




AI-Written, AI-Checked, **AI-Readable:**

the new standard for PDF
client reporting in private banking

**Automate narrative, strengthen controls,
make reports usable by humans and machines:
without losing trust, traceability or human judgement**

Structured data in, audit-ready reports out.



1. Introduction: Why AI-Ready Client Reports Pay Off Now

For decades the client PDF has been the final artefact of the reporting process: a carefully paginated, branded and legally durable representation of approved information. That role remains essential. The document is now entering a second life on top of it.

Relationship managers upload reports into assistants before a client call. Clients ask conversational tools to explain a performance number. Compliance teams compare a report with previous ones. Operations teams search historical reports to understand what changed. In parallel, banks are exploring AI-generated commentary to reduce manual effort and improve consistency. The PDF is becoming an interface between the bank, the client, the relationship manager, the archive and the next generation of AI systems.

This changes the design question. A bank should no longer ask only whether a report looks correct to a human reader. It should ask whether the report can be written safely with AI, checked systematically before release and understood reliably by AI after publication.

THE CENTRAL THESIS

The next generation of client reports must be AI-written, AI-checked and AI-readable, without losing traceability, accessibility or human control.

About ReportingSoft

ReportingSoft SA, based in Lausanne, automates client and business reporting for banks, asset managers and wealth management firms. Its platform, Reporting Maestro with the self-service template layer Business Report Designer, consolidates data from core systems such as Avaloq, Triple'A, Finnova, Olympic and T24, produces branded documents at scale and manages controls, review, distribution and archiving in one governed process.





Patrick Cornaz

Founder and CEO, ReportingSoft SA

Patrick is ReportingSoft's CEO, with a successful track record of more than 10 years in managing software and consulting companies. He is also a recognised technical project manager and technical architect with 15 years of experience, including 4 years at Accenture and 6 years at Actuate. Patrick has a special focus on business intelligence, Java/J2EE/web applications and systems integration within the financial services industry.

Your contact: info@reportingsoft.ch · +41 21 791 54 87

What leading banks should aim for

- ✓ **Faster reporting cycles:** Generate first-draft commentary from approved figures and context, while keeping the reviewer in control.
- ✓ **Higher pre-publication assurance:** Combine deterministic controls with AI-based comparison of text, tables, charts, periods, currencies and benchmarks.
- ✓ **A more useful document after publication:** Preserve headings, table semantics, chart context, metadata and provenance so authorised AI tools can cite the report rather than guess from it.
- ✓ **A defensible operating model:** Capture data snapshots, prompt and rule versions, human decisions, exceptions and the exact document that was distributed.



2. Context: One Report Must Now Serve Readers, Controls and Machines

Private-banking reporting is moving from a one-directional publication process to a controlled information loop. PDF remains the natural container for formal client communication because it preserves pagination, visual identity and a stable representation across platforms. It can be archived, signed, distributed and reproduced consistently, which matters when a bank must demonstrate exactly what was communicated to a client at a specific moment.

But the report is no longer consumed only page by page. It may be summarised, compared, indexed, translated, queried or embedded in a broader client dialogue. The same file can become an input to an RM briefing assistant, a compliance review, a portal chatbot or an evidence search years later. In that environment, the quality of the underlying structure becomes as important as the visual layout.

A report now has three audiences




The human reader

Needs clarity, hierarchy, explanation, premium design and a document that supports a useful conversation.



The control environment

Needs reproducibility, approvals, evidence, distribution accuracy and a defensible record of what happened.



The machine reader

Needs logical structure, unambiguous context, metadata, table relationships and explicit links between narrative and data.



THE STRATEGIC QUESTION

Are your client PDFs designed only for the person looking at the page, or also for the authorised systems that will use the information next?

Why visually perfect PDFs still fail AI

A document can look flawless on screen while remaining ambiguous, incomplete or misleading when interpreted by software. The common failure modes are consistent across institutions.

