



WHITE PAPER

Turning BCBS 239 Compliance Into Competitive Advantage

Collate Helps Global Banks Aggregate and
Report Risk Data to Surpass Peers



The world's largest financial institutions face unprecedented risks. To help ensure they can meet these challenges, [BCBS 239](#) requires Global Systemically Important Banks ([G-SIBs](#)) to address 14 "Principles for effective risk data aggregation and risk reporting".

BCBS 239 doesn't rely on clearly defined metrics but calls on global banks to demonstrate their adherence to the spirit of 14 key principles.

Adopted in 2013 in response to the 2008 financial crisis, the Basel Committee on Banking Supervision's (BCBS) principle-based framework can no longer be ignored. Notably, however, like many new regulations around the world and in the U.S., BCBS 239 doesn't rely on clearly defined metrics but calls on G-SIBs to demonstrate their adherence to the spirit of its 14 principles.

Supervisors have increasingly applied BCBS 239 expectations to Domestically Systemically Important Banks (D-SIBs) and other significant institutions. The European Central Bank's (ECB) 2024 Risk Data Aggregation and Risk Reporting (RDARR) guidance effectively makes BCBS 239 a bank-wide standard touching risk reporting, financial reporting, and model development. Recent ECB and Bank for International Settlements (BIS) publications have also turned the originally principle-based framework into more concrete expectations around lineage, data quality, and material risk data element governance.

Banks that fail to do so could face serious consequences:

- Regulatory sanctions and capital surcharges for delayed or incomplete reporting
- Reputational damage if supervisors lose trust in the numbers
- Operational disruption as teams scramble to manually assemble fragmented data
- Competitive disadvantage compared to institutions that demonstrate greater resilience and agility.

Industry reviews show that only 2 of 31 G-SIBs have ever been assessed as fully compliant with all 14 BCBS 239 principles, while more than 40% remain materially non-compliant — highlighting how difficult and ongoing this work remains.

How can banks ensure they're able to deliver the accurate, complete, and timely risk data demanded by BCBS 239?

So, how can banks ensure they're able to deliver the accurate, complete, and timely risk data demanded by BCBS 239? It's critical to ensure your data is scrubbed and organized with clean metadata, so you can automate data lineage, implement common taxonomies and standards, automate quality control, and create integrated reporting platforms.

This guide to addressing BCBS 239's 14 principles — separated into 4 key areas — offers insight into what's required, what actions need to be taken, and how Collate's unified metadata and governance capabilities can be an essential component in leveraging the process to build a foundation for resilience, agility, AI-readiness, and competitive advantage.

Four Key Areas for BCBS 239 Compliance:

- I. Principles 1 and 2: Overarching Governance and Infrastructure
- II. Principles 3–6: Risk Data Aggregation Capabilities
- III. Principles 7–11: Risk Reporting Practices
- IV. Principles 12–14: Supervisory Review, Tools and Cooperation

Principles 1 and 2: Overarching Governance and Infrastructure

These general management principles call for institutions to establish strong governance and accountability, and build data architecture and IT infrastructure that fully support risk data aggregation and risk reporting even in times of stress or crisis. Clear ownership and standardized definitions enable banks to respond instantly and credibly to regulatory questions. Supervisors now explicitly tie these principles to the “three lines of defense” model, requiring transparent, auditable ownership structures for first-line data producers, second-line risk and compliance functions, and third-line internal audit.

The Challenge

Every critical dataset must have a clear owner, standardized definition, and auditable approval trail. Regulators increasingly ask not just “what are your results?” but “who signed off, and when?” In most financial institutions, however, data ownership is fragmented across lines of business, legacy systems complicate unified aggregation, and non-standard terminology can cause reporting inconsistencies.

What Banks Need to Do

To address the first two BCBS 239 principles, banks need to assign accountable owners for each critical dataset, publish a global Risk Data Glossary federated by domain, and configure workflows to enforce approvals for new or changed assets.

Banks must also identify and govern material risk data elements, defining their ownership, lineage, quality thresholds, SLAs, and monitoring requirements — as required by ECB RDARR and emerging global supervisory expectations.

Leading institutions now align their internal taxonomies with regulatory schemas (such as ECB BIRD or ISO 20022) to ensure consistency across jurisdictions and accelerate supervisory responses.

