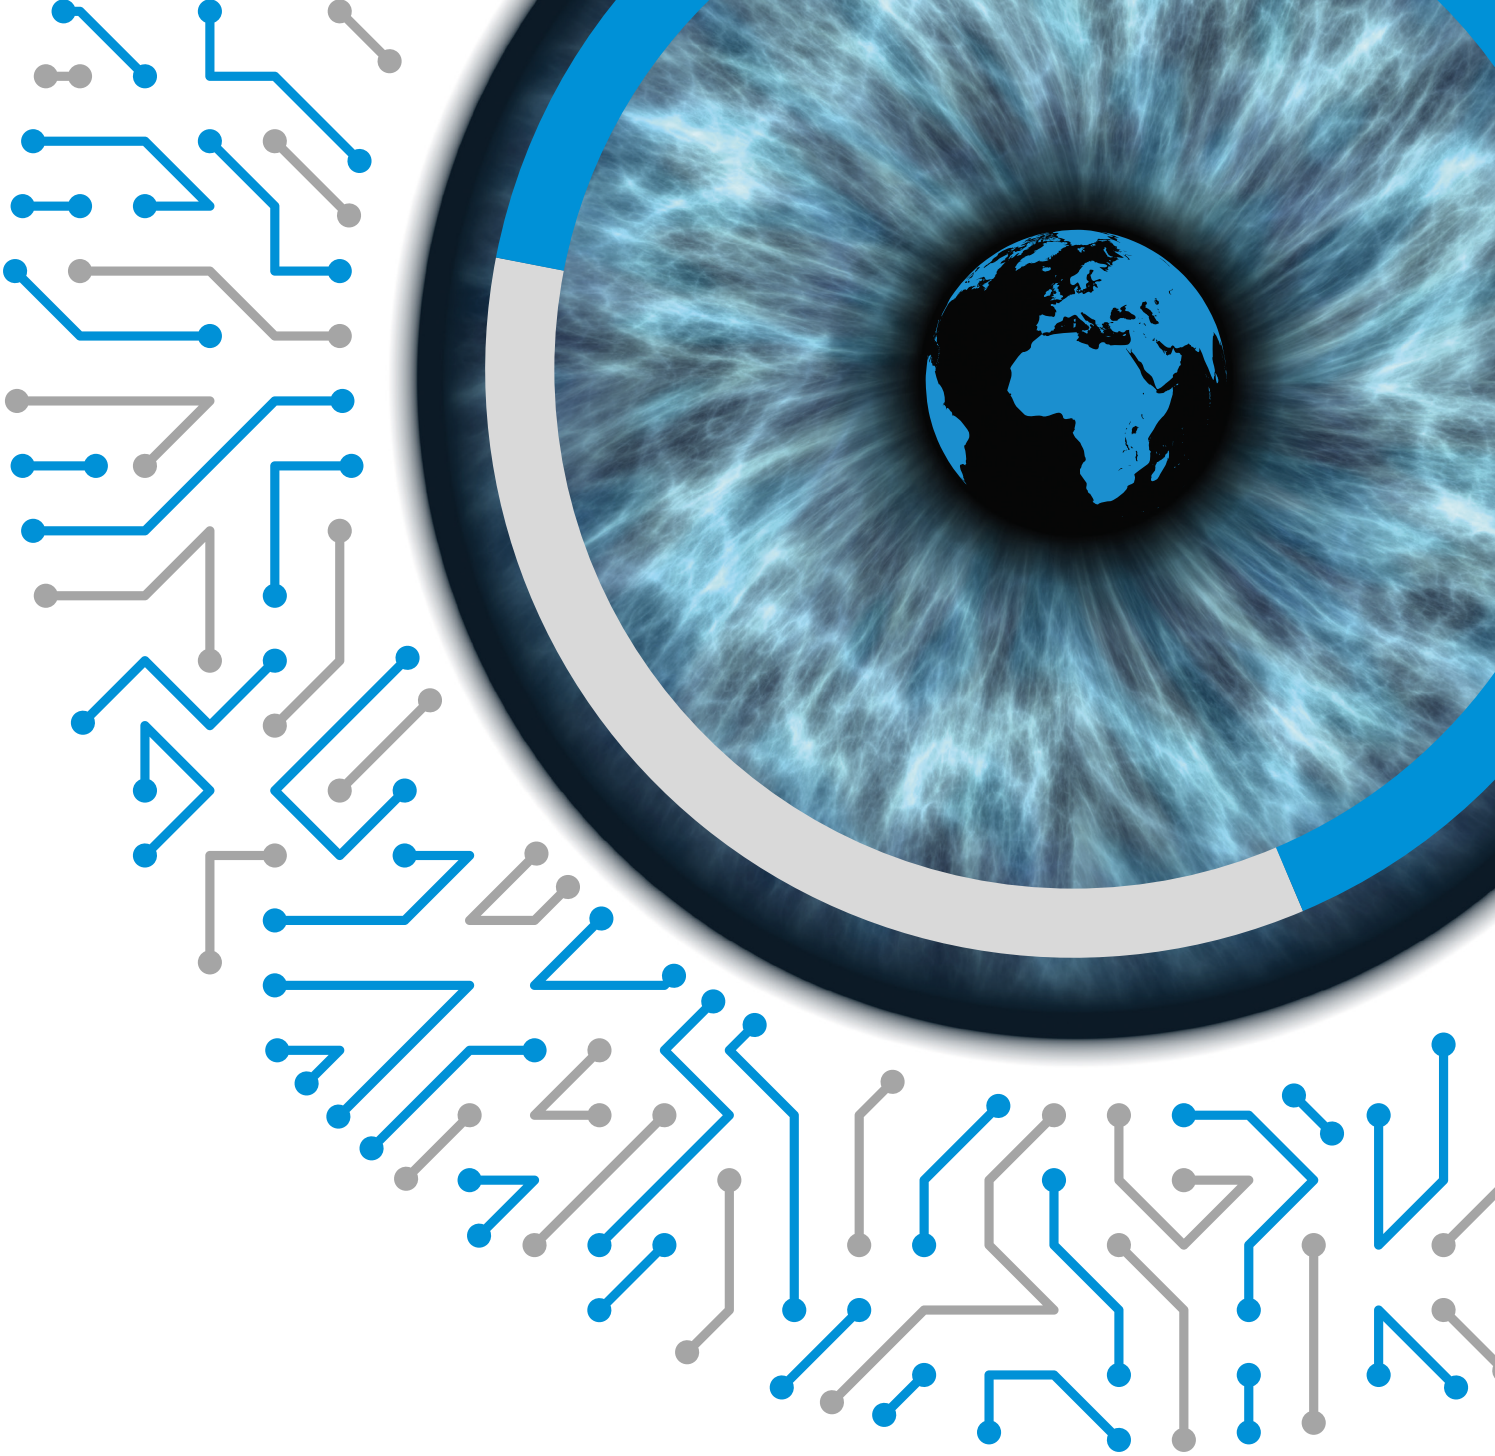

# REGTECH100

Profiles of the **REGTECH100**, the world's most innovative RegTech companies that every leader in the regulatory industry needs to know about in 2026

REGTECH ANALYST

FINTECH GLOBAL



GLOBAL
REGTECH
SUMMIT

20 MAY | LONDON

# Join the world's largest gathering of RegTech Leaders & Innovators

www.GlobalRegTechSummit.com

# REGTECH100

The **REGTECH100** is an annual list of 100 of the world's most innovative RegTech companies selected by a panel of industry experts and analysts. These are the companies every financial institution needs to know about as they consider and develop their mission critical RegTech and digital transformation strategies.

There's plenty of interest and hype about RegTech in the marketplace, but much of it is superficial, incoherent or self-serving and fails the needs of decision-makers in incumbent financial institutions who require independent, facts, figures and analysis.

The **REGTECH100** list will help senior management and compliance professionals evaluate which digital solutions have market potential and are most likely to succeed and have a lasting impact on the industry.

---

## CRITERIA

A range of factors was considered by the Advisory Board and RegTech Analyst team to make the final selection including:

- Industry significance of the problem being solved;
- Growth, in terms of capital raised, revenue, customer traction;
- Innovation of technology solution offered;
- Potential cost savings, efficiency improvement, impact on the value chain and/or revenue enhancements generated for clients;
- How important is it for financial institutions to know about the company?

