A-Team Innovation
Awards 2026

A-Team Innovation Awards 2026 Winners' Report

www.a-teaminsight.com
@ATeamInsight
ATeam Insight

A-Team Innovation Awards 2026

The results are in. We are thrilled to unveil the winners of the 2026 A-Team Innovation Awards, a tribute to the visionaries reshaping the financial landscape.

Now in its sixth year, these awards have become a benchmark for excellence. Our 2026 recipients particularly stood out for their ability to turn emerging technologies into high-value realities for capital markets. This year's celebration honours the projects and personnel across the vendor and practitioner communities who are setting new standards in Data Management, Trading Technology and RegTech,

Sifting through the submissions for our 40+ categories was an immense challenge. The calibre of ingenuity seen over the last 12 months reflects an industry in rapid evolution.

Winner categories included Most Innovative Financial Technology Executive, Best AI for Operational Resilience, Most Innovative use of Generative AI and Most Innovative Use of Open Source/Cloud Technologies.

We extend our sincere thanks to every firm that shared their breakthroughs with us. Our deepest appreciation also goes to our independent advisory board, whose collaboration with our editorial team ensured a rigorous and fair selection process.

To our winners: your contributions are driving the industry forward, and you should be immensely proud of standing out in such a competitive field.

Are you a solution provider redefining the future of financial data or regulatory reporting? Don't miss your chance to be recognised next year. **Click here** to be alerted when A-Team Group announces entries for the 2027 awards.

.

Andrew Delaney
President & Chief Content Officer
A-Team Group

I'd personally like to thank our advisory board for their time and valuable input.

ANTHONY TA
IS Program Director &
Emerging Tech Leader
Societe Generale

JOSEPH SABELJA
Former Executive Director -
AI/ML Product Strategy
JP Morgan

JANE GAVRONSKY
Chief Operating Officer
FINOS Foundation

KATHRYN ZHAO
Former Head of Electronic Trading
Cantor Fitzgerald

PETER JACKSON
Global Head of Data Office
Schroders
Director, Chief Data and Analytics Officer
Carruthers & Jackson

JULIA BARDMESSER
Adjunct Professor
NYU Stern School of Business
Founder and CEO
Data4Real, LLC

AKHIL KHUNGER
VP Quantitative Analytics
Barclays

ARUN MAHESHWARI
Executive Director - Head of Model Risk control, Legal and Compliance
Morgan Stanley

KRISHAN SHARMA
Senior Vice President - Risk & Analytics
Citi

interop.io

io.Intelligence

Bring AI Copilots to User Workflows Today

AI that accesses context. Coordinates across applications.
Takes action inside workflows your users already trust.