The Collate Solution

With Collate, data ownership is captured automatically in metadata, approvals are enforced by workflows, and versioned audit trails are maintained by default. Specifically:

- **Domain-based data ownership** lets banks assign accountable stewards by domain (e.g., credit risk, liquidity risk).
- **Collate's semantic metadata graph (built on the open source OpenMetadata project)** unifies technical lineage, business definitions, policies, classifications, and quality signals across databases, warehouses, Business Intelligence (BI) tools, and pipelines using open standards.
- **Support for Resource Description Framework (RDF) and Data Catalog Vocabulary (DCAT)** provide W3C Standards that allow banks to model definitions, reports, and governance evidence using open linked-data standards, and to align internal structures with regulatory taxonomies when required.
- **Governance workflows** let banks build in structured approvals, versioning, and role-based annotations.
- **Data contracts** (with version histories) let banks define schemas, SLAs, and usage rules with embedded audit trails/version history to track approvals and be audit ready.

Collate's capabilities support all three lines of defense by providing a single, authoritative source of truth for ownership, definition changes, lineage, and decision trails.

Together, these capabilities support all three lines of defense by providing a single, authoritative source of truth for ownership, definition changes, lineage, and decision trails.



Principles 3–6: Risk Data Aggregation Capabilities

(accuracy, integrity, completeness,
timeliness, and adaptability)

This area is where Collate’s capabilities make the biggest impact. These principles require banks to automate the generation of complete, accurate, and up-to-date data by business line, legal entity, asset type, industry, region and other groupings to meet normal and stress/crisis reporting requirements. Banks must also adapt to meet a broad range of on-demand, ad hoc risk management reporting requests.

The Challenge

Building a complete set of metadata, ingested from all the relevant databases, and tracking where everything comes from isn’t easy, but it’s critical to implementing BCBS 239. Broadly, banks’ existing static systems inherently resist new regulatory demands, and today’s batch processing data pipelines can delay reporting. Banks also need to ensure complete data coverage across all their systems, and avoid error-prone manual validation processes.

Supervisors increasingly expect banks to identify and govern the material risk data elements that support risk reporting and risk management, and to ensure these elements have clearly defined ownership, lineage, and quality expectations. Regulators also emphasize “shift-left validation,” meaning banks should detect and block data issues as early as possible in the lifecycle rather than relying on manual reconciliation downstream. As banks adopt AI and model-driven risk processes, regulators expect the same level of accuracy, integrity, and traceability for the data used to generate model inputs, features, and training datasets.

What Banks Need to Do

To ensure their BCBS 239 compliance is accurate, complete, and timely — and can be adapted to meet evolving requirements — banks must onboard all their data systems. This can include core systems using native connectors, and extended coverage with custom connectors for legacy applications. They must configure automated, real-time quality tests to address critical datasets, and enable “shift-left” validation of data earlier in their processes to block incorrect data in real time. Banks must also enable full audit trails (where all activity is tracked in server logs, including version history) and lineage tracking across sources, transformations, and reports, including Structured Query Language (SQL), Extract-Transform-Load (ETL), streaming, and BI layers. Granular, role-based access controls must apply to sensitive datasets, with an audit log for all data-access events. And they need extensible metadata models to cover ever-evolving regulatory requirements.

Banks should formally designate their material risk data elements across systems and ensure they are governed with consistent definitions, lineage visibility, quality thresholds, and validation rules — aligned to supervisory expectations without relying on vendor-specific terminology. For AI and model governance, banks must trace how material risk data flows into models, how features are constructed, which datasets are used for training, and how upstream data issues could influence model outputs or risk assessments.



The Collate Solution

Collate enables banks to do metadata aggregation right, a key component in BCBS 239 compliance. With Collate, accuracy, completeness, timeliness, and adaptability are embedded into the metadata layer — adding agility and resilience under both normal and high-stress conditions. Specifically:

- **100+ native connectors** for databases, warehouses, BI tools, and streaming platforms provide immediate coverage of the entire data estate.
- **Collate's custom connector framework** works with both legacy and proprietary systems.
- **Automated data quality checks** help ensure completeness and accuracy.
- **Column-level lineage tracking** across SQL, ETL, BI, and streaming systems provides complete visibility across sources, transformations, and reports.
- **Role-based access controls** apply granular permissions to sensitive data sets.
- **Collate's open standard and extensible metadata model**, built on the open source OpenMetadata project, supports new attributes as needed to address on-demand reporting requests.
- **Tagging and governing material risk data elements** directly within the semantic metadata graph, including lineage, definitions, quality rules, ownership, and SLAs in one place, ensuring they are governed consistently and traceably.
- **AI-assisted lineage automation and metadata classification** reduce manual effort and accelerate the bank's remediation cycle.
- **Feature lineage, model input flows, and training data provenance** support model risk management and explainability, ensuring AI-driven processes meet BCBS 239 expectations for accuracy and integrity.

