

From Co-Pilots to Governed Digital Employees — The Next Chapter in Finance-Native Agentic AI

A Message from Our Team

Q2 2026 marks the most significant platform expansion in Scalata's history. What began as the production release of Grid AI has expanded into a comprehensive agentic finance operating system — adding AI assistant capabilities, a full policy & guardrail governance engine, a rebuilt Deep Research Agent with multilingual search, enhanced user and chat history management, real-time cost analytics, major Data Marketplace expansions (Snowflake, Databricks, Fabric, Nasdaq Data Link), Dictionary Management and model training capabilities, a new Document and PowerPoint Canvas with multi-template export, and a full multi-sheet Table AI environment with native financial formulas and Python code execution.

These capabilities have been developed in close collaboration with our enterprise clients and informed by rigorous monitoring of the regulatory landscape — including revised US Model Risk Management guidance, DORA operational resilience requirements, and the broader market shift from AI experimentation to governed, workflow-native deployment.

We are grateful to our clients, partners, and team for the extraordinary effort that has made this quarter possible. The platform that emerges from Q2 2026 is the foundation for the next chapter: a governed digital workforce for regulated financial institutions.

— The Scalata.ai Team

Product Videos & Platform Highlights

This quarter, Scalata released a series of new marketing videos showcasing how financial institutions can leverage finance-native generative AI to automate complex workflows, accelerate analysis, and deploy governed AI across the enterprise. These videos provide a visual walkthrough of key platform capabilities and demonstrate how Scalata is transforming financial operations through intelligent automation.

Featured Videos

- [From Co-Pilots to Digital Employees](#) Discover how AI is evolving beyond simple assistants into governed digital employees capable of executing complex financial workflows with enterprise-grade controls.
- [Grid AI: Transforming Financial Research & Analysis](#) See how Grid AI enables analysts to reason across dozens of documents simultaneously, accelerating due diligence, credit analysis, and investment research while maintaining full source traceability.
- [Document AI & Intelligent Report Generation](#) Explore how Document AI automates the creation of credit memos, investment reports, and regulatory documentation with customizable templates, citations, and AI-assisted authoring.
- [Table AI & Financial Modeling](#) Learn how Table AI brings AI-powered spreadsheet analysis, financial formulas, Python execution, and multi-sheet workbook capabilities together in a governed environment designed for financial professionals.

Market Traction & Client Success

Scalata.ai's annual recurring revenue (ARR) has grown to \$5M, with an active pipeline of approximately \$1.2 million in deals under active negotiation across three continents. Over 30 new enterprise opportunities have been generated across the US, EU, Japan, and Brazil, with major deployments progressing in regulated banking and credit markets.

Our SOC 2 Type II compliance certification continues to serve as a critical differentiator in enterprise procurement, enabling accelerated onboarding with tier-1 financial institutions that require independently verified security and governance controls.

Strategic Growth: Fundraising and Partnerships

- Deepened strategic partnership with the MISMO community for mortgage lending in the USA, participating in the MISMO Marketplace and Built On Program — reinforcing Scalata's role in setting agentic AI standards for the US mortgage data ecosystem.
- Progressed partnership discussions with a major global banking institution — one of the largest players in finance, AI, and total banking solutions — marking a pivotal step toward enterprise-scale deployment of Scalata's agentic credit platform inside tier-1 banking infrastructure.
- Attended a major Japan conference focused on Nasdaq-listed markets and capital market innovation, opening new commercial frontiers in Asia and establishing Scalata's presence as a global leader in compliance-first agentic AI.

Team Growth

- Welcomed a new Head of Sales in the USA, bringing deep expertise in fintech enterprise sales, coming from IBM and Bloomberg.
- Added a Junior Customer Success Specialist to strengthen client support and accelerate onboarding for our growing roster of enterprise clients.
- Added two Senior Software Engineers to accelerate platform development
- Added a new Head of Japan in Tokyo and a Sales representative
- Added New Head of Latin America in São Paulo, Brazil, and Sales Representative
- Added two sales representatives in London, UK
- Added six Advisors in the USA on the sales side

AI Assistant Features

Additional Highlights

The Q2 materials explicitly call for adding AI assistant features. While the report does not provide a single exhaustive feature list under that exact heading, the documented upgrades below capture the assistant-related enhancements that were clearly described.

Grid AI was productionized, with reasoning across 50+ documents in parallel, extending assistant usefulness for high-volume review and synthesis.

Page-level citation highlights were added to support transparent verification of outputs, helping users inspect where conclusions came from in the source material.

Python Agent entered formal Beta, enabling custom scripts and advanced data manipulation inside the platform for more sophisticated assistant-led workflows.

Remittance Report Generation and enhanced Table AI expanded the assistant's role in reconciliation, structured reporting, and complex data-processing workflows.

Admin-configured digital employee personas introduced controlled assistant behavior, with fixed scope, model, and output settings defined by administrators rather than end users.

Public web corroboration describes Grid AI as combining parallel document querying, structured outputs, shareable grids, chat-enabled cells, governance, monitoring, and agentic workflows, reinforcing the report's description of a more operational assistant experience

Q2 2026 marks the launch of Scalata's next-generation AI Assistant layer — a persistent, context-aware conversational interface embedded directly into every surface of the platform. Rather than a standalone chatbot, the AI Assistant is deeply integrated into Grid AI, Table AI, Document AI, and all data workflows, functioning as a real-time co-pilot that understands the user's current task, active documents, data context, and institutional policies.

Contextual Inline Assistance

The AI Assistant monitors the user's active workflow context and proactively surfaces suggestions, warnings, definitions, and next-step recommendations. When working inside a credit memo, the assistant can flag covenant breaches, identify missing disclosures, and suggest comparable precedent language drawn from the firm's policy library.