CONTEXT

React to live user context in real time across every app

ACTION

Trigger workflows across systems — no custom wiring

TRUST

Embed AI inside the tools your teams already use

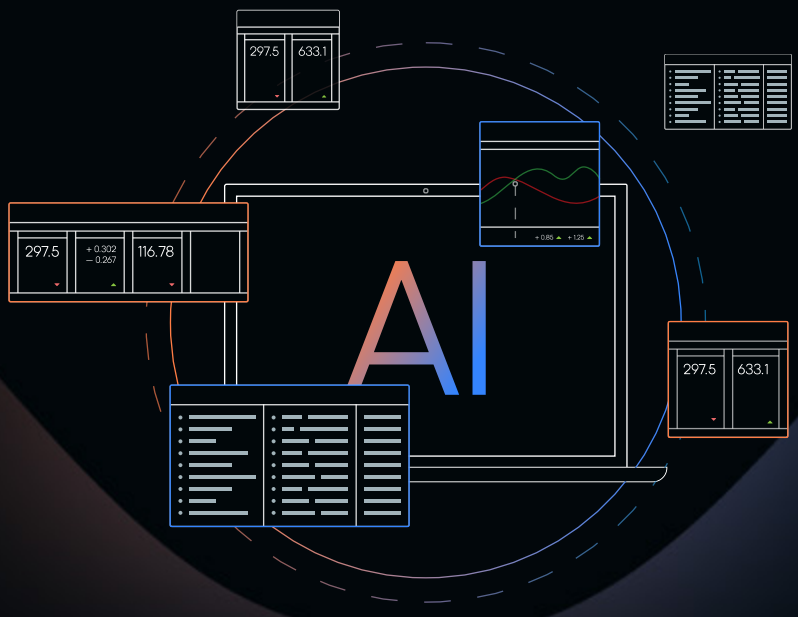

From the industry's trusted interoperability provider.
Your teams are ready for AI. Are your systems?

CONTACT US

WINNER

MOST INNOVATIVE FINANCIAL TECHNOLOGY EXECUTIVE

interop.io

Bob Myers leads product vision and strategy across the interop.io product suite. Over the past year, he has played a central role in shaping the company's AI direction, spearheading the launch of io.Intelligence — an AI framework that helps financial institutions bring AI into the workflows users already trust.

Bob recognized that the same principles behind interoperability — connecting different technologies into one seamless user experience — could provide a practical foundation for enterprise AI. Under his leadership, interop.io has advanced a governed, workflow-based approach that helps firms introduce AI into real-world environments faster, prove value early, and scale with confidence.

With the latest expansion of io.Intelligence, including io.Assist, AI Web, and AI Server, Bob has helped position interop.io at the forefront of practical AI adoption in financial services.

interop.io

interop.io helps financial institutions build modern application platforms without replacing the systems their teams already rely on. Its flagship product, io.Connect, brings together web, desktop, and legacy applications so they can share context, launch actions, and work together as part of a seamless user experience. This gives technology teams a practical foundation for platform evolution, faster delivery of new capabilities, and more effective AI adoption inside governed enterprise environments. Leading financial institutions, including Citi and JPMorgan, work with interop.io to support trading, investment, and advisory operations with greater flexibility and control.

interop.io

Bob Myers, Chief Product Officer & Head of io.Intelligence

BOB MYERS OF INTEROP.IO NAMED MOST INNOVATIVE FINANCIAL TECHNOLOGY EXECUTIVE

Bob Myers, Head of io.Intelligence and Chief Product Officer

Bob Myers, Chief Product Officer of interop.io and Head of io.Intelligence, has been named Most Innovative Financial Technology Executive in the A-Team Innovation Awards. The recognition reflects a body of work centred on making AI practical and governable inside the applications traders and operations staff already use, rather than alongside them, and on building the interoperability foundations agents will need to operate across those applications safely.

In conversation with TradingTech Insight, Myers sets out the thinking behind the workflow-first approach, what a first production copilot actually looks like on a buy-side or sell-side desk, how the agent-identity problem is being tackled in practice, and why he believes the industry's current tilt towards per-user assistants is the wrong default.

TTI: Your appointment as CPO came alongside continuing to lead io.Intelligence. What does the dual mandate signal about where interop.io is heading, and what does the CPO remit open up that leading io.Intelligence alone wouldn't?
BM: AI is moving faster than any technology shift we've been through before, and most companies are still working out how to absorb that pace organizationally. At interop.io we set up io.Intelligence as a dedicated business unit so we could move quickly on research, protocols, and implementation patterns alongside the business-as-usual roadmap commitments we've already made to clients.

I joined as SVP of Product Management and stood io.Intelligence up shortly after. The critical design point from day one was that io.Intelligence has to feed back into the broader product strategy rather than run on its own track. Two failure modes we wanted to avoid: being unable to innovate because of BAU commitments, and creating divergent strategies between the AI work and the core product line.

The CPO remit is what closes that loop. It's about making sure we don't end up with two competing strategies, and that everything we're learning on the io.Intelligence side becomes part of the DNA of every product we ship. My honest view is that over time, the idea of a separate io.Intelligence business unit should become unnecessary. AI is going to be part of everything we do, and the organizational structure should eventually reflect that.