With Collate, accuracy, completeness, timeliness, and adaptability are embedded into the metadata layer.

Principles 7–11: Risk Reporting Practices

(reporting accuracy, comprehensiveness, clarity, usefulness, frequency, and distribution)

These principles require concise risk management reports to clearly, accurately, and completely convey aggregated risk, including all material risk areas within the organization to facilitate informed decision-making. Reports should balance risk data, analysis, and meaningful explanations and interpretations. The required frequency of report distribution depends on the needs of banking supervisors, how quickly the risk can change, and the importance of the report to sound risk management — and the frequency must be increased in times of stress. The reports must be delivered to regulators quickly and confidentially.

The Challenge

In the real world, reports are often outdated or inconsistent, and can suffer from semantic misalignment across domains as well as poor traceability back to the source data. From an operational perspective, many institutions struggle to create and securely distribute multiple reports at the appropriate frequency, especially during a crisis.

Supervisors increasingly expect banks to demonstrate how each reported figure can be traced back to its underlying material risk data elements, including all transformations, adjustments, thresholds, and validation steps applied along the way.

Many institutions also struggle with KPI and metric inconsistencies because definitions differ across business units; this is often cited by regulators as a core obstacle to producing clear, comparable, and trustworthy risk reports.

With the rise of AI-assisted risk analytics, supervisors want evidence that model-generated outputs, summaries, or scenario projections can be traced back to governed data sources with the same reliability as traditional risk reporting.

What Banks Need to Do

Ensuring the accuracy and completeness of BCBS 239 reporting isn't easy. Among many other things, it requires banks to implement continuous quality validation using ongoing metadata tests to monitor reporting pipelines and build dashboards to track failures, coverage, and trends. Banks must ensure auditable lineage to reports by capturing transformations and calculations so they can be traced back to raw sources, while documenting the validation rules applied along the way. They must also standardize definitions across multiple sources by aligning the glossaries in BI tools and removing ambiguity in KPIs across business lines. Finally, automated distribution workflows are needed to enable approval-based report delivery while tracking receipt of the reports.

Banks should explicitly maintain standardized, organization-wide business definitions for metrics and risk indicators, ensuring that reports and dashboards draw from consistent semantic definitions rather than local or ad hoc variations.

Risk reporting processes should integrate lineage checkpoints that verify how material risk data elements flow into aggregated metrics, ratios, or exposure calculations. This ensures reporting teams can answer supervisory questions about data origins and transformations.

Banks increasingly need to embed governance context directly inside their business intelligence tools — surfacing lineage, quality status, and definitions within Tableau, Power BI, or Looker — so report creators cannot inadvertently use untrusted or outdated data.

BCBS 239 requires banks to implement continuous quality validation, ensure auditable lineage to reports, standardize definitions across multiple sources, and make use of automated distribution workflows.

The Collate Solution

Once your metadata is all scrubbed and clean, Collate is critical to creating accurate, frequent, and audit-ready risk reporting. Collate lets you tie governance evidence to every data asset. Specifically:

- **Continuous quality validation** enables ongoing metadata tests with dashboards and alerts to track data quality failures, coverage, and trends.
- **Glossary alignment** ensures reports use consistent definitions.
- **End-to-end lineage** tracking shows transformations from raw data to final metrics and applies document validation rules at each step. By surfacing lineage, quality status, and definitions directly inside BI tools via Browser extensions, Collate embeds governance “in the flow of work,” reducing reporting errors and eliminating manual reconciliation cycles.
- **Audit-ready approval workflows** automatically enforce Chief Risk Officer (CRO) sign-off for report releases.
- **Built-in Data Insights reporting** generates metadata-driven compliance reports monitoring readiness, failures, and coverage.
- **A single authoritative source for metric definitions** ensures all BI tools and reporting pipelines use consistent terminology across business units and jurisdictions.
- **Explainability and regulatory-grade lineage for AI-assisted reporting**, ensuring that any model outputs can be traced back to governed data sources.