---

## PROCESS

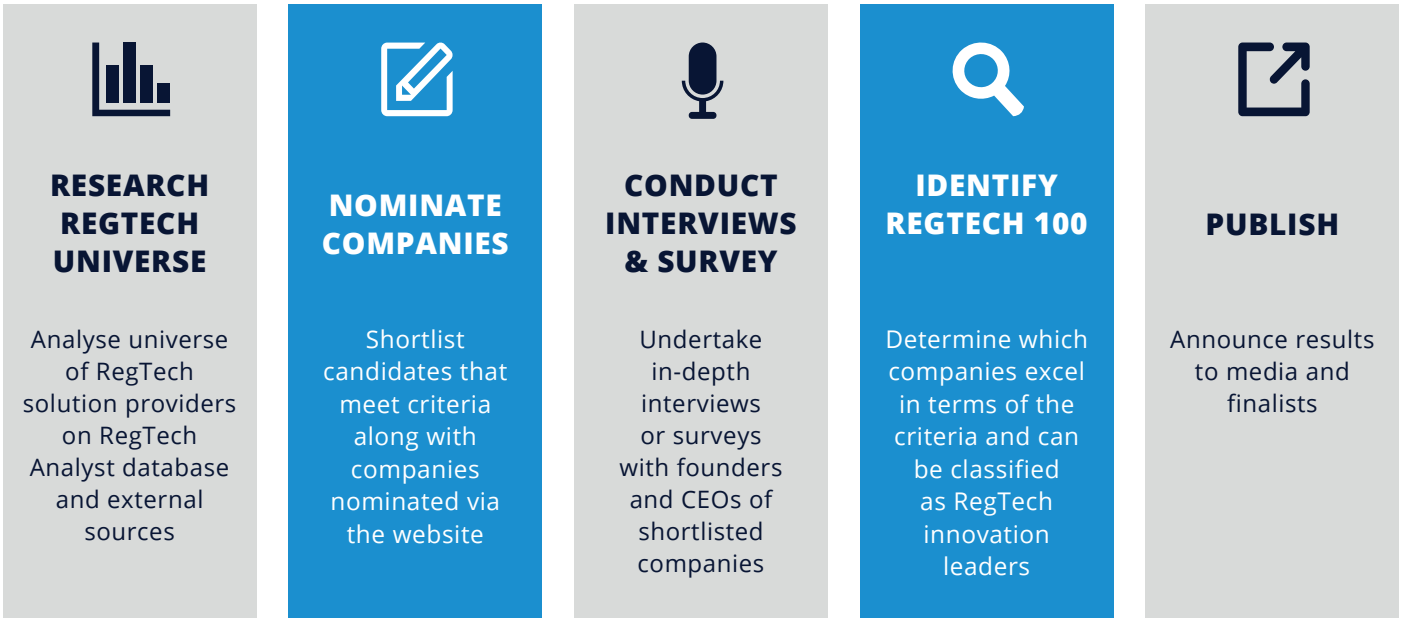

REGTECH ANALYST

# COMPANY RESEARCH PROFILE

BEHAVOX

Unified Surveillance, Unified Archive, Unified Controls.

Founded 2014
London, United Kingdom
www.behavox.com
marketing@behavox.com
Employees: 250-500

Regions of operation:
North America, EMEA, APAC, and Latin America

**KEY EMPLOYEES:**

**Erkin Adylov**
Founder & CEO

Subsectors: **Communications Surveillance, Trade Surveillance, Archiving & Data Retention, Control Room and Conflict of Interest, Policy Management**

Regulation/Legislation: **SEC, FINRA, NFA, Federal Reserve (SR 11-7), FCA, ESMA, MAS, IIROC, JSDA, SEC 17a-4, MiFID II, EU MAR, DORA, GDPR, FCA SYSC, PDPA.**

## OFFERING

Behavox is a Controls Platform for global banks, asset managers, hedge funds, and commodity firms. It unifies communications surveillance, trade surveillance, compliant archiving, and policy management on a single AI-native stack. Surveillance delivers the detective controls, policy management provides the directive and preventive controls, and compliant archiving underpins the evidence, retention, and record-keeping.

Each product can run independently, but together they operate as one ecosystem on one data layer, one policy framework, and one case workflow. This integrated approach reduces alert volumes, manual effort, and total cost of ownership versus managing multiple point solutions.

Behavox offers proven, agentic AI for controls — not experimental chatbots. Our purpose-built LLMs have been in production for more than three years and have passed dozens of internal audit, regulator, and model validation reviews, giving firms a controls platform that is both state-of-the-art and battle tested.

Behavox helps firms consolidate their tech stack away from legacy tools and reduce total cost of ownership — on a platform they can explain, evidence, and defend in front of regulators.

## PROBLEM BEING SOLVED

Regulators increasingly expect a coherent, end-to-end control framework with clear links between regulatory obligations, internal policies, and controls. In practice, most firms have built controls incrementally by region, product, and risk type, creating a patchwork of tools that are hard to govern, explain, or change at speed.

Controls for communications, trading, archiving, and policies typically sit in separate systems with separate logic. The same conduct or market abuse risk must be defined several times, often with inconsistencies. When an issue arises, teams stitch together chats, emails, voice, and trades manually, which is expensive and error-prone. Expanding coverage to new channels, products, or regions usually means hiring more people, not improving the quality of controls.

Legacy rules- and lexicon-based tools generate large volumes of low-quality alerts, while many "AI" offerings rely on generic LLMs that fail model validation. Compliance leaders are left with high cost, high operational risk, and rising expectations, but no unified, AI-native controls platform that aligns detection, policies, and records in a way they can confidently stand behind.

## TECHNOLOGY

Behavox has been in RegTech since 2014 and is now used by more than 100 leading institutions globally, including one central bank and one regulator. Backed by over $200m in R&D investment, the platform is purpose-built for controls: ingesting and normalizing data across channels, applying AI-driven detection, and managing policies, cases, and audit trails on a single data and policy layer.

Behavox partnered with Google Cloud to build a cutting-edge technology stack that is designed for very large scale. Some customers use Behavox to monitor communications in 15 languages across more than 70,000 employees. Many customers are moving toward full-population coverage with Behavox across all employees, all channels, and all languages — because AI makes that level of surveillance both affordable and effective.

At the core of our products is Behavox's own family of purpose-built LLMs and agentic AI components. We do not bolt on a generic LLM; we apply specific AI to specific control problems. Models are tuned for specific use cases, with extensive documentation that makes outputs explainable, transparent, and governable.

## PRODUCT DESCRIPTION

Behavox offers an ecosystem of products built for controls. Customers can deploy individual modules or run them together as an integrated controls stack. Quantum provides AI-native communications surveillance across voice, chat, and email in 15 languages, increasing true positive detection while reducing cost.

Polaris delivers trade surveillance that can operate alone or alongside Quantum, with agentic AI automatically pulling related chats, emails, voice, and archive records into a single case to reduce manual investigations. Intelligent Archive gives firms a unified, compliant archive with consistent retention, search, and trade reconstruction. Pathfinder manages the full policy lifecycle and links policies and controls to underlying risks and regulations, helping firms move from purely detective controls to directive and preventive controls.

