auditability by integrating three key capabilities, including content intelligence, subjective-contextual decision support, and continuous oversight.

Let's understand how AI is transforming regulatory oversight

I. Unlocking Content Intelligence. Training AI on All Regulatory Touchpoints

Regulatory decisions rely on a sprawling web of documents, such as SOPs in PDFs, emails from applicants, transcripts of meetings, official circulars, scanned affidavits, case law, precedent-setting decisions, and handwritten annotations. All these remain scattered across systems, folders, and inboxes. When reviewing something as critical as a payment bank license, regulators have to manually sift through this disorganized mass, reconstructing context from fragments. This approach translates into slow processes, inconsistent rulings, and institutional knowledge that fades when experts leave.

GenAI solves this through content intelligence. It trains training LLMs on historical documents across all formats, such as PDFs, emails, internal memos, and public rulings. GenAI systems can:

- Ingest and **contextually label documents** across case history, rulebooks, and decisions
- Extract **judgment parameters** used in past approvals, such as financial health, promoter integrity, strategic relevance
- Convert past subjective decisions into **searchable, retrievable institutional knowledge**



LLMs can interpret the semantic logic behind decisions. Why was one fintech application rejected in 2018 but a similar one passed in 2020? A GenAI agent trained on case transcripts and regulatory commentary can identify those considerations, board structure, or revised supervisory frameworks, which influenced those outcomes. The real advantage is regulators don't have to 'start over' with each new case. They operate with historical depth, without requiring decades of personal experience.

II. Confident, Auditable Decisions. Subjective Guesswork Translated into Traceable Judgments

Today, regulators must assess not only whether an applicant meets formal conditions but also whether their appointment, structure, or expansion serves systemic stability, reputational safety, or ethical norms. These are subjective, multifactorial calls, especially in edge cases where past precedent exists but isn't codified in checklists.

Traditionally, such decisions rested on domain experts' memory or ad-hoc consensus. But, this approach introduced critical flaws, including:

Inconsistent Exceptions: Similar cases were treated differently based on who reviewed them

Black Box Reasoning: No clear trail explaining why an application was rejected or approved

Knowledge Silos: Non-transferability of decision intelligence to newer officers

GenAI addresses all these barriers by creating a dynamic decision-making companion that captures not only the outcome but also the reasoning path. AI agents:

- Suggest **confidence scores** for each application based on SOP adherence and similarity to past approved/rejected cases
- **Surface relevant precedents like** 'This mirrors the 2019 XYZ bank case where conditional approval was granted with additional oversight'
- Automatically generate **decision memos** with justifications mapped to the regulator's own documented principles

The outcome is a regulator who not only decides, but does so with explainable, defensible logic, supported by references and citations. This is not about removing human discretion. In this model, every decision becomes part of a learning loop. It's about ensuring that discretion is informed, transparent, and aligned with an organization's accumulated experience. The results are the decisions that stand up to scrutiny today while building institutional intelligence for tomorrow.



III. Oversight 2.0. Real-time, Context-aware Governance

Approval isn't the finish line, it's just the beginning of regulatory responsibility. Traditional monitoring methods (quarterly audits, periodic check-ins) find it difficult to match with today's dynamic financial landscape, where risks can emerge overnight. With GenAI, oversight becomes **continuous, contextual, and intelligent**. Here's how it makes a difference:

GenAI agents monitor incoming updates (e.g., new filings, KYC updates, board changes) and flag anomalies or non-compliance

Every new update is cross-checked against the **decision justification history**, ensuring the terms of approval are being met

AI auto-generates **summary alerts** for red flags (e.g., a new director's appointment that bypasses the declared succession plan)

This ensures that regulators don't just make the right decision once; they maintain the integrity of that decision throughout the entity's lifecycle. Furthermore, by linking all touchpoints, including applications, approvals, and post-approval changes, into a **unified Agentic Workdesk**, regulators avoid the typical 'system hopping' between silos. Case files, document intelligence, stakeholder inputs, and compliance history are all integrated, searchable, and audit-ready in a dynamic case management journey.

The Era of Agentic Regulatory Systems

GenAI is not just a tool for summarizing documents. It has become a cognitive partner, one that understands judgment, recalls institutional values, and offers guidance with nuance. The next phase is Agentic Regulatory Management, wherein there will be AI agents that will interpret rules and learn from context, simulate scenarios, and even recommend a regulatory stance with traceable logic. Such agents, trained on the full corpus of a regulator's institutional history, will:

Highlight policy precedents that support or contradict a decision

Explain edge-case logic, like 'This approval deviates from SOP because of XYZ systemic benefit'

Suggest mitigants, like 'conditional approval with quarterly disclosures' (backed by similar past cases)

AI serves as a trusted advisor, not a replacement, ensuring each judgment call is informed by history, aligned with values, and defensible in its reasoning.

Bringing the Vision to Life with Newgen with Newgen's Agentic Regulatory Management System

Three Pillars of Smarter Regulation:

1. Smart Content Engine

- Reads between the lines of PDFs, emails, case files and handwritten notes
- Organizes decades of regulatory history into actionable intelligence

2 Decision Confidence Builder

- Scores applications against historical patterns
- Recommends approvals/conditions with precedent-backed reasoning

3. Transparency Architecture

- Maintains full decision audit trails
- Explains AI suggestions in regulator-friendly terms

By institutionalizing intelligence and empowering decision-makers, NewgenONE Agentic Regulatory Management ensures that regulators don't lose the wisdom of past cases or the integrity of their frameworks. In doing so, it closes the loop between **knowledge, judgment, and accountability**, bringing regulatory management into the age of intelligent governance.



About Newgen

Newgen is the leading provider of an AI-first unified digital transformation platform with native process automation, content services, customer engagement, and AI/ML capabilities. Globally, successful enterprises rely on Newgen's industry-recognized low-code application platform to develop and deploy complex, content-driven, and customer-engaging business applications on the cloud. From onboarding to service requests, lending to underwriting, and for many more use cases across industries, Newgen unlocks simple with speed and agility.

For Sales Query

AMERICAS: +1 (202) 800 77 83
CANADA: +1 (202) 800 77 83
AUSTRALIA: +61 290 537174
INDIA: +91 11 407 73769
APAC: +65 3157 6189
MIDDLE EAST: +971 600 521 468
AFRICA: +27-11-461-6497
EUROPE: +44 (0) 2036 514805

info@newgensoft.com
www.newgensoft.com



July 2025