Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

BCBS 239 Explained: The 14 Risk Data Principles

> Reliable Context is a kernel of truth: a system of record for organizations that fuse private and public information and must defend every value they act on. One canonical record per real thing, a cited source and an observation date behind every value, unknowns declared rather than left blank, and any past state of the record reproducible.

Published by Reliable Context · 2026-10-29 · bcbs 239 · bcbs 239 principles · risk data aggregation and reporting · bcbs 239 compliance data lineage · basel risk data aggregation

Original article: https://reliablecontext.com/articles/bcbs-239-risk-data-aggregation-principles-explained



In this article

1. The short answer
2. Why it was written
3. Scope and dates
4. The 14 principles
5. Worked example: one counterparty, two booking systems
6. What supervisors report in practice
7. Mapping each principle to a checkable property
8. Self-assessment checklist
9. Sources

The short answer

BCBS 239 is the Basel Committee on Banking Supervision's *Principles for effective risk data aggregation and risk reporting*, published in January 2013. It sets out 14 principles in four groups: overarching governance and infrastructure (Principles 1-2), risk data aggregation capabilities (3-6), risk reporting practices (7-11), and supervisory review, tools and cooperation (12-14) (BCBS 239, ¶20 and Annex 2). Banks identified as global systemically important banks (G-SIBs) in 2011 or 2012 were to meet them by January 2016. More than a decade later, the Committee's latest progress report found that only two of 31 assessed G-SIBs were fully compliant with all the principles (BCBS progress report, November 2023).

The principles describe outcomes, not a data model. This article quotes each one, summarises what supervisors report, and turns each principle into a property an examiner can test. BCBS 239 is an international standard; what binds a bank depends on its national supervisor's implementation. This is educational, not legal or regulatory advice.

Why it was written

The document opens with the lesson of the crisis that began in 2007: banks' "information technology (IT) and data architectures were inadequate to support the broad management of financial risks. Many banks lacked the ability to aggregate risk exposures and identify concentrations quickly and accurately at the bank group level, across business lines and between legal entities" (BCBS 239, ¶1). It cites a parallel initiative to build the Legal Entity Identifier (LEI) system as "a key building block for improvements in the quality of financial data" (¶6).

Two definitions frame the rest. **Risk data aggregation** means "defining, gathering and processing risk data according to the bank's risk reporting requirements", including "sorting, merging or breaking down sets of data" (¶8). **Reconciliation** means "the process of comparing items or outcomes and explaining the differences" (BCBS 239, fn 17 and Annex 1). A difference found but not explained is an incomplete reconciliation.

Scope and dates

- The principles "are initially addressed to SIBs and apply at both the banking group and on a solo basis". National supervisors "may nevertheless choose to apply the Principles to a wider range of banks, in a way that is proportionate to the size, nature and complexity of these banks' operations" (BCBS 239, ¶13).
- Banks identified as G-SIBs by the Financial Stability Board in November 2011 or November 2012 "must meet these Principles by January 2016"; G-SIBs designated later have three years from designation (¶14).
- For domestic systemically important banks (D-SIBs), "it is strongly suggested that national supervisors also apply these Principles ... three years after their designation as D-SIBs" (¶15).
- They also cover "all key internal risk management models" (¶17) and outsourced processes (¶19).

Banks "should meet all risk data aggregation and risk reporting principles simultaneously" (¶22), and may omit information only if it would not have changed a board or senior management decision; "supervisors expect banks to be able to explain the omissions" (¶23).

The 14 principles

I. Overarching governance and infrastructure

Principle 1, Governance. Risk data aggregation and reporting "should be subject to strong governance arrangements". They must be "fully documented and subject to high standards of validation", independently validated, and "unaffected by the bank's group structure" (¶29).

Principle 2, Data architecture and IT infrastructure. Infrastructure must support aggregation "not only in normal times but also during times of stress or crisis". Banks should have "integrated data taxonomies and architecture across the banking group", including metadata and "single identifiers and/or unified naming conventions for data including legal entities, counterparties, customers and accounts" (¶33). A footnote adds that banks "do not necessarily need to have one data model; rather, there should be robust automated reconciliation procedures where multiple models are in use" (fn 16).

