

Qlik Customer Day 2026

Your questions, answered

Qlik, Snowflake and MCP - questions from Vilnius (6 October) and Riga (8 October)

Prepared by Māris Svilāns, The Infotrust · 12 October 2026

Contents

1. The big picture
 2. Qlik AI - how it works
 3. On-premise, MCP and the roadmap
 4. Trust, security and correctness
 5. Platforms and architecture
 6. Licensing and support
 7. Questions on the customer stories
- Before publishing
- Sources

Thank you for the questions sent before and during Qlik Customer Day in Vilnius (6 October) and Riga (8 October). They are grouped by theme below

1. The big picture

What was the main message of the day?

The closing session summed it up in three sentences:

- **The model is the commodity. Your context is the capital.** Running a capable AI model became about 280 times cheaper between 2022 and 2024, and the score gap between the first and the tenth model is 5.4% (Stanford AI Index 2025). Every company can use the same models. What differs is the business context: agreed definitions, joined data and people who can tell when a number is right. Much of that already sits in existing Qlik apps.
- **Let value decide which decision comes first** (Ramunas Mazeliauskas and Raivis Mackevics). Start where the value is biggest, prove it, then extend what pays off.
- **You can keep moving forward without moving everything** (Simona Klebonaite and Justinas Petkevicius). AI can be tested on Qlik Cloud or on Qlik Sense running on your own servers.

What does an organisation need to get value from AI? Is deploying the platform enough?

No. The welcome session described three things companies have built around their dashboards over the years, and AI depends on all three:

1. **Context:** agreed definitions (what "margin" means), systems joined together, known data gaps fixed.
2. **People:** those who know both the business and the data, and can tell when a number is right and when it only looks right.
3. **Habits:** people already explore data on their own and expect fast answers. Asking an AI assistant continues that habit.

Two further conditions:

- **Clean up the app landscape first.** Apps that are out of date, duplicated, undocumented or unused give AI the wrong context. An answer that is confidently wrong, or shows numbers someone should not see, costs more than no answer.
- **Start from a value case.** The morning keynote used a value framework: understand the context, find where value is won or lost, build a value case, then build the capability, change how decisions are made, and measure the result.

On why context matters: in a 2026 test (Cube, the research paper arXiv 2604.25149), three AI models answered 100 business questions. Adding 4 KB of business definitions raised accuracy from 46-51% to 68-69%, while the choice of model made little difference.

How will data analytics look in 3-5 years, and what should we watch out for when developing a new ERP?

What we expect:

- More questions asked in plain language, and agents that act on the answer. The AI session showed this with Qlik's Predict and Automate agents: train a model on late payments, score open invoices, send the high-risk ones to collections.
- Business definitions kept in one semantic layer that BI tools and AI agents both read. Qlik has announced its own semantic model as coming soon.

- Shared data products on one data platform, instead of each tool preparing its own copy.
- AI models replaced every few months. The architecture should allow switching models without rebuilding.

For a new ERP:

1. Decide early how data will leave the ERP (APIs, change data capture, Synapse Link). Some SaaS ERPs restrict direct database access. More and more often it will be open format – Apache Iceberg.
2. Keep the history from the old system accessible and mapped to the new structures, so trends do not break at go-live.
3. Assign owners to master data (customers, products, chart of accounts) before go-live.
4. Keep change logs and event logs switched on. Process mining and AI agents rely on them later.
5. Keep management reporting outside the ERP, in an analytics layer where data from several systems can be combined.

2. Qlik AI - how it works

Does Qlik develop its own AI models, or rely on third parties?

Both, for different purposes:

- Qlik Answers and the agents use third-party foundation models through Amazon Bedrock.
- Qlik Predict trains machine-learning models on your own data, inside your tenant. These models are not shared.
- In 2024 Qlik acquired Kyndi's natural-language technology to improve answer quality.

What runs behind Qlik Answers? Is it an API to popular LLMs?

Qlik Answers uses models on Amazon Bedrock, including Anthropic's Claude. Requests go through Qlik's own gateway, which selects a model for each task. The model provider has no access to customer requests or data.

The model does not calculate the figures. It interprets the question, the Qlik engine calculates the KPI with the app's definitions and the user's access rights, and the model receives the result rather than the raw rows. As the AI session put it: the model asks, the Qlik engine calculates.

How can these models or agents be fine-tuned?