TTI: The workflow-first approach is a clear principle, but many firms have already invested in standalone copilots and chatbots. Where does that legacy tend to get in the way, and how do you guide clients who've started from a different place?
BM: There's no one-size-fits-all approach to AI, and I'd actually argue it's sensible to start with a copilot or assistant experience. It gets you familiar with the capabilities and forces you to work through the nuances like cost and how to handle probabilistic versus deterministic systems.

The Menlo Ventures late-2025 State of Enterprise AI report put roughly 86% of enterprise AI spend on chat-based copilots, so that's where the industry actually is. MCP has made those chatbots much more capable, especially for pulling data from multiple sources. But interop.io's focus has always been the enterprise desktop and the end user, and that user is moving between 20 or 30 different applications to do their job. Unlocking real value from AI means getting it connected into those existing workflows, not fragmenting it further.

The legacy that tends to get in the way is fragmented AI: an assistant per application, each one disconnected from the next. Vendor solutions are the hardest case, because vendors are inherently motivated to deliver their own independent AI experience inside their product. Your users then end up with 30 applications and 30 different AI experiences, which is a bad outcome for a lot of reasons.

For clients who started from a different place, the guidance is roughly the same as modernizing any desktop experience. Catalog what's working today, figure out where your users actually get value, and think hard about the UX from the end-user perspective. What's going to change their ability to do work more efficiently or more accurately? Start from there rather than from whichever AI tool showed up first.

Workflow-first thinking also sets you up for what's coming with agents. Once you've invested in understanding your workflows and connecting AI into them at the desktop layer, you have the foundation agents actually need to operate across applications rather than trapped inside any single one.

TTI: The io.Intelligence materials describe copilots going into production in weeks rather than months. What does a first production copilot typically look like at a buy-side or sell-side firm, and where does the timeline most often get stretched?

BM: The first piece of advice is to trim the scope. AI is exciting and you can build neat things quickly, but there's a big difference between demoware and a system you're comfortable rolling into production. By trimming the scope, you're also trimming what you need to get right in the first release.

You learn a lot on that first rollout. Some of those lessons are already known from good software development and deployment practice: you need observability from day one. This telemetry should cover LLM interactions and token usage, the tools the copilot is invoking, and any failed states users might hit. None of that should be bolted on later.

The scope trim also applies to what the copilot does. I wouldn't start with the most general-purpose assistant that could answer anything. Go deeper on a specific task. A trading example: in a lot of firms, forming the ticket entry in an OMS for fixed income securities is painful, and today that often lives in an Excel spreadsheet with a VBA macro. It's very doable to build a copilot with tools that let a trader describe the trade they want in natural language, convert it into the correct syntax, and open the OMS or trade entry system with the ticket pre-populated. The copilot isn't making the trade. It's removing friction from a specific, well-scoped task. That's the kind of first rollout I'd target.

On where the timeline stretches: a lot of it is the unglamorous project management of getting system access in place. Depending on the maturity of the organization,

LLMs may not even be provisioned yet through something like AWS Bedrock or Azure Foundry. That's normal plumbing, but it's plumbing that adds time.

The other piece is testing. Anyone who has spent real time with AI knows how quickly you can put these experiences into a bad state. When rollouts slow down, it's usually because people are genuinely worried about the quality of what users are about to see, and that's a reasonable place to spend more time. Our answer there is almost always the same: reduce the scope further, add guardrails, and limit the tools the assistant has access to.

TTI: io.Intelligence is designed so copilots operate within existing user permissions. In a bank with entitlements fragmented across OMS, EMS, market data and CRM, how is the agent-identity question being tackled in practice — and where does the line sit between what interop.io handles and what the client owns?

BM: This is an emerging discipline, and we've simplified it deliberately. The agent's identity is essentially a shadow of the user it's working with, and more specifically a subset of that user's permissions.