II. Risk data aggregation capabilities

Principle 3, Accuracy and Integrity. Banks should generate "accurate and reliable risk data", aggregated "on a largely automated basis so as to minimise the probability of errors". Controls on risk data "should be as robust as those applicable to accounting data"; risk data "should be reconciled with bank's sources, including accounting data where appropriate"; and a bank "should strive towards a single authoritative source for risk data per each type of risk" (¶36). A "dictionary" of concepts is a precondition (¶37).

Principle 4, Completeness. Banks should "capture and aggregate all material risk data across the banking group", available "by business line, legal entity, asset type, industry, region and other groupings", including off-balance-sheet exposures (¶41). Data should be "materially complete, with any exceptions identified and explained" (¶43).

Principle 5, Timeliness. Banks should produce "aggregate and up-to-date risk data in a timely manner" while still meeting the accuracy, completeness and adaptability principles. Critical risks include "the aggregated credit exposure to a large corporate borrower" and "counterparty credit risk exposures, including, for example, derivatives" (¶46).

Principle 6, Adaptability. Banks should meet "a broad range of on-demand, ad hoc risk management reporting requests", including in stress and for supervisory queries. The example given: aggregate country credit exposures "as of a specified date based on a list of countries" (¶50).

III. Risk reporting practices

Principle 7, Accuracy. Reports "should be reconciled and validated". The minimum includes "an inventory of the validation rules" and exceptions reports for data errors (¶53).

Principle 8, Comprehensiveness. Reports "should cover all material risk areas within the organisation", including single name, country and industry concentrations (¶57).

Principle 9, Clarity and usefulness. Reports should be "clear and concise" and "tailored to the needs of the recipients". The bank "should develop an inventory and classification of risk data items" with reference to the concepts used (¶67).

Principle 10, Frequency. The board and senior management set report frequency, which "should be increased during times of stress/crisis". A bank "should routinely test its ability to produce accurate reports within established timeframes" (¶70).

Principle 11, Distribution. Reports "should be distributed to the relevant parties while ensuring confidentiality is maintained" (¶72-73).

IV. Supervisory review, tools and cooperation

Principle 12, Review. Supervisors "should periodically review and evaluate a bank's compliance with the eleven Principles above", including through requests "within short deadlines" that test rapid aggregation (¶75).

Principle 13, Remedial actions and supervisory measures. Supervisors should be able to require timely remediation, using tools that may include "capital add-ons" under Pillar 2 and limits on risks or growth (¶79-80).

Principle 14, Home/host cooperation. Supervisors in different jurisdictions should cooperate on review and remediation (¶83).

Principles 12-14 address supervisors. For a bank they define the test: a specific aggregation, requested at short notice.

Worked example: one counterparty, two booking systems

Aldermoor Bank plc is a fictional bank; Kestrel Shipping Holdings and Kestrel Tankers are fictional companies, and the identifiers are invented placeholders, not valid LEIs. Amounts are in EUR millions.

A supervisor asks Aldermoor for its total credit exposure to the Kestrel group as of 30 September 2026, within 48 hours: a Principle 6 ad hoc request about a ¶46 critical risk. Two booking systems hold the positions:

System	Legal entity booking	Counterparty name as stored	Internal ID	LEI on record	Exposure
Corporate lending	Aldermoor London	Kestrel Shipping Holdings Ltd	C-48812	`XXXX00KESTRELEXMPL00`	Loan drawn 120.0; undrawn commitment 30.0
Derivatives	Aldermoor Frankfurt	KESTREL SHIPPING HLDGS	CP-0091	none	Net exposure after collateral 18.0
Corporate lending	Aldermoor London	Kestrel Tankers AS	C-51207	`XXXX00KESTRELTANKR00`	Loan drawn 25.0

The LEI is "a unique 20-character alphanumeric code", and "every LEI is unique and can represent only one entity"; its reference data answers both "who is who" and "who owns whom" (GLEIF, Introducing the LEI).

Step 1, name-based aggregation. Grouping on the stored name finds only C-48812: 150.0 (120.0 drawn plus 30.0 undrawn). The derivatives record (different name string) and Kestrel Tankers (separate legal entity) are missed. The undrawn commitment is correctly included (¶41 covers off-balance-sheet exposures), but the total fails Principle 4 (completeness).

Step 2, identifier resolution. CP-0091 has no LEI. An analyst matches it to `XXXX00KESTRELEXMPL00` using the registered address and company number in the LEI reference record, retrieved that day. The match is stored as a dated, attributed assertion with its evidence, not as an overwrite. Single-counterparty exposure becomes 168.0 (120.0 + 30.0 + 18.0).