## TRACTION/GROWTH

- **Clients** – J.P Morgan, BNY, TD Securities,Danske Bank, Brevan Howard, Invesco, PWP, TJC, SMBC, Vitol
- **Partners** – Google Cloud, LeapXpert, Custodia, Baringa, Nomura Research Institute, TeleMessage

# Why Trade Surveillance and Communications Surveillance Must Be Unified - and the Efficiency Gains This Unlocks

*By Erkin Adylov, CEO and Founder, Behavox*

Most firms still monitor trades and communications in silos, creating duplicate work, inconsistent controls, and slow reconstructions that drain budgets and weaken risk coverage.

Different vendors. Different data. Different rules. Different teams.

That model was tolerable when systems were smaller and regulatory expectations narrower. But a siloed,

rules-based approach breaks down when you are trying to monitor all employees, all channels, all major products, in all key geographies. Under that load, fragmented surveillance is not just inefficient – it is increasingly hard to defend to regulators and investors.

Behavox's view is simple: trade surveillance and communications surveillance must be unified on top of a unified, compliant archive, powered by AI that understands both sides of the picture. That is the design behind Behavox Polaris (trade surveillance), Behavox Quantum (communications surveillance), and Behavox Intelligent Archive.

## One archive, one policy, one alert stream.

In the unified model:

- All data sits in one archive – orders, trades, RFQs, market data, emails, chats, and voice. You can search across any of them from a single place.
- Detective policies for market abuse and conduct risk are defined once and applied consistently across trade and comms, with alerts presented in a unified alert management layer.
- Alerts become cross-referenced – a comms alert can automatically pull the related trades and close obvious false positives; a trade alert can pull the relevant communications and, again, close what is clearly benign.

We take it as a given that both Polaris and Quantum must use AI-based detection, not just rules and lexicons. The real question is how to use agentic AI to push detective controls much further than a single model on a single dataset.

## Where AI really moves the needle

We see three areas where AI can eliminate structural pain and unlock meaningful efficiency.

1. **Integration and field mapping**
   Today, getting a trade surveillance system live takes months and requires expensive specialist resources. Most of that time is spent on field mapping: aligning a firm's data to the vendor's schema. This is exactly the kind of repetitive, pattern-matching work AI should do. In Polaris, AI proposes the mappings; humans review and confirm in a UI. Integrations that used to take months can be delivered in weeks.

2. **Search and investigation for non-quants**
   Trade data has historically been the domain of quants and technologists. With AI, that changes. An investigator should be able to ask: "How many times in the past six months has this trader done a similar trade?" Polaris can translate that into SQL, run the queries, analyse the results, and present a summary, along with the underlying queries and data. In effect, you have a "mini-quant" available to every reviewer.

3. **Contextualising and closing obvious false positives**
   The easiest and most powerful agentic use case is context. An unusual P&L or trade pattern is a problem if it happens before relevant news breaks; it is not necessarily a problem if the news is public and the desk is simply faster than the market.

AI can pull the news, read and time-stamp it, compare it to the trade, retrieve relevant chats or emails, and close an alert as a false positive – while documenting why. Alerts can be routed to QA queues or L1 review where needed. Behavox supports rich alert states and QA workflows so firms can calibrate how much they allow AI to auto-resolve.

Unifying trade and communications surveillance on a single archive, and using AI not just for scoring but for integration, investigation, and context, is how firms move from fragmented, expensive surveillance to controls that scale with the business and stand up in front of regulators and investors ●

# How purpose-built AI is transforming compliance from detective controls into preventive controls

*As AI rewrites the playbook for financial crime and conduct risk, Behavox founder and CEO Erkin Adylov explains how purpose-built models are shifting compliance from fragmented, detective controls to unified, preventive ones.*


BEHAVOX

Unified Surveillance, Unified Archive, Unified Controls.

### Q1. Behavox is best known for communications surveillance. Why are you evolving into an end-to-end controls platform, and why now?

**A:** Because our customers pushed us there.

Over the past three years, we've proven that AI can transform communications surveillance: four to five times more true positives, far fewer false positives, and better outcomes with regulators. Once that became the baseline, customers started asking a very simple question: **"If AI can fix comms surveillance, why stop there?"**

They don't want one system and one story for comms, a different one for trades, another for archiving, and a separate stack for policies. They want **fewer systems, simpler controls, and a single narrative** they can explain to regulators, boards, and investors.

So evolving into an end-to-end controls platform is not a distraction from surveillance; it's the logical next step. We're taking the same AI stack, the same data layer, and the same governance discipline and applying them to trade surveillance, record-keeping, and policy management. The benefit to customers is a stronger overall control framework and fewer fragmented tools to defend in exams, investigations, or enforcement discussions.

### Q2. Many institutions still struggle with fragmented systems and data silos. How does unified data and unified logic change what's possible for the control framework?

**A:** You can't build serious, defensible controls on top of fragmented data and inconsistent rules.

> ***"In most institutions, comms surveillance, trade surveillance, archiving, and policy management all sit in different systems. "***

The same risk—MNPI misuse, market abuse, conflicts of interest—is defined separately in each tool. When something happens, teams stitch together chats, emails, voice, trades, and records by hand. It's slow, expensive, and error-prone.

At one large "megabank", they had exactly this problem: a patchwork of legacy systems that worked in individual regions but **never scaled globally.** Regional variations, incompatible data models, and poor integration meant they could not get a single, consistent view of risk or controls.