Multi-Turn Conversational Interface

Users can open an assistant panel alongside any Grid AI session, Document Canvas, or Table AI workbook and conduct multi-turn conversations about the data, model outputs, regulatory interpretations, or report structure. The assistant maintains full session memory and can reference prior questions, uploaded documents, and previously generated outputs within the same conversation thread.

Prompt Templates & Workflow Accelerators

Admins can configure institution-specific prompt templates and workflow accelerators. Teams working on CLO surveillance, SBA underwriting, or remittance reconciliation can launch pre-built assistant flows with one click, reducing time-to-insight by eliminating repetitive prompt construction.

Role-Aware Response Calibration

The AI Assistant adapts its response depth, vocabulary, and format based on the user's defined role within the digital employee framework. A credit analyst receives different default response depth than a portfolio manager or compliance officer, ensuring outputs are appropriately calibrated without extra configuration.

Policies & Guardrails — Managed Governance

Additional Highlights

Scalata's Q2 2026 materials place policy enforcement and workflow guardrails at the center of enterprise AI deployment. The report states that the company implemented GenAI-supported Policy & Guardrail Engines that translate organizational policies into machine-enforced controls in real time, shifting governance from guidance documents to operating controls inside production workflows.

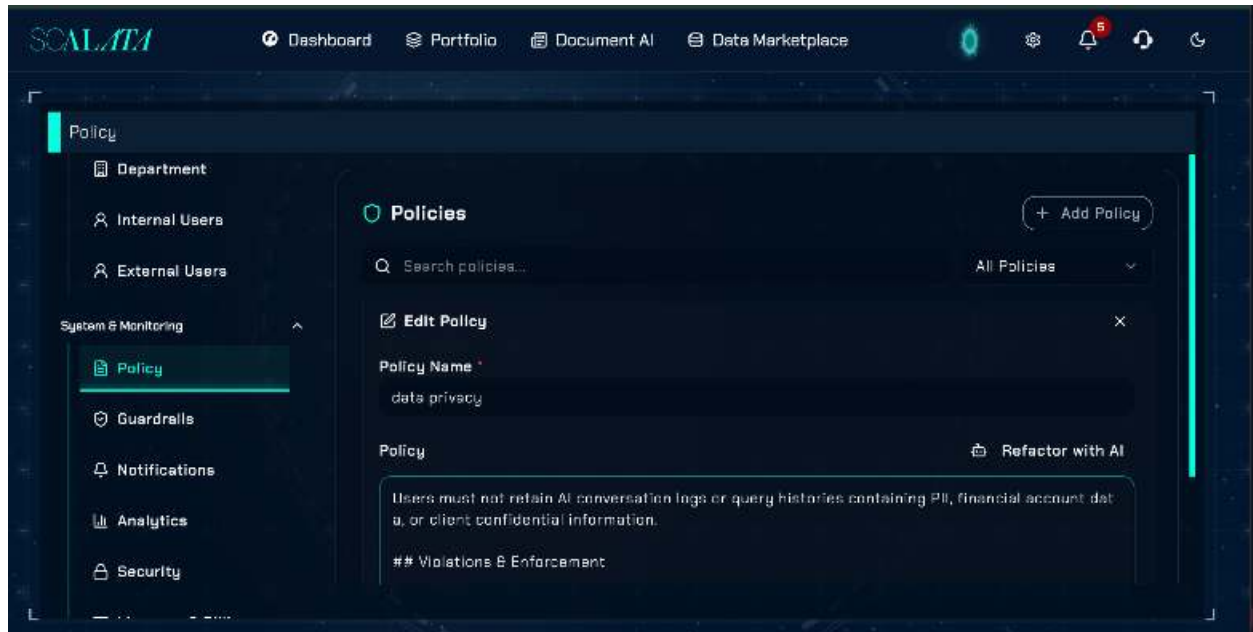
Organizations can author policies natively or import compliance PDFs, then use AI to refactor those rules into enforceable language. The report indicates that Sanitize and Block controls are applied throughout the workflow and across every workflow stage, including data ingestion, credit scoring, underwriting, risk analysis, surveillance, and investor reporting. This is framed as a way to create explainability, auditability, and production trust, rather than allowing open-ended prompt use in sensitive environments.

The governance model is supported by immutable audit logs or audit trails for every key action. Public-facing website language also reinforces this direction, highlighting Responsible AI, permissions and user controls for compliance, workflow control, enterprise license management, API control, and data monitoring as platform features as of 2026-06-25.

Scalata's Policy & Guardrail Engines have been elevated from a supporting capability to a first-class platform product in Q2 2026. Every agentic workflow, digital employee action, data extraction, report generation, and external API call is now subject to real-time policy enforcement — without interrupting the user's workflow.


Policy Authoring Studio

Institutions can author governance policies natively inside the platform using a structured policy editor, or import existing compliance PDFs and regulatory documents. The Gen AI–supported authoring engine automatically refactors narrative compliance language into machine-enforceable rule sets, validated for internal consistency and tested against synthetic edge cases before activation.



Guardrail Actions: Sanitize, Block, Redact, and Escalate

Four guardrail actions are now available across all workflow stages: Sanitize (clean outputs before delivery), Block (prevent prohibited actions from executing), Redact (mask sensitive data fields in outputs), and Escalate (route flagged items to designated human reviewers with full audit context). Each action is configurable at the policy level with institution-defined triggers, thresholds, and override approval paths.



Workflow-Stage Coverage

Guardrails are applied across the full agentic pipeline: data ingestion, document extraction, credit scoring, underwriting, risk analysis, surveillance, investor reporting, and external API outputs. This ensures no stage of an automated workflow can bypass governance controls, even in multi-step agent chains.

Immutable Audit Logs

Every guardrail trigger, policy decision, override approval, and agent action is captured in immutable, tamper-evident audit logs. Logs are exportable in standard formats for model risk management reviews, internal audit, and regulatory examination. Log entries are timestamped, user-attributed, and include the full reasoning trace of the guardrail decision.