In the interop.io model, the desktop is the initial sandbox. The assistant can at most do what the user is already permitted to do in a given application, and usually a lot less. The architects of the assistant experience should embrace least privilege at every step, within reason. I don't think it's realistic to require user approval for every single tool call the AI makes. You want a well-defined set of actions the agent can take, scoped so it's not over-permissioned for things it doesn't need to do, for both cost and effectiveness reasons.

A concrete example: Take a portfolio analytics application. As a user, you only see the clients you're entitled to see. That's typically not the whole book. When an agent takes action inside that same application, its view is bounded by your entitlements. If you ask about a client's portfolio, you can be confident there's no data leakage, because the agent's identity is tied to your account in that front-end application. Hallucinations and other issues are still things to protect against, but the data-access side is handled by the entitlements the application already enforces.

On where the line sits between what interop.io handles and what the client owns: our focus is the front-end enterprise desktop. We're not wrapping backend APIs or providing data warehouse access. The client owns the entitlement model in their applications, their LLM provisioning, and the governance and audit infrastructure around their deployment. What we handle is the interoperability fabric that lets the agent see the applications the user has open, understand what each one can do, and invoke the right tool inside those applications under the user's identity. Combining the enterprise-grade permissions the user already has with tight scoping of what the assistant can do is what gives us confidence that the co-pilot pattern can be deployed safely.

TTI: You've spoken about the need to avoid repeating the "sins of the past" with MCP. Which specific mistakes from earlier integration waves are you most concerned the industry might repeat?

BM: Honestly, MCP is a pretty good counter-example to the sins of the past. In earlier integration waves, the pattern was that everyone built their own thing first, and then the industry tried to standardize after the fact, which is the hard way

to do it. With MCP, the industry collectively recognized early that a standards-driven ecosystem was going to be critical for AI, and most of the serious players have embraced that. I was at an MCP Dev Summit recently and the community is genuinely collaborating well.

That said, it's moving very quickly, and the risk I'd watch for is vendors eating themselves trying to carve out a niche on top of the protocol. Parallel protocols that fragment the ecosystem, proprietary extensions that undermine portability, or vendor-specific flavors of MCP that effectively re-create the fragmentation problem the standard was meant to solve. That's the specific sin I'd hate to see repeated.

Our approach is simple: keep supporting open-source standards, contribute to them, and add value through expertise in applying those standards to capital markets specifically. Alongside the protocol work, there's also a lot of rapid innovation happening around the practitioner-level challenges of using AI: accuracy, wrangling probabilistic outcomes users can't act on, and the guardrails every serious deployment needs. That's where a lot of the good work is going to happen over the next twelve months.

TTI: Every trading-tech vendor, LLM provider and SI now has an agentic AI story. How do you frame the case for anchoring an AI strategy on the interoperability layer when a bank is weighing that against its OMS vendor's or cloud provider's framework?

BM: The good news here is that we've already seen this play out practically. Clients who invested in the interoperability layer before AI arrived are the ones moving fastest now.

A quick definition is useful. The interoperability layer is a metadata-rich framework that defines which applications exist on the desktop, what each one is capable of doing, and what data they handle. On top of that sits the messaging fabric between applications, and the visual layer where users, and now agents, can participate alongside those applications. Done right, it's the underlying context model for the user's working environment.

What we've seen is that clients with that foundation in place can turn their interop-enabled desktop into an AI-centric desktop in weeks. You introduce an assistant alongside 30-plus applications, and the AI knows what those applications are, what they do, which tools it can invoke against each one, and that happens in a location that's already governed and observable. They didn't need to build 30 different AI experiences, one per app. They got something closer to a universal remote control for their entire desktop.

That's only going to matter more. As writing code gets cheaper, applications proliferate, and increasingly so will agents. You need a fabric that lets AI understand the context the user is working in quickly and act in it safely.

OMS vendors and cloud providers have strong frameworks at their own layer, but they're still disconnected from the last mile, which is the enterprise desktop the end user actually works in. That space is changing quickly. There are genuinely powerful concepts coming online from operating-system-level providers around co-working and assistants, and we'll be ready to participate alongside them. But enterprise work has always been done on the edges, across custom systems and complex

workflows, and that isn't changing. Anyone who invested in interoperability should feel pretty smart right now. Anyone who hasn't, this is exactly the kind of horizontal investment that positions you for the next two, three, four waves of AI capability, with stable deterministic patterns backing up the AI experience.

