Suade

pwc

AGENTIC AI FOR REGULATORY REPORTING

FROM DATA TO DECISION-MAKING

This whitepaper is written for senior leaders across regulatory reporting, finance, risk, data, and technology who are considering how AI can be safely and practically introduced into regulatory reporting workflows.

It is particularly relevant for teams assessing how standardised data, interoperable architecture, and human-in-the-loop controls can support more scalable, explainable, and resilient reporting.

Table of Contents

Executive Summary

Banks face growing pressure from several directions at once. Reporting requirements are becoming more detailed, regulatory frameworks are evolving across jurisdictions, and supervisors are exploring more structured, machine-readable ways of collecting information.

Data is sourced from multiple systems, transformed through layers of logic, and reconciled under tight deadlines. That makes reporting costly to change, difficult to explain, and hard to scale across jurisdictions. This fragmentation is increasingly at odds with evolving supervisory expectations that increasingly require traceability, consistent definitions, and reusable datasets across use cases, treating data less as a reporting by-product and more as a controlled regulatory asset.

Banks should consider two responses to this challenge:

- **Data standardisation:** Banks need a clearer internal model for how regulatory data is defined, sourced, transformed, and controlled. Open standards and common input layers can reduce duplication, improve transparency, and provide a stronger foundation for regulatory agility. Data standardisation is becoming a design principle for enterprise data architecture, shaping how data is structured, governed and reused across finance, risk, and regulatory domains.

- **Utilisation of agentic AI:** Once data is structured and governed, AI agents can support multi-step reporting workflows by validating inputs, applying logic, surfacing exceptions, compiling evidence, and helping teams make decisions with greater speed and consistency. The use of AI is also expanding beyond banks themselves. Supervisors are moving towards more continuous, data-driven supervision, where AI and advanced analytics are used to interrogate granular datasets, monitor outcomes in near real time and challenge inconsistencies across reporting regimes.

Crucially, AI will only work at scale in regulatory reporting if banks first address the data standardisation challenge. To manage regulatory fragmentation, they also need an interoperable semantic layer that can support multiple frameworks, rather than a one-off solution tied to a single reporting model.

1. Regulatory reporting of the future

Regulatory reporting is shifting from a report-driven to a data-driven model. Today, many reports are still aggregated and constructed at the end of the process. Going forward, reporting will need to be derived from structured, granular data that is consistent across multiple regulatory uses.

This transition is being formalised through initiatives such as IReF (Integrated Reporting Framework) and BIRD (Banks' Integrated Reporting Dictionary) in the EU. In the UK, initiatives from the Bank of England and the FCA are reinforcing the same direction of travel. Together, they signal a structural redesign of regulatory reporting in Europe. These approaches seek to move away from static, siloed templates towards standardised, machine-readable data models, where a single granular dataset can support multiple regulatory outputs.

The challenge facing banks today

Regulatory reporting has become as much a data problem as a compliance problem. Reporting teams are expected to deliver increasing volumes of data across multiple regimes, with tight deadlines and rising expectations around data quality, transparency, and auditability.

In many firms, the process remains fragmented. Data is sourced from multiple internal systems, transformed through local logic, and passed between teams that do not always share a common view of the reporting process. This increases operational cost, slows down change, and makes it harder to explain how final numbers were produced.

At the same time, the operating model for reporting is changing. The question is no longer just how to submit the next report. It is how to build a reporting architecture that can adapt to new requirements and support more automation over time.

2. Why data standardisation matters

Banks cannot use AI effectively on top of fragmented, poorly defined data. Before reporting can become more automated or intelligent, firms need a structured way of defining, managing, and governing the data that underpins their reporting obligations.

This is where industry initiatives such as BIRD are relevant. BIRD helps banks identify which data to extract from internal systems and how to process it using redundancy-free concepts, transformation rules, mappings, and validation logic.

Three advantages of a standardised foundation

- Clearer definitions and lineage across the reporting process
- More efficient change management when rules evolve
- A stronger basis for automation and AI-enabled workflows

"Standardising data is not only a route to efficiency. It is increasingly a supervisory expectation. Regulators increasingly expect data to be accurate, traceable, and governed throughout its lifecycle, with the ability to reconcile across finance, risk, and regulatory views. This is consistent with the BCBS 239 principles on risk data aggregation and reporting, which emphasise accuracy, completeness, timeliness, and traceability."

Murat Abur, *CIO and co-founder,* Suade

Despite this trend across the industry, the challenge of handling many local standards and frameworks remains.