Governance Dashboard

A new Governance Dashboard provides real-time visibility into policy coverage across all active workflows, guardrail trigger frequency by policy category, pending escalations, override rates, and trend analysis by department or digital employee. Compliance teams can monitor institutional AI governance posture without querying individual audit logs.

Deep Research Agent — Multilingual & Long-Form Reports


Additional Highlights

The Q2 report includes an instruction to add an improved deep research agent with multilingual API search and long-report generation. However, the retrieved excerpts do not provide detailed functional specifications, benchmark numbers, or a confirmed list of supported languages. Accordingly, the safest reading is that Q2 materials indicate directional expansion of deep research capabilities rather than a fully documented public specification.

What can be stated with confidence is that the broader Grid AI stack now supports parallel reasoning across 50+ documents and produces page-level citation highlights, both of which materially support research workflows by improving evidence collection, synthesis speed, and source verification. That combination is highly relevant for analysts, strategy teams, and operating teams that need to move from fragmented document review to structured, traceable outputs.

Public web material is directionally consistent with this interpretation. Scalata's Grid AI article highlights interactive document analysis, chatting with document subsets, structured insights, and inference-driven workflows, but it does not explicitly confirm multilingual research depth or long-form publishing mechanics. For that reason, multilingual support and long-report generation should be described as capabilities indicated in the Q2 report, with implementation detail still limited in the retrieved source material.

Scalata's Deep Research Agent has been substantially upgraded in Q2 2026, moving from a document-scoped summarization tool to a full autonomous research capability that can plan, execute, and synthesize multi-source financial research across languages, geographies, and data formats.



Autonomous Research Planning

The agent now generates structured research plans from a single user prompt. Given a research objective — such as analyzing covenant structures across a CLO portfolio, reviewing a target company's regulatory history, or benchmarking credit spreads across EM issuers — the agent decomposes the task into sub-questions, maps required data sources, executes parallel searches, and synthesizes a coherent long-form output without manual step-by-step orchestration.

Multilingual API Search

The research agent now integrates with multilingual search APIs covering English, Japanese, Spanish, Portuguese, French, German, and Italian. Financial research, regulatory filings, central bank publications, and market data from non-English sources are automatically retrieved, machine-translated with financial terminology preservation, and incorporated into the research synthesis. This is critical for Scalata's expanding operations in Japan, Brazil, and the EU.

Long-Form Report Generation

Research outputs can now be formatted as full long-form reports — structured documents with an executive summary, methodology section, section-by-section analysis, data tables, source citations with page-level references, and an appendix of raw retrieved excerpts. Reports are

generated in Scalata's Document Canvas and can be exported to DOCX, PDF, or added directly to a presentation canvas as slide-ready content.

Source Traceability and Citation Engine

Every claim, data point, and conclusion in a Deep Research report is linked to its source with page-level citation highlights. Users can click any citation to open the source document at the exact passage, supporting internal peer review, compliance sign-off, and audit trail requirements for research-driven investment decisions.

Chat History & User Management

Additional Highlights

The report explicitly says to add improvements on chat history management and user management, but the extracted text does not provide release-note level specifics such as search, retention controls, pinning, SSO, SCIM, or detailed admin-console functionality. As a result, this section should stay grounded in the higher-confidence evidence available.

Public sources describe interactive Grid AI usage, including chat with document subsets, as well as internal chat with users, team cooperation, deal cooperation, permissions and user controls, and enterprise license management. Together, these signals suggest that daily collaboration and controlled multi-user operation are active product priorities in the current platform.

Within the report itself, the stronger evidence comes from governance-oriented user controls. Admin-governed personas and RBAC-aligned structures imply more disciplined user management and more structured chat or assistant usage patterns. The practical implication is improved consistency across teams, better administrative control, and tighter alignment between collaboration workflows and institutional governance requirements, even though detailed feature mechanics were not explicitly present in the extracted report text.

Q2 2026 introduces a major upgrade to Scalata's conversation management and user administration layer, giving enterprise administrators full control over session history, access patterns, and user-level governance.

Persistent Session History with Smart Search

All AI conversations are now persisted with full context — including uploaded documents, generated outputs, and agent actions taken during the session. Users can search across their entire conversation history using natural language queries, filter by date range, document type, workflow category, or digital employee persona used. Relevant prior conversations can be reopened and resumed, with the full document and context state restored.

Session Summarization & Knowledge Extraction

At session close, users can generate an automatic AI-powered summary capturing key decisions, conclusions, pending actions, and flagged items from the conversation. These summaries can be saved to the user's personal knowledge base, attached to a workflow record, or exported to DOCX for inclusion in internal memos or client deliverables.

Admin User Management Console

A new centralized Admin Console provides full lifecycle management of platform users: provisioning, role assignment, department scope configuration, digital employee access controls, and deprovisioning. Admins can manage user permissions at the feature level — controlling access to Python Agent, Deep Research, Data Marketplace, API configurations, and governance override capabilities — without requiring developer-level configuration.

SSO and Directory Integration

Scalata now supports enterprise SSO integration with Okta, Microsoft Entra ID, and SAML 2.0 providers. User attributes from corporate directories — including department, title, and security clearance — are automatically mapped to Scalata RBAC roles, reducing provisioning time and eliminating manual role assignment for large institutional deployments.

Session Audit and Compliance Export

Compliance teams can export full session histories by user, department, or date range in structured formats suitable for regulatory examination. Session exports include the full conversation transcript, all documents accessed or generated, guardrail events triggered, and any override approvals issued during the session.



Platform Cost Analytics & Model Spend Management

Additional Highlights

The Q2 materials explicitly indicate additions to platform cost analytics, including model spending and per-user management. The attached report excerpt does not provide numeric dashboards, formulas, thresholds, or example savings figures, so the most accurate framing is qualitative rather than metric-heavy.

Even with those limits, the report's direction is clear: enterprises are being given greater visibility into usage and spend, enabling AI deployment to be governed not only for compliance and quality but also for operating efficiency and budget management.