Qlik Answer quality is improved through the apps themselves:

- descriptions on master dimensions and measures
- a business glossary and business definitions
- clear field names and date fields
- Qlik's agent skill `qlik-ai-readiness-optimizer`, which checks an app for these points.

To use a model of your own choice, there are two routes: your own AI assistant connected to Qlik over MCP or analytic connectors (OpenAI, Azure OpenAI, Amazon Bedrock and others) in load scripts and chart expressions.

Are Qlik's AI capabilities similar to competitors', or is there something specific?

Most BI vendors now offer a chat interface. What is specific to Qlik:

1. Calculations are done by the associative engine, so an AI answer matches the dashboard and respects section access.
2. Existing master items, scripts and KPIs serve as AI context without being rebuilt.
3. Qlik connects to external AI assistants (Claude, ChatGPT, Copilot and others) and to data platforms (Snowflake, Databricks and others). The keynote cited 60+ MCP endpoints and 700+ connectors.
4. Data products carry a Trust Score, lineage and data quality rules.
5. The agents cover the step from prediction to action (Predict and Automate agents).

Why choose Qlik over Power BI? What makes Qlik better, especially with AI?

Power BI is a good choice for organisations standardised on Microsoft. Qlik fits better when:

- users explore data across many dimensions; the associative model shows both related and excluded values
- data comes from many systems that have to be combined
- AI answers have to match the dashboard numbers and respect access rules
- the organisation wants a free choice of cloud, data platform and AI assistant
- many users work at the same time; in a [BARC benchmark published by Qlik in October 2025](#), Qlik Cloud Analytics responded about 3 times faster than Power BI at 50 concurrent users (engine performance, not AI)

For a company with years of Qlik development, reusing the existing logic for AI usually costs less than rebuilding it in another tool.

How fast is Qlik Answers now? It used to be slow.

Since June 2026 Qlik Answers has a Fast mode for simple questions, answered from master measures. Qlik's documentation gives 10-15 seconds as typical. For "why" questions, Think more switches to step-by-step reasoning, which takes longer. Answers can also draw on documents and show the sources. A Fast mode question that fails does not count against the quota.

Do users of published Qlik apps have access to a contextual AI chat?

- **Qlik Cloud:** yes. Qlik Answers opens from the navigation bar with the current app selected, answers from app data and documents, and can create charts and sheets (saved as "AI Generated" sheets). It can also be embedded in other portals.
- **Qlik Sense on your own servers:** Stretch Qconnect App+ adds a chat panel in the hub, in each app and in the script editor. Requests run through the [Stretch Qconnect MCP server](#) with the user's own Qlik rights.

Who is the intended audience for chatting with data?

Mainly people who do not build dashboards: managers who need one figure, staff in the field, and occasional users. Until now, the value of Qlik reached the people who open dashboards. With AI it can reach anyone with a question, provided the context is clear and access is governed. Analysts get fewer ad hoc data requests, and developers use the same tools to build apps and master items.

Do you plan to build "vibe check" reports, i.e. on-demand reports, into the platform?

Qlik Answers creates charts and sheets on request. Over MCP, an assistant such as Claude can create sheets, charts and whole apps. On Qlik Sense on-premise, Stretch Qconnect builds and styles

sheets from a prompt; in the demo a what-if question was answered in the chat and then added to the dashboard as a slider. Qlik lists an Analytics Agent for building apps as coming soon.

On-demand views are safe when they use certified master measures. A new formula created on the fly should be checked before others rely on it.

Is there a vision for improving Qlik development, such as data models, with AI?

Yes:

- Qlik Answers and the Qlik MCP server can create apps, edit load scripts, reload, and add master items, sheets and charts. In the AI session an Excel file became a Qlik app in about five minutes.
- Qlik Talend Cloud pipelines can be written in YAML or generated from a prompt (available since June 2026). Qlik has announced agentic data engineering as coming soon.
- Qlik publishes agent skills for Claude and other assistants.
- In migration projects, AI-assisted conversion of Qlik script logic into tested SQL models cuts the effort per app from weeks to days.

The AI session listed three checks that remain with people: a developer reviews the load script against best practice, an analyst validates every KPI before people rely on it, and the result follows company design standards.

There was little about data lineage. Is AI helping with lineage?