- ✗ **Ambiguous reading order:** Multi-column layouts, sidebars, headers and footnotes are extracted in the wrong sequence.
- ✗ **Tables reduced to visual geometry:** The viewer sees rows and columns, but the extraction layer does not preserve headers, units or relationships.
- ✗ **Charts without machine context:** A chart image may have no description, no data series, no scale, no period and no explanation of what matters.
- ✗ **Repeated page furniture:** Logos, page numbers and disclaimers mixed into content on every page pollute retrieval and summaries.
- ✗ **Disconnected narrative:** The paragraph describes performance, but nothing links the sentences to the relevant portfolio, benchmark, period or figures.

OCR is not the answer. Optical character recognition can recover visible characters from a scanned or flattened page. It cannot reliably reconstruct the business meaning that was lost when the document was created. OCR may identify "3.4%" and "MSCI World", but it may not know whether 3.4% is portfolio return, benchmark return, contribution, volatility or a footnote threshold. It may not distinguish the current quarter from since-inception results, may lose the hierarchy of grouped columns, and may recognise a disclaimer without understanding which section it qualifies.


The most reliable approach is to preserve semantic information while the report is being assembled. The reporting engine already knows which object is a heading, which data series



belongs to a chart, which period and currency apply, which table cells are headers, which text is a disclaimer and which client or portfolio the document concerns. Capturing that meaning at source is more reliable than inferring it later from page coordinates. The W3C makes the same point for accessibility: correct reading order is best established in the authoring process, before conversion to tagged PDF, rather than reconstructed after publication.

A USEFUL RULE

Do not ask downstream AI to reverse-engineer meaning that the reporting system already knew at generation time.




A Reporting Maestro client statement: structure and meaning are defined at generation time.



3. Write, Check, Read: Three AI Roles That Cut Effort and Risk

AI plays three distinct roles across the reporting lifecycle. Each is a separate capability with its own controls, and most programmes currently use only the first.



Write: First Drafts in Minutes, Every Sentence Grounded

The opportunity is not to let a general chatbot write the report. It is to generate a controlled first draft from approved data, rules and context. Client-report commentary is expensive because it combines analytical work, editorial judgement and review, and the same patterns recur across many portfolios: explain the main performance drivers, identify changes in allocation or risk, compare results with the benchmark, highlight material events, adapt the message to client language and reporting period. Generative AI can reduce the mechanical part of this work, but only when the inputs and the permitted language are tightly governed.

BANKABILITY CONDITION

Every generated sentence should be traceable to an approved data snapshot, a defined rule set, authorised context and a recorded human review decision.

Check: Find Costly Errors Before the Client Does

For many banks the most credible first use of AI is not generation. It is finding problems before the report reaches the client. Reports already pass through many deterministic controls: totals reconcile, mandatory sections are present, files are generated, page counts are within range. The most expensive defects often sit between those rules: a sentence



contradicts a table, the commentary names a different benchmark from the one displayed, a currency is omitted, a chart is present but not explained, the tone crosses from factual reporting into recommendation. These are comparison and context problems, which is a natural area for controlled AI assistance.

POSITIONING LINE

Before AI writes the report, let AI help prove that the report is internally consistent.

Read: Cited Answers From Every Published Report

Once released, the PDF enters multiple controlled environments: the client portal, the relationship manager's workspace, the document archive, a compliance case file or a secure enterprise search platform. AI can answer questions about a report, prepare a client-call briefing, compare the current period with the previous one or retrieve an explanation given during a past market event. These use cases depend on more than extracting text. A reliable answer must know which client and portfolio the document concerns, the period and currency, the meaning of each table, the scope of each disclaimer and the evidence behind a cited sentence. It must also respect the bank's entitlement model, because machine readability never replaces access control.