Because the source material does not confirm whether these capabilities include chargeback, showback, alerts, budget controls, or optimization recommendations, those specifics should not be asserted without further product documentation.

Enterprise AI deployments require rigorous cost governance. Scalata's new Platform Cost Analytics module gives administrators full visibility into model consumption, spending trends, and per-user cost attribution — enabling informed decisions about model selection, workflow optimization, and budget allocation.

Real-Time Model Spend Dashboard

The Cost Analytics Dashboard provides real-time monitoring of token consumption and estimated spend across all platform models, broken down by model provider, model version, workflow category, department, and individual user. Spend is updated continuously and displayed in both raw token counts and dollar-equivalent cost using current provider pricing.

Per-User Spend Attribution

Every model call is attributed to the initiating user, their assigned department, and the workflow category it belongs to. Admins can view per-user spend over any date range, compare consumption across departments, and identify the highest-volume workflows or users for optimization review.

Model Selection & Budget Controls

Admins can configure model routing rules that automatically select the appropriate model tier based on task complexity, workflow type, or cost threshold. Budget limits can be set per user, per department, or per workflow category. When a budget threshold is approached, the platform issues an alert; when it is reached, the system can auto-downgrade to a lighter model tier or block additional calls pending admin approval.

Cost Allocation & Chargeback Reporting

For institutions managing AI spend across multiple business units, Scalata supports cost allocation tagging and chargeback reporting. Workflows and sessions can be tagged to a cost center, project code, or client engagement, enabling accurate internal chargeback reporting aligned to institutional finance processes.

Data Marketplace Enhancements

Additional Highlights

Q2 reporting points to broader Data Marketplace enhancements, especially more pre-built models and pre-configured APIs or connectors. The report explicitly references Snowflake, Databricks, Microsoft Fabric, and private sources such as Nasdaq Data Link, positioning these additions as ways to reduce time to deployment for regulated finance workflows.

Public website material corroborates the broader Data Marketplace positioning through references to universal data ingestion, external API integration, multimodel connectivity, analytics agents, and a Credit Data Marketplace with API and LLM. However, separate public announcement pages confirming the precise connector status or launch timing for each named integration were not surfaced in the provided synthesis, so the table below should be read as report-cited enhancements rather than independently verified connector releases.

Overall, the marketplace message is one of deployment acceleration: broader pre-configured data access, more reusable models, and closer alignment with the data environments financial institutions already use.

The Scalata Data Marketplace has expanded significantly in Q2 2026, with new pre-built data models, an expanded library of pre-configured data connectors, and support for private enterprise data links alongside public financial data APIs. The Marketplace now functions as the central data orchestration layer for all platform workflows.

New Pre-Configured Cloud Data Platform Connectors

Scalata now ships pre-configured, production-ready connectors for the major enterprise data platforms:

Snowflake — direct connector supporting SQL pushdown, schema browsing, and secure data sharing. Users can query Snowflake tables directly within Grid AI and Table AI workflows, with full credential management handled through the platform's secrets vault.

Databricks — native integration with Databricks SQL Warehouses and Delta Lake tables. Supports both interactive queries and scheduled data pulls, with Unity Catalog metadata surfaced directly in the Scalata table browser.

Microsoft Fabric — connector for Fabric Lakehouses and Warehouses, supporting both OneLake and semantic model access. Financial teams using Fabric as their enterprise data platform can surface their standardized models directly inside Scalata workflows.

Private Financial Data Connectors

In addition to cloud platform connectors, Scalata has added high-value private financial data links:

Nasdaq Data Link (Quandl) — pre-configured API access to Nasdaq's premium financial, economic, and alternative data datasets, including equity fundamentals, macro indicators, bond reference data, and alternative datasets.

Private API Framework — institutions can register proprietary internal data APIs (risk systems, loan origination platforms, fund administration systems) and surface them as first-class data sources inside Scalata workflows, with the same governance and audit controls applied to all external data connectors.

Expanded Pre-Built Model Library

The Data Marketplace now includes an expanded library of pre-built financial data models covering: CLO tranche waterfall models, SBA 7(a) underwriting scorecards, mortgage prepayment and default models, ESG scoring frameworks, credit spread benchmarking models, and multi-factor equity screening models. Each pre-built model is deployment-ready and configurable with institution-specific parameters.



Dictionary Management & Model Training

Additional Highlights

The report explicitly says to add new dictionary management and model training capabilities, but the retrieved excerpts do not provide operational detail on terminology dictionaries, retraining loops, model registries, fine-tuning methods, or human-in-the-loop review processes.

Given those source limits, the most responsible synthesis is that Q2 materials indicate expanded capabilities intended to improve domain adaptation and consistency in finance-specific outputs. In practical terms, dictionary management would support more reliable interpretation of specialized business language, while broader model-training workflows would help organizations tailor outputs to their internal logic and reporting conventions.

For customers onboarding new datasets or scaling repeatable analysis, the likely value is improved semantic consistency and better adaptation to institution-specific terminology. That said, the attached excerpt does not document the precise implementation approach, so this section should remain directional rather than technical.

Q2 2026 introduces Scalata's institution-specific Dictionary Management system and a structured model fine-tuning framework — enabling financial institutions to train platform models on their proprietary terminology, document conventions, data schemas, and institutional knowledge bases.

Financial Dictionary Studio

Institutions can build and maintain custom financial dictionaries inside the platform — defining proprietary terms, acronyms, product names, regulatory references, and data schema mappings specific to their business. The dictionary is applied across all platform surfaces: AI-generated definitions, document extraction, table column labeling, research synthesis, and report generation are all informed by the institution's custom terminology.

Model Training on Institutional Knowledge

Platform administrators can initiate supervised fine-tuning runs using curated examples from the institution's own document library — including historical credit memos, investment committee reports, underwriting decisions, and covenant compliance analyses. Fine-tuned models are deployed as institution-specific inference endpoints accessible only within the client's environment, ensuring proprietary training data never leaves the institution's data boundary.