Step 3, group aggregation. The LEI ownership data for Kestrel Tankers AS names Kestrel Shipping Holdings Ltd as its direct parent. Group exposure becomes 193.0.

Step 4, reconciliation to accounting. Principle 3 asks for reconciliation with accounting data. The ledger shows the London loans at 145.0, agreeing with 120.0 + 25.0. The ledger carries the Kestrel derivatives at a gross positive fair value of 22.0; the risk figure is 18.0. The 4.0 difference is explained by collateral received under the credit support annex. The reconciliation is complete because the difference is explained, not because it is small.

Aggregation method	Kestrel exposure	Principle at risk if reported
By stored name	150.0	4 (completeness), 7 (report accuracy)
By LEI, single entity	168.0	4, if the board's limit is set at group level
By LEI with parent link	193.0	None identified; reconciled with explained difference

The 43.0 gap between the first and last rows is the kind of error the 2013 introduction describes. Once both systems carry the same LEI, as ¶33 asks, the aggregation is automatic (Principle 3) and fast (Principle 5). The site's LEI standard page covers the identifier itself.

What supervisors report in practice

The Basel Committee's November 2023 report rated 31 G-SIBs on each principle from 1 (not adopted) to 4 (fully complied with), using data as of June 2022:

- "Of the 31 banks assessed, only two banks are fully compliant with all the Principles", and "there is not a single Principle that has been fully implemented across all banks" (progress report, executive summary).
- Average ratings for Principles 5 (timeliness) and 7 (report accuracy) deteriorated between 2019 and 2022 (§2.1).
- "Several banks still lack a common taxonomy and complete data lineage, which further complicates banks' ability to harmonise systems and detect data defects" (§2.2.2).
- Supervisors use onsite examinations, deep dives and "fire drills" that "test banks' ability to aggregate and report risk data under simulated stress scenarios", and should consider "more forceful measures" such as capital add-ons.

The report also states that "not all jurisdictions actually require compliance with the Principles since certain jurisdictions did not transpose them into their national framework" (fn 1).

In the euro area, the ECB's *Guide on effective risk data aggregation and risk reporting* (May 2024) sets out its "minimum supervisory expectations" for significant institutions, applying national transpositions of the Capital Requirements Directive. It "does not impose new requirements" and is to be read with BCBS 239 (ECB Guide, §1). It is more concrete than the 2013 text:

- Data owners are named for key risk indicators and their **critical data elements**, defined as data elements "used to calculate the key risk indicators" with a direct or significant impact on their value (§3.2-3.3).
- An integrated data architecture includes "complete and up-to-date data lineages on data attribute level (starting from data capture and including extraction, transformation and loading)". Lineage is defined as the movement and transformation of data "from front (capture) to end" (§3.4).
- Data quality indicators have tolerance levels. A "register" of data quality issues records severity, root cause, quantitative impact, deadlines, and "a date for effective remediation (including appropriate evidence)". Manual workarounds require controls such as "audit trailing of changes, data overrides and sign-offs" (§3.5).
- On timeliness, institutions "will not be able to react to changes in a timely manner if a monthly or quarterly risk report needs more than 20 working days to produce" (§3.6).

Mapping each principle to a checkable property

Each principle can be restated as a property of a system of record that an examiner can test with a query rather than a presentation.

#	Principle	Checkable property of the system of record
1	Governance	Every risk data element has a named owner; an independent validation report is on file

#	Principle	Checkable property of the system of record
2	Data architecture and IT infrastructure	Counterparties, entities and accounts carry one identifier group-wide; cross-model reconciliations are automated and stored
3	Accuracy and integrity	Each aggregate traces to source records; every change records who, when and why; ledger reconciliations have no unexplained differences
4	Completeness	A query lists any in-scope entity, book or off-balance-sheet product missing from the aggregate, with the reason
5	Timeliness	Each critical figure carries an as-of time and a production time, measured against a stated target
6	Adaptability	An exposure can be re-aggregated by any defined dimension "as of a specified date" without new code
7	Report accuracy	Each report figure reconciles to risk data, and the validation-rule inventory and exceptions log are retrievable
8	Comprehensiveness	The report inventory maps every material risk type to at least one report
9	Clarity and usefulness	Report items link to dictionary definitions; recipient feedback is recorded periodically
10	Frequency	Frequencies for normal and stress conditions are recorded, and test runs are logged with times
11	Distribution	Distribution lists, access controls and receipt confirmations are recorded
12	Review	Any past report can be reproduced as it was at the time, from retained inputs
13	Remediation	Each finding has an owner, deadline and evidence of closure in an issue register
14	Home/host cooperation	Entity-level data can be produced for a host supervisor without manual re-extraction

Rows 3 and 12 carry most of the weight. An aggregate that cannot be traced to its inputs, or re-run as of a past date, cannot be demonstrated after the fact; that is the practical difference between a warehouse holding the latest load and a record keeping the history behind each value (data warehouse vs system of record).