Yes, in three ways:

1. The [Qlik Cloud catalog shows field-level and dataset-level lineage and impact analysis](#), with owners: "What happens if we change this field?" Qlik Talend Cloud adds column-level lineage for pipeline projects.
2. Over MCP, an assistant can call Qlik's lineage tool and answer "where does this number come from?" in plain language.
3. In migrations, AI reads old Qlik scripts and reconstructs how data flows. Often this is the first time the logic is documented.

3. On-premise, MCP and the roadmap

When, if ever, will AI products be available for Qlik Sense on-premise?

Qlik's own generative AI (Qlik Answers, the agents and the Qlik MCP server) is available in Qlik Cloud already.

Stretch Qconnect MCP server runs next to Qlik Sense on-premise and connects it to Claude, Azure OpenAI or a model hosted in your own data centre. It was demonstrated in the AI session.

What changes do you expect over the next three years in the models that work with the data? Is the Qlik roadmap only for cloud?

Models will keep changing every few months, and the gap between them is small. The context the models work with (definitions, KPIs, access rules) is what lasts.

Whichever route you take, keep definitions and access rules in one semantic layer and connect AI through open standards, so that a model can be replaced without rebuilding. Gartner expects that by 2028, 60% of agentic analytics projects relying solely on MCP connections will fail without a consistent semantic layer.

Are MCP connectors public or available for purchase? Can IT set them up locally?

- **Qlik Cloud:** the Qlik MCP server is included in Qlik Cloud subscriptions. A tenant admin enables it and grants the permission to users. It works with Claude, ChatGPT and other MCP clients. MCP usage counts against the same monthly AI question quota as Qlik Answers.
- **Qlik Sense on-premise:** Stretch Qconnect is a licensed partner product that your IT installs on your own infrastructure (Linux with Docker, about 1-2 hours). Open-source community servers also exist; check their security model and support before using them with company data.
- **Snowflake:** Snowflake offers a managed MCP server, and a Qlik MCP Server native app is available in Snowflake.

What does MCP for on-premise cost, and what do we need to test it?

MCP on-premise is licensed per Qlik environment, the price depends on which functions are enabled, so it is quoted per environment.

For a test you need:

- a trial licence, arranged through us
- a Linux server with Docker, next to or near Qlik Sense
- company sign-in (SSO, for example Entra ID)
- an AI client: Claude with custom connectors, Azure OpenAI or a self-hosted model. With a cloud AI client, the MCP endpoint must be reachable from that service, while Qlik itself stays behind the firewall.

Would Qconnect AI with the side-panel chat also help people who only look at analytics?

Yes. The App+ chat panel sits in the Qlik hub and in each app, so users ask questions where they already work. Access can be limited to reading and asking. Answers come from the Qlik engine with the user's own rights. In the demo, "Why is Poland behind target?" returned ranked drivers, which were highlighted on the user's sheet, and a what-if question was answered in the chat without changing the app.

Are there built-in Qlik tools to manage MCP connections, such as access levels and audit?

Qlik Cloud:

- users need the Qlik MCP permission in their role, and a tenant admin approves each new AI client
- every tool call runs as the signed-in user, so spaces, roles and section access apply
- every tool call is logged as an audit event (tool, user, client, app, response time, errors)
- rate limits protect the tenant, for example 300 tool calls per user per minute

Stretch Qconnect AI (on-premise):

- there is no shared service account; each request runs as the authenticated Qlik user through SSO, and stops if the identity cannot be confirmed
- an admin console sets which tools each group may use and which actions need approval
- load-script changes are version-controlled and can be rolled back

Is there a published Claude skill with best practices for building Qlik applications?

Qlik publishes official agent skills in its open-source repository (qlik-oss/agentik-skills on GitHub), including `qlik-ai-readiness-optimizer` (field naming, master item descriptions, date fields and synonyms that help AI use an app) and `qlik-automation-builder`. Community skills for load scripts also exist.

4. Trust, security and correctness

What prevents data leakage, for example our data appearing elsewhere because a model was trained on it?

- **Qlik Cloud AI:** Qlik states that customer data is not used to train or fine-tune models, and that this is covered in its data processing agreement with AWS. Data is held in memory for a few seconds in the processing region and not written to disk. EU tenants stay within the EU, and the newer agentic features need an admin opt-in before any cross-region processing. Qlik is certified to ISO 42001, the AI management standard.
- **Snowflake Cortex AI:** models run **inside Snowflake's security perimeter**, and Snowflake does not use customer data to train models. Admins can restrict AI processing to EU regions or switch cross-region processing off.
- **Claude over MCP:** on Anthropic's commercial plans (Claude Team, Enterprise and the API), inputs and outputs are not used for training by default. On personal plans it depends on the user's privacy setting, so company data should go through a company plan.