Domain Adaptation for Specialist Finance Verticals

Pre-built domain adaptation packages are available for structured credit, CLO analytics, commercial real estate, private credit, insurance, and asset management. Each package includes a curated base dictionary, document taxonomy, and fine-tuning dataset drawn from publicly available domain documents, enabling faster time-to-accuracy for specialist workflows without requiring institutions to supply initial training data.

Enhanced Document AI — Canvas, PowerPoint & Multi-Template Export

Additional Highlights

Document AI is one of the clearer areas of visible expansion in Q2 2026. The report explicitly calls for enhanced document AI with a new canvas to build PowerPoint presentations with customized canvas layouts for each slide, along with a new document canvas to export in different templates on top of the widget Document AI-managed report flow.

Later report language reinforces this by describing “Canvas / Document & PowerPoint Creation” as a new canvas-style workflow for rich document generation and presentation output, including managed-agent reports, Excel workbooks, and enterprise-ready deliverables. In parallel, Grid AI expanded support for DOCX and PPTX alongside spreadsheets, enabling finance-grade documents and presentation-style outputs with citations, traceability, and governed collaboration.

Public website corroboration aligns with that direction through references to automated report generation, customizable document templates, comprehensive memorandums, enhanced reporting efficiency, and export functionality. Taken together, the Q2 message is that Document AI is evolving from extraction and summarization into governed content production across reports, presentations, and reusable executive-ready templates.

Document AI in Q2 2026 has evolved from a managed report generation engine into a full authoring and presentation platform. The new Document Canvas and PowerPoint Canvas capabilities bring professional-grade publishing workflows into the Scalata platform — built on top of the existing AI-driven document analysis and generation infrastructure.

Document Canvas — Rich Authoring Environment

The new Document Canvas provides a full-featured rich text authoring environment embedded within the platform. Users can draft, edit, and structure long-form documents — credit memos, research reports, investor letters, compliance write-ups — with real-time AI assistance at every stage. The canvas supports: section-by-section AI generation, inline editing with tracked changes, live data pull from Table AI and Grid AI outputs, and a citation panel for source management.

Multi-Template Export Engine

Documents created in the canvas can be exported in multiple professional templates without reformatting. Available export templates include: formal credit memo, investment committee presentation, investor letter, regulatory submission, internal risk report, and client-facing summary. Templates are pre-styled to institutional brand guidelines and exportable as DOCX, PDF, or HTML.

PowerPoint Presentation Canvas

A new dedicated PowerPoint Canvas enables users to build full presentation decks from AI-generated content, data tables, and charts — entirely within the Scalata platform. The canvas provides slide-by-slide editing with:

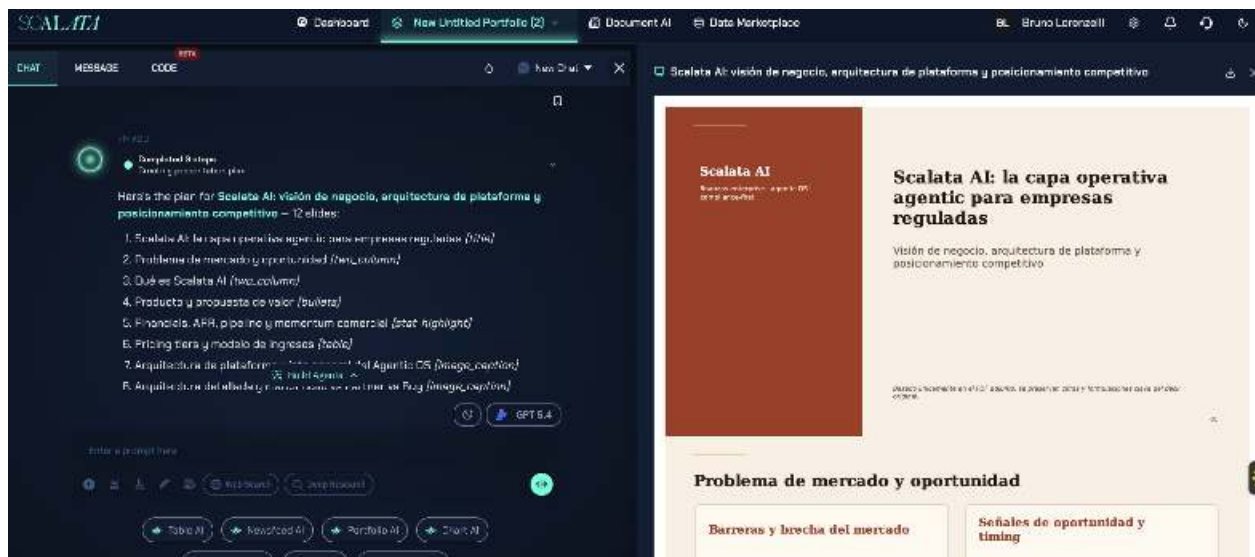
- Per-slide layout customization: title slides, data-heavy layouts, text-and-chart split, full-bleed image, and summary tables

- AI slide builder: generate slide content from a prompt or pull structured data from Grid AI sessions

- Custom slide canvas: drag-and-drop component placement, font and colour overrides, logo and brand asset insertion

- Slide speaker notes generation: AI can generate speaker notes for each slide based on the slide content

- Master template management: admins define institutional slide masters applied globally across all presentation exports



AI-Managed Report Widget — Enhanced

The existing managed report widget has been enhanced with new output modes: structured JSON for downstream API consumption, markdown for knowledge base ingestion, and a new side-by-side comparison mode for generating parallel analyses across multiple documents or scenarios. Report generation now supports conditional formatting rules — tables and data blocks in managed reports automatically apply colour scales, threshold highlights, and data bars consistent with the output template.