With a unified controls platform, you define a risk **once** and apply it consistently across surveillance, archive, and policy.

Behavox does this through **AI Risk Policies**—machine-readable policies that encode regulatory obligations and internal standards directly into the system. When a scenario fires, the platform can automatically pull in the relevant communications, trades, records, and policy references into a single case. In a hedge fund context, that might mean seeing a PM's wall crossings, research interactions, trade blotter, and messages with the street in one place.

Once you have that integrated view, you can see which policies are effective, where gaps exist, and what needs to change. Unified data and logic turn the control framework from a patchwork into something **coherent, explainable, and much easier to defend.**

## Q3. Where do you see AI having the biggest impact in trade surveillance, and how does Polaris illustrate that?

**A:** In trade surveillance, AI has two very concrete jobs: **context** and **speed of deployment.**

The first is context. A single trade alert on its own never tells the whole story. With **Polaris**, we use agentic AI workflows to pull in contextual data automatically: related chats and emails, voice, control room communications, news, and corporate events. The AI reviews that bundle and helps decide whether something looks like a genuine issue or an obvious false positive.

That's exactly what human surveillance teams do today, but they do it manually and slowly. Very often, closing an alert is as simple as matching a trade to a wall-crossing record or a news timestamp. AI can do that work at scale, and even close **obviously benign alerts** on its own, so humans focus on **QA and judgement**, not chasing basic context.

The second is onboarding and integration. Today, getting a trade surveillance system live can take months or even years, largely because of painstaking **field mapping**—aligning each firm's data model to the vendor's schema. It's labour-intensive, repetitive work.

We believe AI should do that. With Polaris, AI helps infer and map different data types and sources from the client environment into the schema needed to generate alerts. We already see that AI can perform this mapping work quickly and accurately, while humans review and confirm. That shifts the model from "armies of people building integrations" to "AI does the heavy lifting, experts check and approve", which is faster, cheaper, and much easier to scale across desks, entities, and regions.

## Q4. Once AI has improved surveillance and reduced noise, why do you see preventive controls as the natural next step?

**A:** The first step was to make **detective controls** actually effective.

When you move from lexicon-based systems to purpose-built AI, you get **four to five times more true positives and dramatically less noise.** You start to see patterns clearly: which desks, PMs, products, channels, and behaviours generate real risk. A UK-based hedge fund CCO put it very directly to us:

"Regulators like that we use AI and Behavox, but our investors love it — because it shows we're serious about identifying risk and safeguarding their capital."

That's the value of investing in high-quality AI detection: regulators see stronger controls, and investors see a firm that takes risk seriously.

Once you have that level of visibility, the question changes from "Can we spot problems?" to "Why are we only reacting?" Preventive controls are the next layer. You use what AI is seeing in surveillance and trading to **inform how policies are written, how attestations are structured, how the first line is supervised, and what gets escalated.** Policies stop being static documents; they become part of a closed loop between obligations, behaviour, and outcomes.

> ***"You still need strong detective controls and evidence, but the emphasis shifts from catching issues late to designing the environment so they are less likely to happen in the first place."***

## Q5. Many vendors are partnering with big LLM providers and "adding Copilot" to their systems. Why did Behavox choose to build its own LLMs, and what does "purpose-built AI for controls" actually mean?

**A:** Because in our world, the AI isn't a convenience feature—it is the control.

Most "Copilot-style" integrations focus on nice-to-have capabilities: summarising, searching, drafting. Those can be useful, but they're not what regulators care about. What matters is the AI that **actually drives detection, prioritisation, escalation, and case outcomes.** That has to be engineered like any other critical control.

We built our own LLMs and AI stack for three reasons.

First, **governance and stability.** If you rely purely on a generic LLM API, you don't control the training data, the update cycle, or the underlying behaviour. That's a problem when model risk, internal audit, or regulators ask for documentation, reproducibility, and change control. With our own models, and with **AI Risk Policies** on top, we can show exactly how the model is configured, what it is optimised for, and how it has changed over time.

Second, **fitness for purpose**. Our models are trained and tuned on conduct risk, market abuse, regulatory language, and control workflows—not on general internet text. They're built to spot specific patterns in trades and communications, link those patterns to policies, and generate evidence you can stand behind in an investigation or exam.

Third, **track record**. We've invested over $200m in R&D and have had AI in production for three years across more than 100 institutions, including a central bank and a regulator as customers. In that time, our models and surrounding processes have gone through internal audit, model validation, monitors, and regulators. That is exactly the kind of scrutiny you want if the AI is part of your control framework.

> ***"There are places where generic LLMs can add value at the edge, but for the core of the controls stack we believe you need purpose-built, owned, and governable AI, not a black box you rent by the token."***

## Q6. What are the main risks of relying on generic LLM tools for compliance, and how should firms think about trusting AI when it is part of their control environment?

**A:** The main risk is mistaking a general-purpose assistant for a **governed control.**

Generic LLMs are extraordinary tools, but they aren't optimised for your risks, your regulations, or your control framework. If you ask a generic model to "help with compliance", you might get clever answers, but when a regulator or model risk committee asks why it produced a particular output, you may not have a defensible explanation.