In practice: use company-approved AI accounts with SSO, and give each user only the permissions their task needs.

When Claude builds an app through Qlik MCP, what data reaches the model: only metadata, or company and customer data too?

Both. The model sees what the tools return:

- **Metadata:** app, sheet, field and master item names, descriptions and the load script.
- **Results:** aggregated KPI values and chart data, and individual field values or sample rows when the assistant searches or samples data.

What limits the exposure:

1. The assistant acts as the signed-in user and sees only what that user's section access allows.
2. Admins decide who may use MCP and which tools are allowed. On-premise, Stretch Qconnect sets tool rights per group.
3. Build work can be done in a development app with test or masked data, and the result promoted afterwards.
4. For the most sensitive data, the model can run inside your own perimeter: a self-hosted model, or Snowflake Cortex.

How do we verify that an AI-generated Qlik app is correct and based only on the source data?

The numbers themselves come from the Qlik engine over the loaded data, so the risk lies in a wrong formula, join or filter rather than in invented figures. Checks to run before anyone relies on the app:

1. **Reconcile with the source:** totals and row counts per key dimension against the source system or an existing trusted report.
2. **Review the load script and data model:** joins, keys, synthetic keys and filters, in the data model viewer.
3. **Validate each KPI** with its business owner, then save it as a master measure with a description.
4. **Keep a set of test questions** with known answers and re-run them after each change.
5. **Check lineage** to confirm each field comes from the intended source.

These match the three checks from the AI session: a developer reviews the script, an analyst validates the measures, and the app follows company standards.

How do we keep semantic definitions from becoming a single point of failure when teams define the same metric differently?

Different definitions are often legitimate: finance margin and sales margin are different measures. The welcome session showed 16 possible versions of "margin in Latvia last quarter". The aim is that each version is named, described and owned.

1. **Name the variants:** "Gross margin (Finance)" and "Contribution margin (Sales)", each with a description and an owner, instead of two measures called "Margin".
2. **Tier the definitions:** company-wide certified KPIs, domain KPIs owned by one unit, and personal measures. AI should use the certified ones by default.
3. **Manage changes:** a business glossary, version control, and impact analysis before a change goes live.
4. **Share definitions as data products with contracts** rather than as copies, so every tool reads the same source.
5. **Test after every change** with questions whose answers are known.

In regulated EU banking, which governance frameworks or anonymisation strategies help move AI from experiments to value without compromising bank secrecy or GDPR?

Approaches used at financial institutions fall into three layers.

Governance process

- One intake for AI use cases, with risk tiering. Under the EU AI Act, credit scoring of individuals is high-risk, with obligations from 2 December 2027 after the 2026 Digital Omnibus postponement. Most analytics assistants are not high-risk.
- Model risk management, an accountable owner per use case, and human sign-off for decisions that affect customers.
- The EBA found no significant contradictions between the AI Act and existing banking law, so existing control frameworks can be extended.

Privacy techniques at data level

- Pseudonymisation following the EDPB's 2025 guidelines, with re-identification keys kept separately.
- In Snowflake: dynamic and tag-based masking, row access policies, aggregation and projection policies (minimum group sizes), differential privacy, synthetic data for development and testing, and data clean rooms for sharing without exposing records.

AI inside the bank's perimeter

- Snowflake Cortex with processing restricted to the EU.
- Qlik AI with permissions inherited from section access.
- MCP with per-user identity and audit logs.

For corporate customers, company data is mostly outside GDPR but remains under bank secrecy, while contact persons and sole traders are personal data. Classifying fields this way lets most corporate analysis run on masked or aggregated data.

Published examples: Danske Bank runs 30+ AI use cases through cross-functional adoption squads with shared standards and governance. Nordea's AI hub sets guardrails per use case, including

hallucination checks and source citation. Landkreditt (Norway) runs its Snowflake platform with dynamic masking and row-level security.

5. Platforms and architecture

Why set up everything in Qlik instead of other popular tools such as GitHub, Kubernetes or dbt?