Multi-Sheet Excel Table AI & Financial Model Canvas

Additional Highlights

The report explicitly says to add new multisheet Excel Table AI and canvas-style financial modeling tools with formulas and Python code. Importantly, the source spelling includes “Canva/canva,” but the surrounding context supports a canvas-style interface rather than a verified integration with the third-party Canva platform, so this should be described carefully.

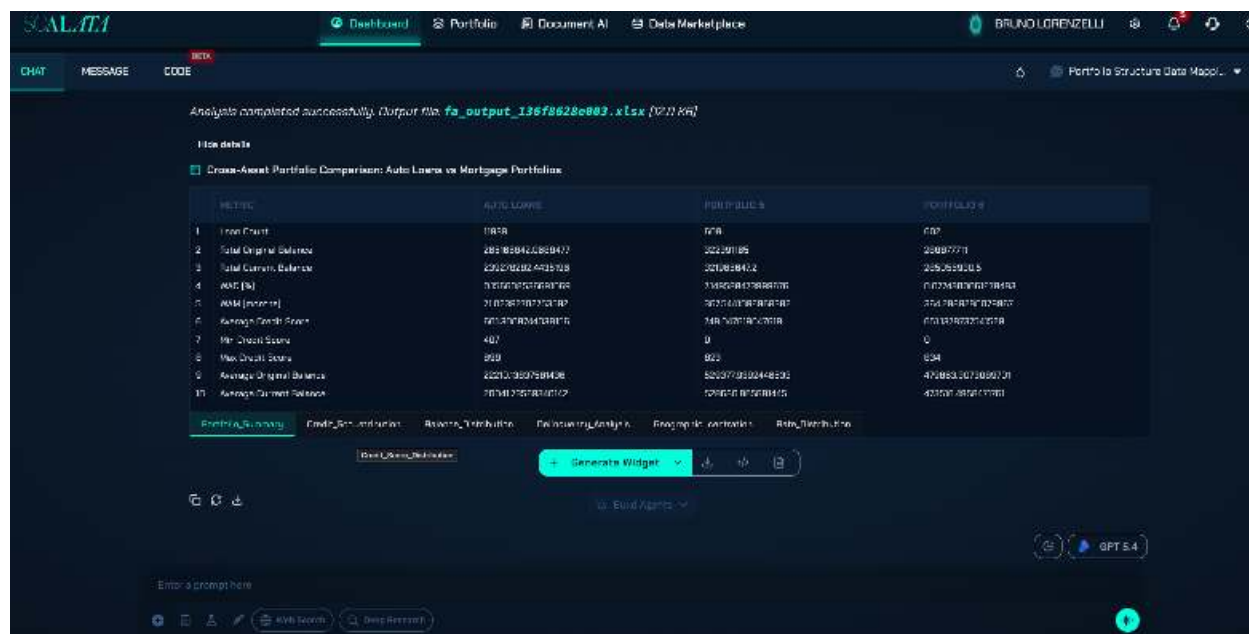
The strongest supporting evidence is concrete. Grid AI now supports multi-workbook Excel models; Python Agent is in formal Beta for custom scripts and advanced data manipulation; enhanced Table AI streamlines reconciliation and structured reporting; and canvas-style workflows extend into Excel workbooks and enterprise-ready deliverables.

Public website language further supports the modeling story through references to automated modeling from data import, chat-to-calculate workflows, quantitative finance capabilities, charts and analytics, and exporting data. For finance and operating teams, the practical result is a more connected path from data ingestion to model creation, scenario analysis, and final presentation packaging within a governed platform environment [scalata.ai/platform-features, fetched 2026-06-25].

Table AI has undergone its most significant upgrade in Q2 2026 — expanding from single-table analysis into a full multi-sheet Excel workbook environment with native financial formula support, Python code execution, and a new Financial Model Canvas for building, auditing, and sharing institutional financial models.

Multi-Sheet Workbook Support

Table AI now supports full multi-sheet Excel workbooks with cross-sheet formula references, named ranges, and structured table links. Users can upload or generate workbooks with up to 20 sheets, define cross-sheet dependencies, and run AI analysis across the entire workbook rather than individual sheets. Sheet-level metadata and data lineage are tracked and displayed in the workbook sidebar.



Native Financial Formula Engine

A financial formula engine has been embedded directly into Table AI, providing native support for Excel-compatible formulas including: NPV, IRR, XIRR, XNPV, PV, FV, PMT, MIRR, RATE,

NPER, DISC, PRICE, YIELD, DURATION, and full financial calendar day count conventions (ACT/360, ACT/365, 30/360). Formula results update in real time as input parameters change, and formula dependencies are displayed in an auditable formula tree.

Python Code Execution in Table Cells

Building on the Python Agent Beta, Table AI now allows users to embed Python expressions directly in table cells or attach full Python scripts to sheet-level transformation layers. This enables advanced calculations — Monte Carlo simulations, factor model computations, custom prepayment curves, or stress test routines — to be executed natively inside the workbook without leaving the Scalata environment. Python cells display their code, inputs, and outputs in a collapsible audit panel.

Financial Model Canvas

The new Financial Model Canvas provides a structured template environment for building institutional-grade financial models:

- Pre-built model templates: LBO model, DCF valuation, CLO waterfall, loan amortization schedule, securitization cash flow model, and credit spread sensitivity table

- Assumption driver panel: centralized input panel for model assumptions with scenario toggle (base, upside, downside, stress)

- Scenario comparison view: run multiple scenarios in parallel and display variance analysis in a side-by-side table

- Audit trail for assumptions: every assumption change is logged with user attribution and timestamp

- Model sharing and collaboration: share live model links with view-only or edit access, with version history and rollback

Export & Integration

Financial models built in the Table AI Canvas can be exported as fully formatted Excel workbooks with preserved formulas, data validation rules, and conditional formatting. Models can also be published to the Data Marketplace as reusable institutional templates, or connected to Document AI for automatic injection of model outputs into credit memos and investor reports.