On models: ¶17 brings key risk models into scope, but US model risk management has separate guidance; Federal Reserve SR 26-2 (17 April 2026) "supersedes and replaces SR letter 11-7".

Self-assessment checklist

- The scope document names the legal entities, reports, models and key risk indicators covered, consistent with ¶13-19 and, in the euro area, ECB Guide §3.2.
- Counterparty and legal-entity records carry an LEI where one exists, and every match to an LEI is dated, attributed and evidenced.
- Reconciliations to the ledger are run, stored, and every difference is explained.
- Production times for monthly and quarterly risk reports are measured, and stress drills are run and logged.
- Any past report can be regenerated as of its original date.

Reliable Context publishes a related guarantee for organizational facts: one canonical record per real thing, with a cited, verbatim, dated source behind every value and any past state reproducible.

This article summarizes Basel Committee and ECB supervisory texts as of 28 October 2026 for educational purposes and is not legal or regulatory advice.

Sources

1. Basel Committee on Banking Supervision, Principles for effective risk data aggregation and risk reporting (BCBS 239), January 2013 (PDF) - text of the 14 principles; definitions; scope and dates (¶13-15); trade-offs and materiality (¶22-23); Annex 1 terms.
2. Basel Committee on Banking Supervision, Progress in adopting the Principles for effective risk data aggregation and risk reporting, November 2023 (PDF) - 2022 assessment of 31 G-SIBs; key challenges; supervisory activities and recommendations.
3. ECB Banking Supervision, Guide on effective risk data aggregation and risk reporting, May 2024 (PDF) - minimum supervisory expectations for scope, data owners, critical data elements, lineage, data quality register and the 20-working-day reference.
4. GLEIF, Introducing the Legal Entity Identifier (LEI) - LEI definition, uniqueness, "who is who" and "who owns whom", ISO 17442.
5. Board of Governors of the Federal Reserve System, SR 26-2, Revised Guidance on Model Risk Management, 17 April 2026 (PDF) - supersedes SR 11-7.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.bis.org/publications/201301-guidelines-principles-effective-risk-data-aggregation-and-risk-reporting.pdf
2. https://www.bis.org/publications/202311-implementation-reports-progress-adopting-principles-effective-risk-data-aggregation-and-risk-reporting.pdf
3. https://www.gleif.org/en/about-lei/introducing-the-legal-entity-identifier-lei
4. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.supervisory_guides240503_riskreporting.en.pdf
5. https://www.federalreserve.gov/supervisionreg/srletters/SR2602.pdf

THE PUBLISHER

About Reliable Context

Reliable Context is a system of record for organizations that fuse private and public information and must defend every value they act on.

The guarantees

One canonical record per real thing: duplicates resolve by rule, ambiguity is shown rather than hidden, and related organizations relate without ever being merged. A verbatim, cited source and an observation date sit behind every value. Unknowns are declared as not found, blocked or not applicable with a reason, never left as an empty cell. Any past state of the record is reproducible, together with what changed since.

Where it runs

It runs inside the customer's perimeter, one database per organization. Permitted sources, model routing, reader clearance and retention are set as the customer's own policies, and every value is exportable with its chain of custody for an auditor or regulator.

What it is not

It is not a vector database or a RAG framework, not a data catalog, MDM or data-quality suite, not a CRM or data-enrichment vendor, and not a chatbot or AI assistant. Where a reader needs one of those, the honest answer is to say so and point them there.

Visit

See Reliable Context.

Reliable Context is in production use in two sectors; more than six thousand of its identity decisions were made by rule and none has been reversed. Articles published here explain general practice, cite their primary and academic sources, and name the jurisdiction of any regulatory rule they describe. They are not legal, regulatory or compliance advice for a specific organization, and requirements change.

Website: https://reliablecontext.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.