Governance is another issue. You often have limited visibility into the model's training data or update cycle. That doesn't align with the documentation, stability, and change-control expectations placed on critical controls. We've even seen generic LLMs give **different answers to the same conduct scenario** when the wording changes slightly. That's unacceptable if your name is on the CCO attestation.

In terms of trust, I think firms should look at three things: **evidence, engineering, and transparency.**

- Evidence: Is the AI already in production in environments like yours? Has it survived internal audit, model validation, and regulatory scrutiny?
- Engineering: Is the stack built specifically for controls, with clear configuration, logging, and AI Risk Policies that tie behaviour back to obligations and risks?
- Transparency: Can you document how it works, challenge its behaviour, and explain it to a regulator?

That's how we've built Behavox: three years of production AI, a stack designed specifically for controls, and a way of working with customers that makes the AI feel like a **well-understood part of the control framework**, not a black box.

## Q7. Looking ahead, how will AI reshape compliance over the next few years, and what role do you see Behavox playing?

**A:** I see three big shifts.

First, **full-population coverage**. AI makes it realistic to monitor all relevant employees, channels, and languages. We already have customers using Behavox to monitor communications in 15 languages across more than 70,000 employees, and many are moving toward full coverage because it's now affordable and effective.

Second, **integrated controls**. Instead of isolated tools for surveillance, archiving, and policy management, firms will move to coherent platforms where data, logic, and evidence are shared. That makes it easier to respond to exams, defend decisions, and explain the control framework to boards, regulators, and investors.

Third, **prevention.** As AI-driven detection gets better, the real value shifts to using those insights to design better policies, better first-line controls, and better training. The line between "monitoring" and "prevention" will blur.

Behavox's role is to be the partner that makes that transition **safe and credible**: the most compliant AI stack in the industry, the most integrated controls platform, and a roadmap driven by what our customers need to solve next, not by marketing trends. Ultimately, we want firms to have effective, efficient, defensible controls so they can focus on running the business and generating returns ●