3. Interoperability and FIRE

As regulatory fragmentation is likely to continue, banks need to optimise for interoperability rather than wait for convergence around a single global reporting model. In practice, that means creating a semantic layer that can sit between internal systems and multiple external reporting frameworks.

The Financial Regulation (FIRE) Data Standard:

- provides a common specification for the transmission of granular regulatory data, acting as a common input layer that can help banks organise data once and reuse it across multiple reporting requirements;
- is designed around open, machine-readable, and legally anchored data attributes that do not depend on a single application or report; and
- is intended to be easy to read, understand, explain, and integrate with other IT systems, while remaining free and open source under the Apache 2.0 licence.

Used in this way, FIRE can help banks

- Reduce duplication across reporting regimes
- Create more consistent mappings from source systems to reporting outputs
- Centralise controls and documentation around a common semantic layer
- Build a stronger foundation for future automation and AI-enabled workflows

The value of FIRE is not that it replaces regulatory frameworks. It gives banks a regulator-agnostic internal language that makes those frameworks easier to interpret, map, govern, and manage over time. It maps a bank's internal data to legally defined terminology and harmonises that data across regulations and geographies.

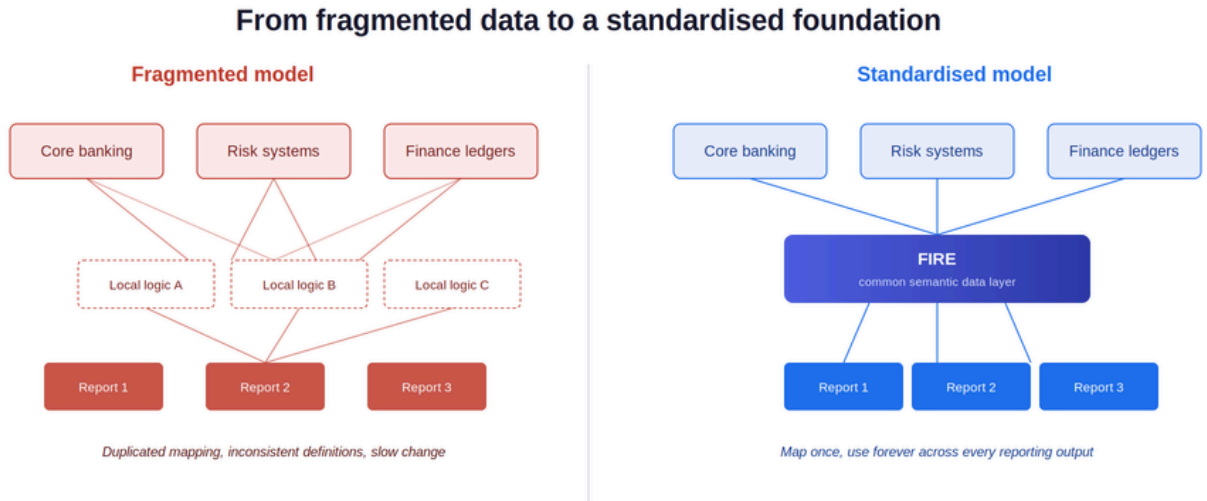

4. Agentic AI for regulatory reporting

If the first part of the challenge is getting regulatory data into the right shape, the second is deciding what to do with that data quickly, consistently, and with confidence. This is where agentic AI becomes relevant: it can execute multi-step workflows, apply logic, escalate issues, and support action across the reporting process.

Regulatory reporting is not one activity. It is a chain of steps that includes sourcing data, checking completeness, applying transformation rules, resolving exceptions, documenting evidence, and preparing submissions for review. Many of these tasks are repetitive and time-sensitive, but they still require judgement about what should be escalated, what can be resolved automatically, and what needs human approval.

Agentic AI offers a different model. Instead of viewing reporting as a set of disconnected manual tasks, it enables a workflow in which intelligent agents can coordinate actions across the lifecycle. One agent might check whether the required source data is present and complete. Another could apply transformation rules and reporting logic. A third might identify anomalies, attach relevant context, and route the issue to the right team. A fourth could compile the evidence needed for review and audit.

"The time to prepare is now. Once the models get good enough, the difference will be stark between firms that already have their data foundations in place and those still catching up. Agentic AI only creates real value in regulatory reporting when it operates on standardised, well-governed data within a clear control framework."

Murat Abur, *CIO and co-founder,* Suade

The real opportunities and benefits are not limited to speed. It is the ability to make the reporting process more adaptive, traceable and oriented around decision-making.