We would not put everything into Qlik. Qlik's advantage is the small number of components: ingestion, modelling and analytics in one platform, and in Qlik Cloud no servers or clusters to operate. For a team whose main need is analytics, that is the shortest route to results.

The cloud adoption session described the turning point: once the same data serves many consumers (other BI tools, AI agents, data science), separate preparations in each tool become a burden. Shared logic then belongs on a data platform such as Snowflake, built as dbt models, versioned in GitHub, deployed through CI/CD and covered by automated tests. Qlik stays on top as the analytics and semantic layer. The session called this the "Modernize" and "Transform" scenarios.

Why use Airflow when Qlik Automate can build reload chains, add warnings, and run parallel or conditional reloads?

If the pipeline is mostly Qlik, Airflow is not needed. Qlik Cloud reload tasks can start when another task succeeds, fails or finishes, and Qlik Automate handles sequential and parallel reloads, conditions, loops and alerts without code.

Airflow is useful when orchestration goes beyond Qlik:

- one schedule for dbt runs, Snowflake tasks, external APIs, file arrivals and Qlik reloads
- pipelines defined as code in Git, with reviews, backfills and SLAs
- monitoring across many tools and teams in one place
- long reload chains that depend on upstream jobs outside Qlik

The two work together: Airflow can trigger Qlik reloads, and Qlik Automate can handle the analytics-side steps.

What are Snowflake's advantages compared with fast-growing competitors like Databricks?

Both are now AI data platforms, but they come from different origins. Snowflake delivers advanced AI as ready-to-use capabilities on governed data. Databricks, built on Spark, gives engineers a powerful toolkit to build and train their own AI.

Snowflake is stronger in:

- AI out of the box - LLM functions callable directly in SQL (Cortex AI), chat with your data in plain language (Cortex Analyst, Snowflake Intelligence), AI agents and document processing, all running where the data already sits
- AI on governed data - the same security, masking and access policies (Horizon) apply to AI as to BI
- SQL and BI workloads with many concurrent users
- Low operating effort - a fully managed service, with warehouses that size themselves
- Data sharing, a data marketplace and clean rooms

Databricks is stronger in:

- Building custom AI - training and fine-tuning own models, full ML lifecycle with MLflow, model serving (Mosaic AI)
- Agent engineering - tools to build, evaluate and tune AI agents on company data (Agent Bricks)

- Heavy Spark data engineering and streaming
- Code-first teams working in Python notebooks

The platforms are converging: both support open Apache Iceberg tables, both offer natural-language analytics and agents, and both now have a managed Postgres for transactional workloads. The difference is in approach - **Snowflake brings AI to the business as a ready capability**, Databricks gives engineers the tools to build it. For organisations driven by business analytics, with a small platform team, Snowflake gets AI into production faster and is simpler to run. For data-science-heavy organisations with strong Spark and ML skills, Databricks is worth evaluating. Qlik works on top of either.

Could you share publicly documented examples of companies that onboarded Snowflake, and the benefits?

Examples with published results:

Company	Industry	Published result
UPB AS	Construction, Latvia	ERP moved from on-prem Axapta to D365 F&O; two legacy data warehouse and reporting platforms replaced by one Snowflake foundation (dbt, ~300 transformations, CI/CD with automated test-to-production deployment). Presented at Snowflake Data & Donuts, Riga, 21 April 2026
NatWest	Banking, UK	80% lower total cost of ownership and carbon footprint after legacy systems were retired; a 40-person ESG team productive in 2 weeks instead of 3 months
Landkreditt	Banking and insurance, Norway	Queries from 30 minutes to about 1 minute, 3x faster development, 6x lower storage cost; dynamic masking and row-level security
Aviva	Insurance, UK	1.3 PB migrated from Oracle in 7 months; 8x faster processing
ABB	Manufacturing	Data from 40 ERP systems and 15 SAP BW instances consolidated; over \$200M a year in inventory savings attributed to the data programme
Siemens	Industrial	On-premise SAP HANA data lake replaced; more than 1.5 billion SAP changes a day replicated from 50+ systems
Penske Logistics	Logistics	AI summarisation model built on Cortex in under 15 days

What are the most impactful use cases where Qlik and Snowflake together delivered measurable business outcomes?