Principles 12–14: Supervisory Review, Tools and Cooperation

The three final principles cover how institutions should support banking supervisors' review and evaluation of BCBS 239 compliance, and the tools and resources they need to work with supervisors in other jurisdictions to spur remedial action.

The Challenge

Ensuring that supervisors have access to the reports they need and the tools required to address the issues they encounter can be surprisingly complex. Missing audit trails can obscure actions covered in the reports, and a lack of structured remediation tracking complicates the process of fixing issues. In addition, in the global banking environment, cross-border compliance requires jurisdiction-specific controls.

Supervisors increasingly expect banks to provide machine-readable evidence of lineage, ownership, and data quality, as well as the ability to rapidly trace any reported figure back to its material risk data elements. Cross-border reporting obligations now require banks to align data structures and semantic definitions to regulatory taxonomies (e.g., ECB BIRD, ISO 20022), even though BCBS 239 does not mandate specific schemas. Banks must demonstrate not only transparency but also the capacity to remediate issues quickly, with evidence of root-cause analysis, accountability, and closure tracking

What Banks Need to Do

To navigate these complex interactions, banks must enable regulator-ready audit log exports, configure remediation dashboards, and define cross-border metadata access policies.

Banks should also maintain structured evidence packages — including lineage diagrams, approval histories, issue logs, and quality validation results — that supervisors can review without manual compilation. Institutions must establish governance models that support jurisdiction-specific metadata access controls, ensuring the right supervisors see the right information under the appropriate regulatory frameworks.

The Collate Solution

Banking supervisors expect full transparency, prompt and trackable remediation, and easy international cooperation. Collate provides regulator-ready evidence and scalable controls by design. Specifically:

Collate provides regulator-ready evidence and scalable controls by design.

- **Audit logs** capture every change, approval, and access event, and can be instantly exported. All activity is tracked in server logs, including version history.
- **Automated remediation workflows** assign tasks for data quality or governance issues and track resolution.
- **Ontology and semantic mapping** align internal data with regulatory taxonomies such as ECB BIRD and ISO 20022, enabling jurisdiction-specific access controls to be enforced consistently.
- **Machine-readable lineage and governance evidence**, enabling supervisors to trace metrics, reports, and controls back to underlying material risk data elements quickly and reliably.
- **Consistent definitions, structures, and mappings** across jurisdictions, even when supervisory authorities impose different reporting requirements, using Collate's semantic metadata graph
- **Complete evidence packages** — including lineage views, approval trails, issue logs, and quality results — allow institutions to respond to supervisory requests rapidly and confidently.



A Challenge — And An Opportunity

BCBS 239 was created to address gaps revealed in the last financial crisis. Today, its principles remain critical for stability and resilience. Living up to the 14 principles can seem daunting, but compliance is not the endgame, it's just the price of entry. With Collate, banks have the opportunity to turn these regulatory obligations into a platform for sustainable competitive advantage across multiple dimensions:

- **Operational efficiency:** Governance workflows, connectors, and data contracts reduce manual effort, improving decision-making, boosting change management, and freeing resources for strategic initiatives.
- **Agility:** Open and extensible metadata equips banks to address new stress scenarios and regulations.
- **AI-readiness:** High-quality, complete, semantically-aligned data powers explainable AI/ML. Collate's semantic intelligence architecture provides a consistent, explainable data layer aligned with supervisory expectations for transparency and auditability.
- **Transparency and trust:** Auditability by design builds credibility with supervisors and stakeholders.

More specifically, Collate helps financial institutions:

- **Prove ownership and governance** with accountable stewards and workflows.
- **Guarantee accuracy and completeness** with native connectors, quality checks, and lineage.
- **Deliver timely reports** with built-in Data Insights
- **Demonstrate transparency and cooperation** with audit logs, remediation, and federated access.
- **Enable AI agents and copilots to reason over trusted, well-defined data** by providing automated enrichment, lineage inference, classification, and policy recommendations.
- **Assure AI models meet regulatory requirements** with a transparent, explainable data layer.

BCBS 239 compliance is not the endgame, it's just the price of entry. Collate helps banks turn these regulatory obligations into a platform for competitive advantage.

Get Started

Ready to get started with transforming regulatory obligations into a competitive advantage? Talk with a Collate expert: sales@getcollate.io



Interested in learning more about Collate?
Email sales@getcollate.io to book a demo
with a product expert.

www.getcollate.io