In that set-up, AI does not simply produce outputs faster. It helps banks prioritise, route, explain, and respond within a controlled workflow. AI does not replace existing control frameworks; it sits within them. In practice, this means integrating AI into model risk management structures, with clear requirements for explainability, monitoring and outcome validation, including checks for bias, robustness, and model drift.

5. Where agentic AI can add value

The most immediate opportunities sit in parts of the reporting lifecycle where complexity, repetition, and time pressure overlap.

Validation and completeness checks

Agents can assess whether the right data has been received, whether required inputs are missing, and whether records align with expected definitions before reporting outputs are generated.

Exception management

Instead of surfacing a generic error, agents can classify issues, attach relevant evidence, and route them to the appropriate owner with context on the likely cause and impact.

Audit trail generation

Agents can record which data was used, which rules were applied, what actions were taken, and where human intervention occurred, so that audit evidence is produced as part of the process rather than after it.

Regulatory change response

When requirements evolve, agents can help identify which data elements, mappings, and controls are affected, reducing the time needed to assess operational impact.

These use cases matter because they are practical. They show how agentic AI can begin to reduce manual effort, strengthen controls, and support faster responses to change without requiring a fully autonomous reporting model from day one. In practice, most firms are likely to start by piloting specific use cases in well-understood areas and only broaden coverage as governance, tooling, and confidence mature.

6. Governance, explainability, and trust

In regulatory reporting, automation is only useful if it remains governable. Outputs need to be explainable, reviewable, and defensible. That means agentic AI cannot operate as a black box. Banks need to understand what an agent did, what data it used, which rules it applied, what actions it took and when a human intervened. Audit trails, decision logs, and escalation paths are therefore core requirements, not optional extras.

Human-in-the-loop oversight remains essential. The purpose of agentic AI is not to remove accountability from reporting. It is to reduce manual workload while preserving control over material decisions, particularly where exceptions, sign-offs, and regulatory interactions are involved.

"There will always be, to some extent, humans in the loop – and that is an important piece of any regulatory reporting architecture. Agentic AI can automate repeatable steps and surface better evidence, but accountability for material decisions must remain with human experts."

Eamonn Clarke, Director, PwC Ireland

The most effective agentic models in banking are unlikely to be fully autonomous systems. They will be controlled set-ups in which AI handles repeatable steps, surfaces issues early and supports humans with better evidence and prioritisation.

In a regulated environment, the key question is not just whether AI can generate an answer. It is whether the bank can explain, evidence and stand behind that answer. In addition, whilst AI can automate reporting execution, human expertise remains critical to govern data changes, interpret new regulations, and validate that reporting outcomes remain compliant and useful for decision-making.