The Common Foundation: Structure at Source Instead of Reconstruction

All three roles depend on semantic structure that is generated intentionally when the report is assembled, then validated, rather than reconstructed afterwards. This is where accessibility, machine readability and reporting provenance reinforce one another. Accessibility standards such as Tagged PDF and PDF/UA provide essential building blocks, but a bank also needs business semantics and reporting evidence that accessibility standards alone do not define.



4. Solution: Automate Commentary and Keep Every Sentence Defensible

✓ 1. Govern the Inputs: Reliable Drafts From a Controlled Package

The opportunity in the Write role is not to let a general chatbot write the report, but to generate a controlled first draft from approved data, rules and context. Give the model a compact, explicit reporting package rather than an uncontrolled data dump: portfolio and benchmark figures, the reporting period, permitted comparisons, a glossary, language rules, approved context and examples of acceptable output. This supports the recurring commentary patterns, including performance commentary, table-to-insight, chart-to-narrative, mandate alignment, approved market context and controlled multilingual output in English, French, German and Italian. Classify what the AI is allowed to say: factual language restates validated information, analytical language interprets relationships the data supports, and advisory or recommendation-like language stays under stricter rules or outside generation entirely. Validate the output against the source package, present it with its evidence, and have a named reviewer accept, edit or reject it. The objective is to focus human attention on judgement and exceptions instead of repetitive drafting.

✓ 2. A Traceable Control Loop: Faster Approval, Full Accountability

Without a traceable control loop, faster drafting simply creates a faster route to an undefendable statement. The process must retain the prompt and rule version used, the data version supplied for the period, the human decision to approve or reject and its reason, the validation evidence linking sentences to source facts, the final distribution record of what was sent to whom and when, and the limits that define which topics require no generation, escalation or compliance review. The reviewer should see the draft together with its evidence rather than as an isolated paragraph, which makes approval faster and materially improves accountability.

✓ 3. AI Plus Deterministic Rules: Evidence-Backed Issue Lists

For the Check role, a robust checker uses both approaches. Deterministic rules stay authoritative for calculations, mandatory fields, thresholds and known policy constraints. AI handles the comparison and context problems that sit between those rules: a sentence that



contradicts a table, a benchmark named in the text but not displayed, a missing period or currency, a chart present but unexplained, a tone that crosses from factual reporting into recommendation, or a template that changed unexpectedly against the previous approved version. The result should be an evidence-backed issue list with severity, location, source facts and a reason, not an opaque quality score.

✓ 4. Structure at Source: AI-Readable Output Without Lock-In

For the Read role, an AI-readable PDF is one generated with enough semantic structure, context, provenance and access control for an authorised system to extract and cite its content with materially less ambiguity: navigable structure, unambiguous data carrying periods and units, grounded narrative, metadata that travels with the file, chart alternatives, citable sections and authoritative entitlements. PDF 2.0 (ISO 32000-2), PDF/UA-2 (ISO 14289-2) and Tagged PDF provide essential foundations, but accessibility conformance alone does not supply the business semantics client reporting needs. The pragmatic architecture is to generate a well-structured tagged PDF as the human document, preserve metadata and stable identifiers, and where appropriate retain a controlled machine-readable package (for example JSON or XML for tables, chart series and provenance) linked to the exact report version. Manage prompts, rules, data contracts and evidence outside any proprietary model interface, so the bank can change model or provider without redesigning the reporting process. AI-readable PDF remains a ReportingSoft working concept, not an ISO certification, and banks should continue to test against the formal standards that apply to their documents.

ARCHITECTURAL RULE

AI may propose, compare and explain. The governed reporting platform remains responsible for data preparation, workflow, approvals, document generation, distribution and the final system of record.



5. The Opportunity Map: Where AI Pays Off First in Client Reporting

The broader value lies in a portfolio of governed use cases, prioritised by credibility, data readiness and commercial pull, rather than by novelty.

RECOMMENDED FIRST WAVE

✓ **AI checker**

A credible control use case with measurable defect reduction, and usually the easiest to approve.