BEHAVOX

# One Platform.
# All Controls. Unified.

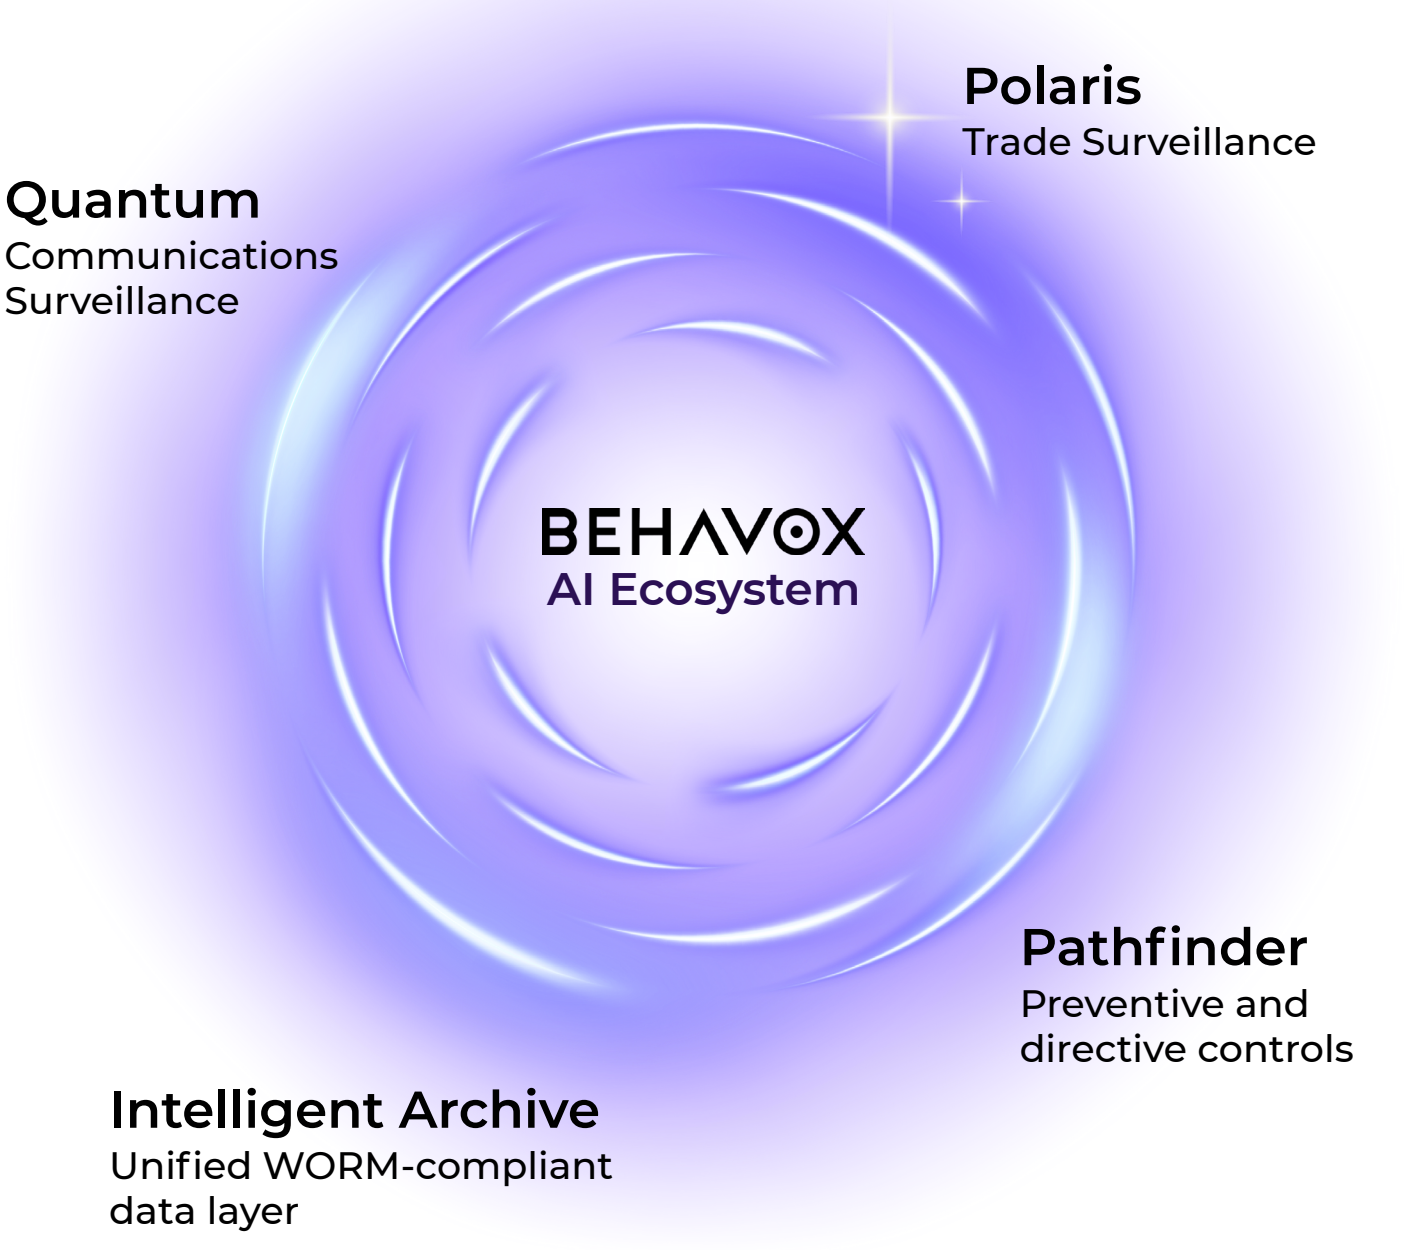


Experience the Future of RegTech

www.behavox.com

## ABOUT US

This summary was produced by RegTech Analyst.

The RegTech Analyst platform offers business intelligence on the RegTech, risk management tech and cybersecurity sectors. RegTech Analyst is the pre-eminent provider of data, research and analysis on the global RegTech market. We cover every trend, every investment and profile every company that provides a technology solution for compliance, risk management or cybersecurity. We deliver essential intelligence for mission-critical business decisions.

For more information, please visit:

www.RegTechAnalyst.com

## ABOUT US

FinTech Global is the world's leading provider of FinTech information services, B2B media products and industry events.

We inform, promote and connect FinTech buyers, sellers, investors and innovators worldwide.

We serve a network of over 300,000 FinTech professionals from market-leading organizations – financial institutions, technology innovators, corporate investors, venture firms and expert advisory firms. We enable them to get the information they need to make better business decisions and to connect and engage with the people and organisations they want to do business with.