The two cover different layers. Snowflake holds the shared, governed data (the gold layer) that every tool can use. Qlik loads the data it needs into its in-memory associative model and adds the business context on top, the "platinum layer" from the AI session. Selections, drill-downs and AI questions are then calculated by the Qlik engine in memory, not sent to Snowflake as new queries. Snowflake compute is used mainly during scheduled, often incremental reloads, which keeps consumption predictable. Direct Query to Snowflake remains an option where data must be live or is too large to load.

Few global examples:

- [Ewals Cargo Care](#) (logistics): new data sources in 5 days instead of 20+, and 75% lower ETL development cost with Qlik data integration into Snowflake.

- [Urban Outfitters](#): 50 GB a day streamed into Snowflake, with dashboards refreshing every 2 minutes.
- [Discover Financial](#): billions of records moved from Oracle to Snowflake with Qlik Replicate.

What are the possibilities of Qlik Sense for visualising an institution's analytics and building KPI dashboards?

Qlik Sense suits public-sector KPI work:

- KPI dashboards where management can go from a headline indicator to the unit, region or case behind it in a few clicks
- data from several systems (finance, HR, case management, registers) combined in one model
- access by role, so each department sees its own data and management sees all of it
- scheduled reports and alerts, for example a monthly report to the board or an alert when a KPI falls below target
- plain-language questions (Qlik Cloud) for people who do not use dashboards daily

A practical start is to agree 10-15 KPIs with clear definitions and owners before building.

6. Licensing and support

How does Qlik Cloud pricing work: users, data volume, consumption, or a combination?

A combination, and the model depends on how Qlik is deployed:

- **User-based:** licences per named user, with different roles for people who build and people who consume analytics. This is the usual model for client-managed Qlik Sense and in many existing contracts.
- **Capacity-based:** the price is driven mainly by the volume of data loaded for analysis, with broad user access. This is the main model for Qlik Cloud Analytics today. AI usage (Qlik Answers questions, including MCP calls) and some features, such as Qlik Predict, depend on the edition.
- **Deployment:** Qlik Cloud (SaaS), client-managed on your own servers, or hybrid, where on-premise apps are distributed to Qlik Cloud and on-premise sources connect through the Qlik Data Gateway.

Qlik Talend Cloud (data integration) is priced on the volume of data moved and on job runs. The right model depends on the number of users, data volumes and where the data has to stay, so it is worth comparing the options on your own figures before choosing.

What support tiers and models does Qlik offer?

- **Qlik Standard support** is included in the subscription, with 24x7 coverage for critical issues.
- **Partner support:** as Qlik partner, The Infotrust provides first-line support in the local language, administration, upgrades and further development, usually as a monthly service plan.

For client-managed Qlik Sense, each release is supported for two years, so regular upgrades are part of running it.

How often is cloud cost revised compared with on-premise cost?

- **Qlik Cloud:** a subscription with the price fixed for the contract term, reviewed at renewal. Capacity can be added during the term. Changes to Qlik's list prices apply at renewal, not to running contracts.
- **On-premise:** licence and maintenance are reviewed at each renewal as well. On top come hardware refresh every few years and an upgrade project at least every two years.
- **Snowflake:** consumption-based, so cost follows usage month by month. The main lever is your own usage, not the price list.

A fair comparison looks at total cost of ownership, including staff time for servers, upgrades and monitoring, not only at licence lines.

7. Questions on the customer stories

The first three questions refer to the commodity-market case from the morning keynote "Cultivating High Value": an intermediate processor that built one model linking market prices, recipes and contracts, giving margin visibility up to 18 months forward.

How do you check forecast accuracy after six months? Do you improve based on the findings?

Yes. Inaccurate cost forecasts were one of the pain points in the case, so accuracy is tracked as part of the solution:

1. **Every forecast is stored with its date**, so it can later be compared with what actually happened.
2. **Accuracy is measured by horizon**, for example 1, 3 and 6 months ahead, with error and bias measures (such as MAPE and bias) per product group and market.
3. **Accuracy is shown next to the forecast**, so users see how far to trust it.
4. **Models are retrained** on a schedule or when accuracy falls below an agreed level. Qlik Predict monitors data drift daily and keeps model versions for comparison.
5. **Findings go back into the inputs**. Most improvements come from better input data (market prices, contract volumes, seasonality) rather than a different algorithm.

Based on the implemented solution, did the company consider fully automated decision-making for contracting?