✓ **Controlled commentary**

Start with factual and analytical text for one report family, then expand step by step.

✓ **AI-readable PDF**

Improve document structure at source and validate downstream retrieval quality.

✓ **RM briefing**

Turn the report into a concise, cited preparation note before a client meeting.

EXPANSION PATH

✓ **"What changed?" analysis**

Compare allocations, performance, risk and narrative across periods.

✓ **Historical knowledge**

Search prior reports and explanations within entitlements.

✓ **Client Q&A**

Answer only from authorised report content, with explicit citations.

✓ **Evidence packs**

Assemble the data lineage, approvals and distribution record for audit.

PORTFOLIO PRINCIPLE

Do not rank use cases by novelty. Rank them by the quality of the available evidence, the consequence of error and the clarity of the human decision point.



6. Implementation: From Pilot to Production in Four Controlled Phases

The best pilot is narrow enough to control, but real enough to expose data, workflow and review constraints. Pilot one report family, not the entire reporting estate.

1. Scope and Boundaries

Choose one recurring report family with stable source data, meaningful narrative effort and a named business owner. Restrict the language scope to begin with, expanding only after terminology and review quality are stable. Assemble a historical validation set of previously approved reports to test extraction, checks, expected commentary and edge cases. Confirm which core systems feed the family.

2. Shadow Operation

Generate and check in parallel with the existing process, before any AI output is allowed into the client document. Define a clear fallback: if evidence is missing, confidence is inadequate or the use case leaves its permitted scope, revert to manual processing. Introduce checking before generation, so the estate's weaknesses are mapped before commentary is drafted.

3. Measurement and Go/No-Go

Quantify the pilot against metrics that matter, not against whether the text looks good. Move to production only when quality thresholds, reviewer ownership, monitoring, incident handling and the evidence package are agreed.



4.

Production and Ongoing Governance

Keep responsibilities explicit across business and reporting, data owners, risk and compliance, IT and architecture, relationship managers, and operations. Manage prompts, rules, data contracts, evaluation cases and evidence outside any proprietary model interface, so the bank can change model, deployment option or provider without redesigning the reporting process. This reduces concentration risk and allows different tasks to be routed to different models under one consistent control framework.

The metrics that matter

Cycle time

From data-ready to approved report.

Reviewer effort

Minutes and edits per commentary.

Defect rate

Material issues found before release.

Grounding

Statements linked to approved facts.

Override rate

Drafts edited, rejected or bypassed.

Consistency

Terminology and tone across portfolios.

Retrieval quality

Answers supported by correct citations.

Audit completeness

Required evidence captured automatically.



7. Success Factors: The Controls That Keep AI Reporting Trusted

The control framework should reflect the consequence of error, the autonomy of the use case and the number and type of clients affected. FINMA Guidance 08/2024 sets out a technology-neutral, proportionate approach to AI risk, covering governance, inventory and risk classification, data quality, testing and monitoring, documentation, explainability and independent review. The NIST Generative AI Profile frames governance, provenance, pre-deployment testing and incident handling as core risk-management activities. Both translate directly into client-reporting controls.

- ✓ **Ground every generated sentence:** The most common failure is unsupported generation, where a plausible explanation cites a driver that is not present in the approved data. Constrain the inputs, require sentence-level grounding, define no-generation topics and keep human approval in the loop.
- ✓ **Apply entitlements before retrieval, not after:** An assistant that retrieves another client's report or a restricted annex is an access-control failure, not a model error. Enforce permissions before the model sees the content, and treat any uploaded document text as data rather than as instructions, so hostile content cannot manipulate a downstream assistant.
- ✓ **Version everything and keep model choice open:** Model or prompt drift lets the same input produce a materially different conclusion after a change. Manage prompts, rules and evaluation cases outside the model interface, run regression tests, gate releases and keep a rollback path. This also limits third-party concentration risk, because tasks can be re-routed without redesigning the reporting process.
- ✓ **Protect the data end to end:** Confidential client information must not reach an unauthorised service or be retained outside policy. Use approved deployment, data minimisation, encryption, access controls, defined retention terms and logging as standing conditions rather than project tasks.