For more information, please visit:
www.FinTech.Global/about

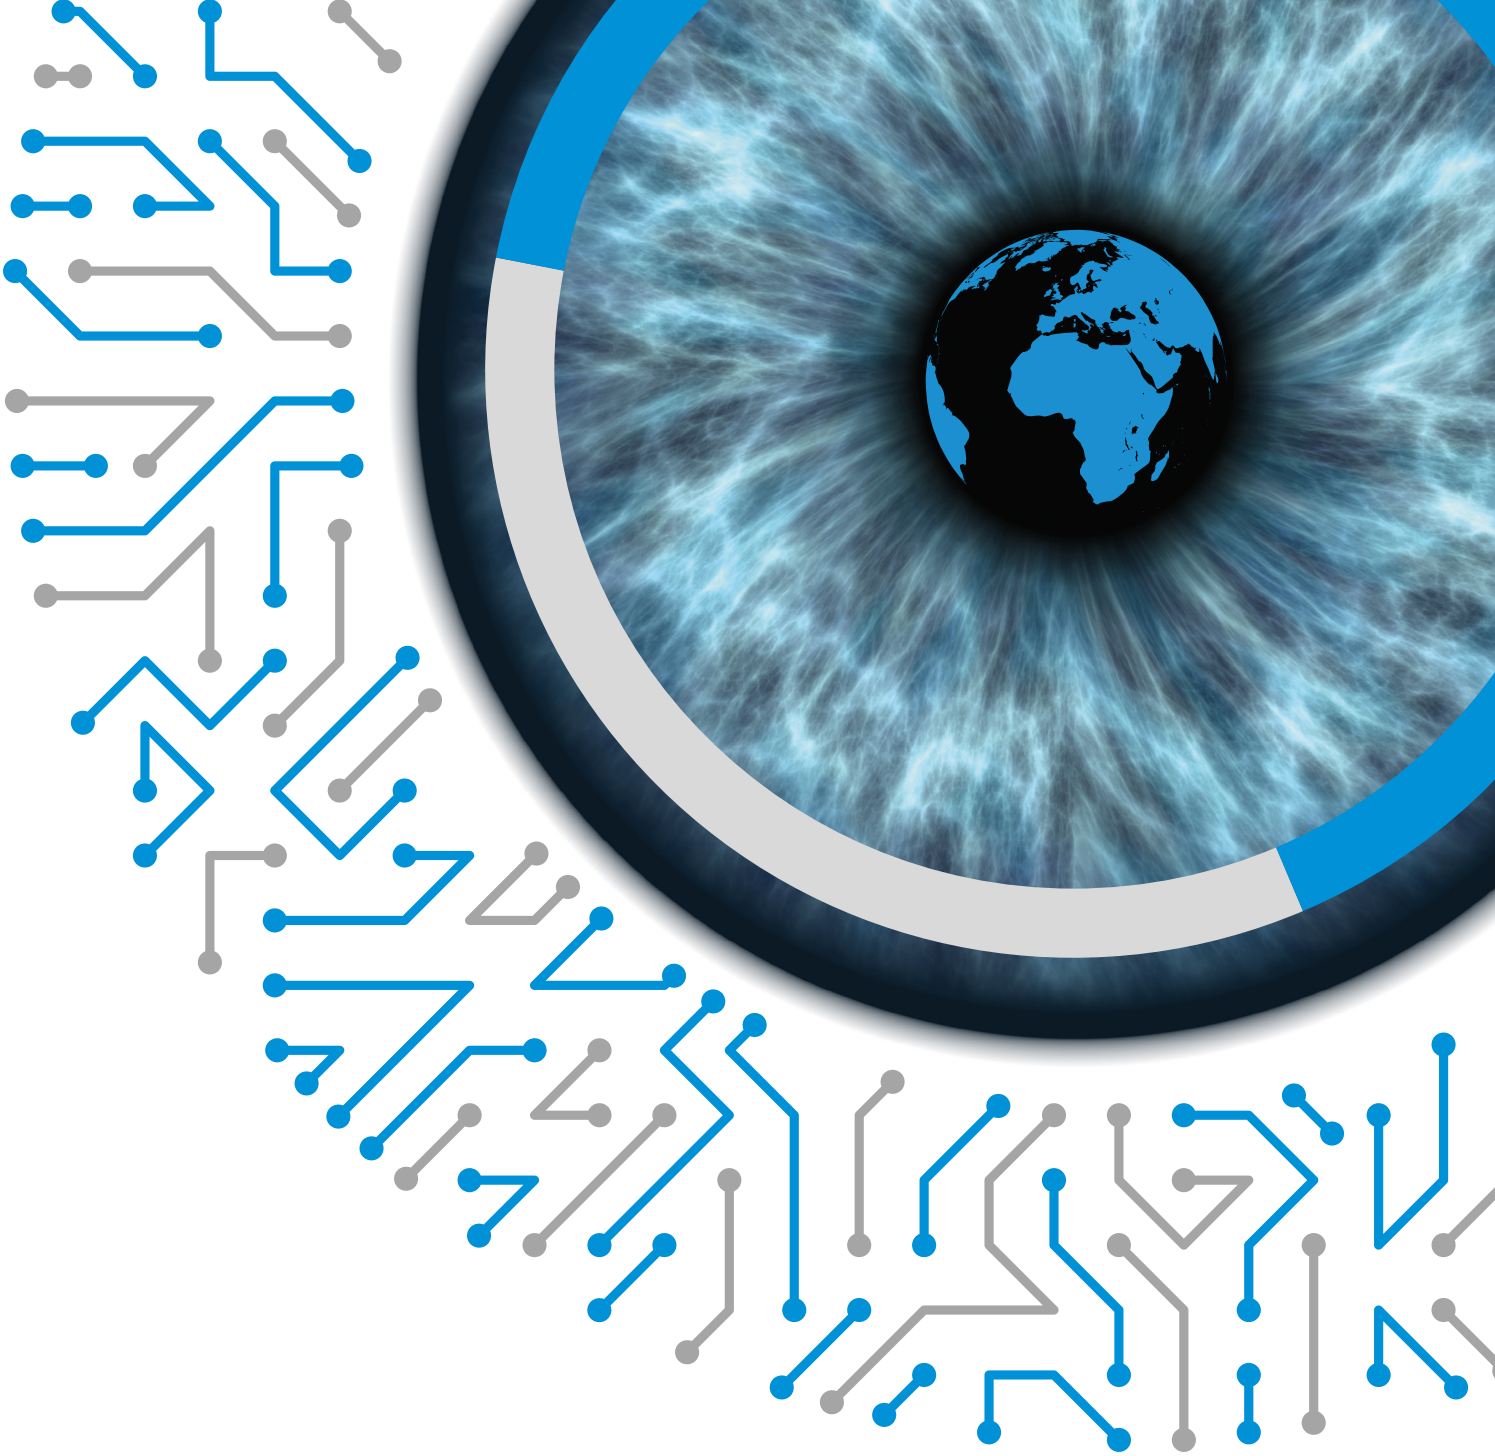

For more information contact info@fintech.global

REGTECH ANALYST

FINTECH GLOBAL