7. What an AI-ready reporting architecture looks like

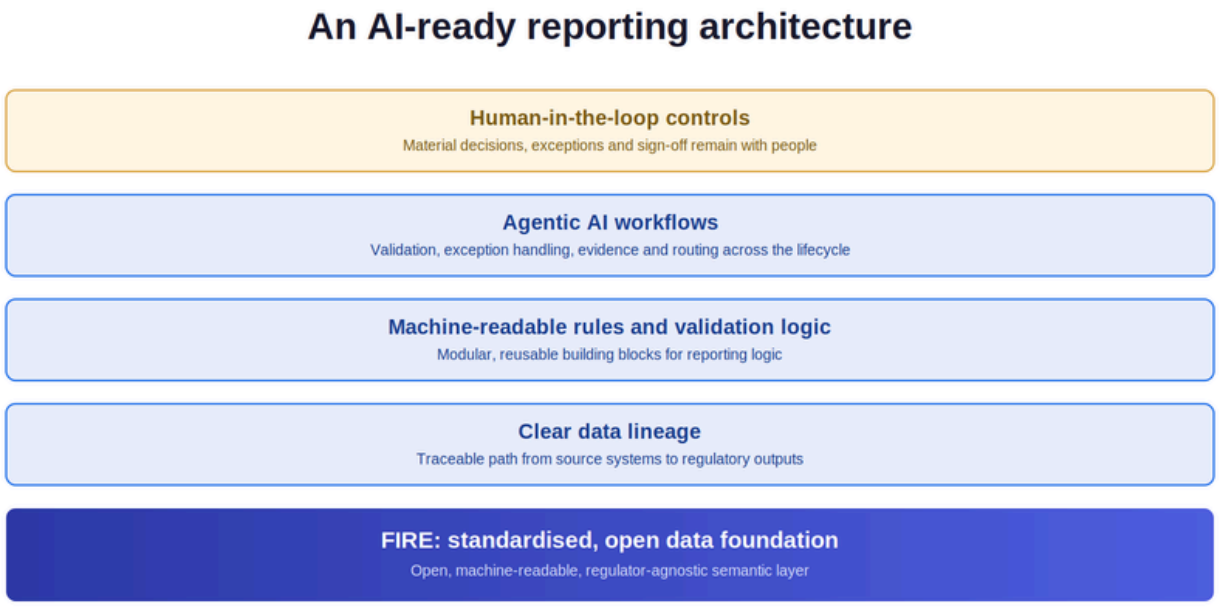

For agentic AI to deliver value in production, reporting systems need to be AI-usable by design. That means modular services, clear interfaces, machine-readable rules, and workflows that are explicit enough for both humans and AI agents to follow.

The most effective architecture is not AI layered on top of reporting. It is a reporting platform designed to support both human and AI-driven workflows from the outset. That includes:

- Modular, reusable building blocks for reporting logic
- Machine-readable transformation rules and validation checks
- Clear data lineage from source systems to regulatory outputs
- Human-in-the-loop controls for material decisions and exceptions
- Open APIs and interoperability with different tools and models

This kind of architecture gives banks more freedom over time. Rather than being locked into a single closed ecosystem, they can adopt AI in a governed way while keeping control over the underlying reporting model.

8. From data to decision-making

The longer-term opportunity is not simply to automate existing reporting tasks at lower cost. It is to reshape regulatory reporting as a more adaptive operating model in which data, logic, controls, and actions are connected more intelligently. That is what "from data to decision-making" means in practice. It means moving away from fragmented manual coordination towards governed decision support. It means using AI not just to process information, but to help route issues, prioritise responses, and provide better evidence at the points where judgement matters most.

For banks, the implication is clear. The institutions most likely to benefit from agentic AI in regulatory reporting will be those that combine strong data foundations with modular, interoperable architecture and robust controls. In that environment, AI becomes more than a productivity tool. It becomes a practical way to improve responsiveness, resilience, and confidence as reporting requirements continue to evolve. Adoption of AI can reduce manual burden without compromising accountability.

"Agentic AI only becomes usable in regulatory reporting when every action is tied back to governed data, reusable logic and enforceable controls. The goal is not just faster reporting, but a system where end-to-end lineage, scenario-level explainability and replayable decisions are standard, not exceptional."

Oliver Margetts, CTO, Suade

Questions banks should ask now

- Do we have a clear view of the data required across reporting regimes?
- Are definitions, transformations and validation rules documented in a reusable way?
- Can we trace data from source systems through to final regulatory outputs?
- Which parts of the reporting lifecycle are suitable for controlled AI pilots?
- Where must human review, sign-off and escalation remain mandatory?
- Can we evidence what an AI agent did, which data it used and why it took a specific action?

Conclusion

Regulatory reporting is becoming a design problem as much as a compliance obligation. Firms that invest in standardised, interoperable data, and a clear reporting architecture will be in a far better position to use agentic AI safely and at scale.

For heads of regulatory reporting and finance, the practical next step is not to “switch on AI”, but to ask whether their current data model, controls and processes are ready for it. Those foundations will determine whether AI becomes a controlled asset in your reporting stack, or another experiment that never makes it into production.

Together, Suade and PwC Ireland combine AI-enabled reporting technology with deep regulatory, data, and operating model expertise to help firms transform regulatory reporting end-to-end.

Authors and Contributors

This whitepaper was developed in collaboration by teams from Suade and PwC Ireland, bringing together expertise across regulatory reporting, data standards, and financial technology.

Authors
Murat Abur,
CIO and co-founder, Suade

John Dwyer,
Partner, PwC Ireland

Contributors
We would like to thank the following contributors for their input, review and expertise throughout the development of this whitepaper:

Oliver Margetts, CTO, Suade
Eamonn Clarke, Director, PwC Ireland

About Suade

Suade is a regulatory reporting platform used by financial institutions globally to automate reporting. Built on an API-first architecture and designed to be AI-native and cloud-native, Suade enables automation, programmability and secure integration with modern AI systems. It provides institutions with a controlled, scalable environment to manage data, calculations and reporting processes within modern finance.

Book a demo with a RegTech specialist

See how Suade helps financial institutions modernise regulatory reporting, improve control and manage change with transparent, data-driven technology.

Find out more on our website: www.suade.org

About PwC Ireland

PwC Ireland works with financial institutions on regulatory transformation, data, risk, and technology challenges, including the implications of AI for banking and regulatory compliance.