- ✓ **Make human oversight operational, not ceremonial:** A reviewer needs time, evidence, authority and a clear reason to intervene. Logging a click on approve is not oversight if the interface hides the source facts, if the volume makes review impossible, or if users are quietly discouraged from rejecting the draft. Design the controls so that exceptions are visible and safe refusal is easy, and capture the source package, rules, versions and reviewer action so the bank can always reproduce why a statement or flag was produced.

The controls at a glance

- 01 ✓ **Ground every generated sentence**
- 02 ✓ **Apply entitlements before retrieval, not after**
- 03 ✓ **Version everything and keep model choice open**
- 04 ✓ **Protect the data end to end**
- 05 ✓ **Make human oversight operational, not ceremonial**



8. Conclusion: Controlled, Not Autonomous, Wins

The winning model is not "AI creates the report". It is a governed reporting system in which AI accelerates selected tasks while the bank retains evidence, judgement and control. The next standard is controlled, not autonomous.

The PDF will remain a central client-communication artefact because it combines visual quality, portability and a stable record, and its role is expanding. It is becoming an input to AI-assisted work before a client meeting, a source for controlled Q&A, a target for automated assurance and a searchable piece of institutional history. Banks should design the document for both the human reader and the authorised machine reader.

OUR POSITION

We help banks produce client reports that AI can write safely, check before sending and understand after publication.



Ten questions for an AI Reporting Readiness Review


- | | | | |
|-----------|---|-----------|--|
| 01 | Can each narrative statement be linked to approved figures, context and a reporting period? | 06 | Are chart captions, descriptions and underlying data available where required? |
| 02 | Are factual, analytical and recommendation-like forms of language governed differently? | 07 | Do metadata and stable identifiers travel with every report version? |
| 03 | Does the reviewer see evidence and exceptions alongside the AI draft? | 08 | Can an authorised assistant cite the exact report section supporting an answer? |
| 04 | Can text, tables and charts be compared automatically for contradictions and omissions? | 09 | Are entitlements applied before retrieval and model processing? |
| 05 | Does the PDF preserve logical headings, reading order and table semantics? | 10 | Can the bank reproduce the data, rules, versions, approvals and final distribution record? |



About ReportingSoft

The strongest AI use cases emerge where data, templates, production controls, approval and evidence are already managed as one reporting process. The difficult part is rarely the final call to a language model. It is making the information complete, validated, contextual, personalised and distributable, then being able to explain exactly what was produced and why.

Reporting Maestro and Business Report Designer are built around those fundamentals: consolidate data from multiple sources, expose approved information to business users, create and manage reusable templates, automate mass production, apply rule-based controls, support review and secure distribution, archive the resulting documents and retain audit information. Integrations include Avaloq, Triple'A, Finnova, Olympic and T24. These capabilities do not remove the need for AI governance. They provide the control plane into which governed AI services can be inserted.



Reporting Maestro: orders, documents, controls and distribution in one governed dashboard.





The next step

Book a 30-minute AI Reporting Readiness Review with Patrick Cornaz to see where your report families stand today: on structure, controls and readiness for governed AI.

Book a 30-minute AI Reporting Readiness Review

Contact: Patrick Cornaz

Website: www.reportingsoft.ch **Email:** info@reportingsoft.ch **Phone:** +41 21 791 54 87

Selected references. FINMA Guidance 08/2024, Governance and risk management when using artificial intelligence. NIST AI 600-1, Generative Artificial Intelligence Profile. Regulation (EU) 2024/1689, Artificial Intelligence Act. ISO 32000-2:2020, Portable Document Format (PDF 2.0). ISO 14289-2:2024, PDF/UA-2. W3C, PDF Techniques and Tagged PDF accessibility guidance.