Platform Evolution — Technical Achievements Q2 2026

Q2 2026 marks the maturation of the Scalata agentic platform from Beta infrastructure to a production-grade, enterprise-ready orchestration environment. Across Grid AI, the Python Agent, governance infrastructure, and digital employee framework this quarter have delivered the foundational capabilities required for tier-1 financial institution deployment.

Grid AI — Production Release

Grid AI has completed its transition from the Beta reasoning layer to full production release. The system can now reason across 50+ documents in parallel, with native support for multi-workbook Excel models and the full spectrum of business formats (DOCX, PPTX, TXT, PDF). Page-level citation highlights make every output verifiable at a glance.

Python Agent — Formal Beta


The Python Agent is now in formal Beta, enabling institution-configured custom scripts and advanced data manipulation workflows to run natively within the Scalata environment. The agent supports library management, scheduled execution, parameter injection from platform context, and full output capture for downstream workflow use.

Remittance Report Generation

Automated Remittance Report Generation automates reconciliation, data extraction, and structured financial reporting from both structured and unstructured source documents. The workflow supports multi-format remittance inputs, cross-document variance analysis, and direct output to Excel workbooks or managed report templates.

Digital Employee Framework

AI personas are now fully admin-configured with defined department scope, model assignment, output format, style prompts, and locked generation settings. This supports reproducible, compliant outputs and materially reduces shadow-AI risk in regulated environments.



Events, Media & Thought Leadership

This quarter, Scalata.ai continued expanding its thought-leadership presence across financial AI, agentic workflows, and regulated deployment. Through new blog content, marketing videos, and upcoming in-person events, Scalata is continuing to position itself at the center of the conversation around governed AI adoption in finance.

Media & Thought Leadership

In Q2 2026, Scalata continued publishing thought leadership focused on the future of financial work, digital employees, workflow automation, and agentic AI infrastructure. These pieces explored how financial institutions can move beyond simple AI chat interfaces toward systems that actively support research, reporting, modeling, and operational execution.

Recent themes included:

The Rise of the Digital Employee

Explores how AI agents are beginning to function as operational participants within financial institutions, supporting tasks, workflows, and decision-making across teams.

Workflow Intelligence Over AI Hype

Examines why the future of financial AI depends not only on model performance, but on the ability to connect data, tasks, controls, and human oversight into one governed operating layer.

Agentic AI for Financial Institutions

Highlights Scalata's approach to building AI systems that can ingest complex financial data, automate research and reporting, and support institution-specific workflows while maintaining auditability and governance.

Scalata also released new marketing and product videos designed to communicate its platform vision more clearly across LinkedIn and other digital channels. These videos focused on Scalata's role in helping financial institutions automate data-heavy workflows, accelerate analysis, and deploy agentic AI in a controlled, compliance-ready environment.

Industry Engagement & Events

In Q2 2026, Scalata is continuing to strengthen its global presence through private events and industry engagements focused on the practical deployment of AI in financial markets.

Scalata.ai Hosted Event — New York

Scalata will host an invite-only event bringing together leaders across finance, AI, and credit markets to discuss the future of agentic AI in regulated financial institutions. The event will focus on real-world use cases, including financial data automation, research workflows, credit intelligence, and governed AI deployment.

The Digital Workforce for Finance — Tokyo, Japan

Scalata will also host an invite-only event in Tokyo focused on GenAI, autonomous agents, and algorithmic trading. The session will introduce Scalata's finance-native AI platform and explore use cases across credit, asset management, trading agents, vibe coding for finance, and regulated in-country AI deployment.

The Tokyo program will include a main workshop, live demonstrations, segment-focused discussions, and networking with financial institutions. Key topics include AI trading agents, credit and portfolio AI, natural-language production code for finance, and Japan-specific requirements around RBAC, audit trails, and data sovereignty.

Navigating the Evolving Landscape

Q2 2026 Market & Regulatory Intelligence Brief

Over the past quarter, the financial technology and AI landscape has continued shifting from experimentation to operational deployment. Across banking, private credit, capital markets, and asset management, institutions are focusing less on isolated AI tools and more on infrastructure that can support workflow automation, governance, auditability, and risk control.

At the same time, regulators and market participants are paying closer attention to the systems behind AI adoption: cloud providers, third-party vendors, cybersecurity, model risk, data quality, and operational resilience. The market signal is clear: financial institutions want AI that can do more than generate answers. They need systems that can operate safely inside regulated workflows.

1. AI in Financial Services Is Moving From Experimentation to Enterprise Deployment

The 2026 Global AI in Financial Services Report highlights continued momentum around AI adoption across financial institutions, AI vendors, regulators, and consumers. The broader market is moving toward responsible AI adoption, with a stronger focus on regulatory clarity, deployment resilience, talent, and governance.

Key Market Themes

AI is becoming part of core financial operations, not just innovation labs.
Financial institutions are prioritizing responsible deployment.
Governance, risk management, and explainability are becoming adoption requirements.
AI vendors are being evaluated on production readiness, not just model performance.

Market Signal

The financial services industry is entering a more serious phase of AI adoption, where the winners will be platforms that combine automation with governance, control, and institutional trust.

2. Model Risk Management Is Being Reframed for Modern AI Systems

In April 2026, U.S. banking regulators issued revised model risk management guidance, replacing older guidance with a more risk-based and tailored approach. The OCC described the revised guidance as focused on model development, use, validation, monitoring, governance, controls, and third-party model considerations.

Key Regulatory Themes

Model governance must be proportionate to the size and complexity of the institution.
Validation and monitoring remain central to responsible model use.
Third-party models and vendor-provided AI systems require stronger oversight.
AI must fit into existing audit, risk, and compliance structures.

Market Signal

Banks are not looking for separate “AI layers” that sit outside governance. They need AI systems that can be documented, validated, monitored, and reviewed within existing model risk frameworks.