No. The goal was that sales, procurement and finance know the margin before signing a contract, based on the same numbers, and can react to market moves within hours. People still make the decision. The result presented was 60% faster decisions on hedging and contracting and up to 4% margin improvement.

Automation is a possible next step for routine cases, for example approving contracts that fall within agreed price and margin limits and routing exceptions to people. For decisions with high financial exposure, a person should stay in the loop.

How did you get the client to engage and reach such a good outcome?

Usage grew by about 600% through an adoption programme with three parts, as shown in the keynote:

1. **One view**: sales, production and procurement decide on the same numbers.
2. **Ownership**: each unit owns its part of the common goal.
3. **New routines**: processes were adjusted to how decisions are now made.

The change described on stage was from a volume mindset to a value mindset: the company now knows the margin before it signs the contract.

Qlik Predict: can I use a combination of models to get higher accuracy than with a single model?

Partly. Qlik Predict trains several algorithms automatically and ranks them, so the best one can be deployed. Several of these algorithms (Random Forest, XGBoost, LightGBM, CatBoost) are themselves ensembles of many models.

Qlik Predict does not stack or blend separately trained models. Two ways to do it:

- Deploy two or more models and combine their predictions in the Qlik load script, for example as a weighted average, then compare the accuracy of the combination with each model alone.
- Build the ensemble in Snowflake ML or Python and load the scores into Qlik.

Better input features (new data sources, seasonality, events) usually improve accuracy more than combining models.

Does process mining work only with detailed CRM statuses, or can it use raw data such as field updates, button actions and system logs?

It works with raw data. As shown in the process mining session, an event log can be built from SRM, ERP, warehouse, OCR, point-of-sale or CRM logs, and a few columns are enough: case ID, activity and timestamp. Change logs and field-update history (for example SAP change tables) are common sources.

The main effort is interpretation: one manual change can trigger many automatic system changes, so part of the work is deciding which events are real business steps.

Button clicks on a user's screen belong to a separate discipline, task mining, which records desktop activity. Most organisations start with system logs.

Which tool was used to capture the process data points, and how does it work, in short?

The Swedbank case uses mpmX, which runs on Qlik. Process mining works from digital traces (event logs) that the bank's systems already record. mpmX turns them into a process model with variants, timings and frequencies. In the process presented, about 30 steps produced more than 12,700 variants over five years.

Did process mining lead to changes in roles or staffing, such as reskilling, redeployment or headcount reductions?

The results presented concerned how the work is done:

- a process design change reduced rework volume by about 85%
- faster consultations improved service levels and contributed to higher conversion
- better workload distribution supported service during peaks, absences and holidays

The solution has views for three groups: process managers use it to evaluate and improve processes and compare countries; team leads use it to manage team performance and capacity; management uses it to compare areas and find cross-functional opportunities.

Sources

Qlik

- [Qlik AI security and Amazon Bedrock](#)
- [Qlik AI platform - no training on customer data](#)
- [How Qlik built AI on Amazon Bedrock \(AWS, Oct 2026\)](#)
- [Qlik cross-region data processing](#)
- [Qlik MCP server tools](#)
- [Administering Qlik MCP](#)
- [Qlik MCP audit events](#)
- [Qlik Answers with client-managed apps](#)
- [Qlik Sense May 2026 release notes](#)
- [Qlik Answers Fast and Think more modes](#)
- [Qlik Connect 2026 announcements](#)
- [Qlik Predict algorithms](#)
- [Qlik Predict drift monitoring](#)
- [Qlik lineage and impact analysis](#)
- [Qlik agent skills](#)
- [Qlik task chaining with Automate](#)
- [BARC benchmark \(Qlik press release\)](#)
- [Qlik at Snowflake Summit 2026](#)

Snowflake and Anthropic

- [Snowflake AI features and data use](#)
- [Snowflake cross-region inference](#)
- [Snowflake differential privacy](#)
- [Snowflake synthetic data](#)
- [Snowflake budgets](#)
- [Idle warehouse spend study](#)
- [Anthropic - is my data used for training](#)

Regulation, banking and process mining

- [EU AI Act Digital Omnibus in force](#)
- [EBA - AI Act implications for banking](#)
- [EDPB pseudonymisation guidelines](#)
- [Danske Bank - adopting AI in regulated industries](#)
- [Nordea - how we use AI](#)
- [Extracting event data from databases \(van der Aalst\)](#)
- [mpmX](#)