3. Operational Resilience and Third-Party Technology Risk Are Becoming Board-Level Issues

DORA continues to reshape how EU financial institutions approach ICT risk, cyber resilience, third-party oversight, and operational continuity. EIOPA describes DORA as establishing an EU-wide oversight framework for critical ICT third-party providers, designed to reduce systemic and concentration risk from reliance on a limited number of technology providers.

Key Market Themes

Financial institutions are increasingly dependent on cloud, data, and AI vendors.
Regulators are focused on concentration risk in critical technology infrastructure.
Third-party risk management is becoming central to AI and digital transformation.
Operational resilience is no longer just an IT concern; it is a business continuity issue.

Market Signal

As AI becomes embedded into financial workflows, institutions will need platforms that can prove resilience, vendor transparency, data control, and operational continuity.

4. Private Credit, AI Infrastructure, and Financial Stability Are Becoming Connected

The European Central Bank’s May 2026 Financial Stability Review raised concerns around private credit exposure and the potential disruptive impact of AI on bank business models. The review also noted that disappointment or volatility around AI-related investments could create pressure in certain market segments, especially where private markets have exposure to AI infrastructure.

Key Market Themes

Private credit remains under regulatory and investor scrutiny.
AI infrastructure financing is becoming a larger part of private markets.
Regulators are watching opacity, leverage, and interconnectedness.
Financial institutions need better visibility into portfolio-level risk.

Market Signal

This creates a growing need for AI systems that can help institutions monitor credit exposure, analyze private-market data, and surface risk across fragmented documents, loan tapes, servicer reports, and portfolio materials.

5. Enterprise AI Is Shifting From Chatbots to Agents That Act

The enterprise AI market is rapidly moving beyond conversational interfaces. Microsoft announced that computer-using agents in Copilot Studio became generally available in May 2026, allowing agents to interact with websites and desktop applications through screen-based actions when APIs are not available. Anthropic has also continued advancing Claude's agentic capabilities, emphasizing stronger performance across coding, multi-step tasks, agents, and vision.

Key Market Themes

AI agents are moving from answering questions to completing tasks.

Enterprise platforms are building agents that can operate across applications.

Natural-language workflows are becoming more actionable.

Permissions, audit trails, and human oversight become more important as agents gain autonomy.

Market Signal

The next phase of enterprise AI will be defined by execution. For regulated financial institutions, the question is not only whether an agent can act, but whether every action can be controlled, reviewed, validated, and audited.

6. Financial Institutions Are Investing in AI Through Strategic Vendor and Cloud Partnerships

Large financial and investment institutions are increasingly forming AI partnerships with major cloud and AI providers. For example, EQT partnered with Google Cloud to roll out AI tools across more than 300 portfolio companies, including access to Google's Gemini Enterprise Agent platform and cybersecurity services. These types of partnerships show how AI adoption is becoming an enterprise transformation strategy rather than a single software purchase.

Key Market Themes

AI adoption is being pushed across entire portfolios and operating companies.

Cloud providers are becoming central partners in AI deployment.

Enterprise AI is increasingly tied to cybersecurity, data infrastructure, and workflow modernization.

Financial institutions want scalable AI systems that can be deployed across teams and business units.

Market Signal

The market is moving toward platform-level AI adoption. Financial firms are looking for systems that can connect data, workflows, governance, and reporting across the organization.

7. What This Means for Scalata.ai

Across Q2 2026, the market is pointing toward the same conclusion: financial institutions need AI systems that are not only intelligent but also operationally reliable, governed, and workflow-native.

The biggest opportunities are emerging around:

- AI governance and model risk management
- Private credit and portfolio monitoring
- Financial data automation
- Third-party and cloud risk oversight
- Workflow execution and validation
- Operational resilience
- Audit-ready AI deployment

While general-purpose AI platforms are advancing quickly, financial institutions need systems designed for the complexity of regulated markets. That means AI that can ingest fragmented financial data, support credit and portfolio workflows, generate research and reporting, validate outputs, and preserve governance from the start.

For Scalata.ai, this reinforces the core market opportunity: building finance-native agentic AI that helps institutions move from experimentation to execution, while maintaining the auditability, control, and compliance required in regulated financial environments.

Looking Ahead: Q3 and Beyond

- As we move into the second half of 2026, Scalata's priorities are focused on deepening enterprise deployments, advancing platform capabilities, and continuing to expand our global footprint:
- Scalata is working on its flagship products that will determine the full-scale global launch and provide a great momentum for the company. It is launching its Vibe coding and full container APPs deployment, something unique in the generative AI space.
- Scalata AI is also launching its full Graph technology and studio design to deploy generative AI workflows more as a mapping and flow management tool
- Scale pilot deployments with enterprise banking and credit clients — converting active pipeline into signed agreements across the US, EU, Japan, and Brazil
- Deepen commercial presence in the UK, Italy, Brazil, and Japan ahead of major regional events

- Expand the Data Marketplace with additional private data connectors and pre-built vertical models
- Advance the PowerPoint Canvas and Document Canvas to full production with institutional template management

Experience Scalata.ai

Ready to see how Scalata.ai can transform your credit analysis workflow? Our platform offers:

- Cutting-edge AI analysis capabilities
- Seamless integration with existing systems
- Industry-leading security and compliance
- Dedicated support team

Book a demo with our team today to experience the future of credit analysis.

<https://scalata.ai/contact-us/>

Scalata.ai is committed to maintaining the highest standards of data security and regulatory compliance, including active pursuit of SOC2 certification and alignment with evolving AI regulations. Our platform is designed to meet the stringent requirements of financial institutions while pushing the boundaries of what's possible in credit analysis.

For more information about our products or to schedule a meeting, please contact our team at info@scalata.ai

Scalata.ai, 295 Madison Avenue, New York, NY 10017

[Unsubscribe](#) [Manage preferences](#)