TTI: Agentic workflows crossing application boundaries raise harder questions about accountability when something goes wrong. Are regulators starting to engage on desktop AI orchestration, and how do you think about auditability holding up to external scrutiny rather than just internal compliance?
BM: Whether it's a regulator or your internal InfoSec team asking the questions, observability has to be the first pillar of any AI strategy. That doesn't fully remove the black-box feeling some AI decisions will always carry, but everything the assistant or agent does needs to be logged, captured, and available for scrutiny. You may end up using AI to help scrutinize it, because the volume of telemetry is going to be significant.

On the specific regulator question: we're not seeing heavy direct engagement on desktop AI orchestration yet. That will come, but the preparation is the same regardless. You need to govern these systems, understand where anomalies are happening, and be able to answer questions about bad actors or bad outcomes with real data.

When something does go wrong, the agent identity question becomes the accountability question. Whatever pattern you're running, multi-agent, sub-agents, whatever, the single most powerful thing you can do is associate a person or a team with ownership of an agent. For assistants operating on behalf of a single user, that user owns the actions the AI takes for them. For shared agents, say one that helps form trade tickets for a desk, a team has to own that agent. When it misbehaves, they're accountable for fixing it. That's not about blame. It's that outcomes improve when ownership is clear, and when we're not shipping and forgetting about tools this capable.

The other piece under auditability is being honest about use cases. There are plenty of AI deployments that shouldn't be AI deployments. Cheap, deterministic scripting patterns existed before any of this and are still the right answer for a lot of work. I'm happy to see AI go first through the door in a new domain, because it lifts value quickly, but you should audit those deployments and be willing to cut use cases as much as expand them. Sometimes what you actually needed was a SQL query on a schedule, and the AI's role is to analyze the result rather than author the query.

This space will evolve, and we're very early. There's going to be both headline risk and practical risk with AI. The foundation worth building now is telemetry plus the organizational structure that maps people and teams to agents. Get those right and you have something that can hold up to external scrutiny, not just pass internal compliance.

TTI: Where do you expect interop.io's clients to sit on the copilot-to-agent spectrum by the end of 2026 — and is there an assumption about agentic AI in capital markets that you think the industry currently has wrong?
BM: Copilots aren't going away. The power of a context-aware assistant that effectively acts as a digital twin for your work only increases the more you invest in

it. That virtuous cycle is real. But by the end of 2026, I expect the balance to have tipped agent-forward.

Here's the example I keep coming back to. It feels great to roll out an assistant to 100 users across a business unit, but if all 100 are asking similar questions over the course of a day, you've deployed AI in a pretty inefficient way. Take overnight variance in accounts or market conditions. If every user walks in the next morning and asks their assistant to research and analyze what happened, you'll get 100 slightly different answers, spend a lot of tokens, and leave a lot of consistency on the table. That's where ambient agents come in. An agent runs once overnight, produces a single work product, and hundreds of people consume it. If someone wants to go deeper or ask questions, the assistant is right there. The baseline research doesn't need to be re-done for every user.

That pattern is where I see the industry heading. A lot of enterprise work really can be done asynchronously by agents on behalf of many users. Given the cost and accuracy dynamics of AI, that pattern may actually be required at scale, not optional.

The assumption I think the industry has wrong is that on-the-fly, per-user assistants are the default shape of enterprise AI. They're one shape. They're not the cheapest, the most consistent, or the most auditable shape for a lot of the work that needs to be done. Scheduled agentic runs that build once and serve many deserve a bigger share of the roadmap than they're currently getting.

On the interop.io side, the playbook is straightforward. Stay in market. That means showing up in the working groups at forums like the Agentic AI Foundation and at the MCP conferences, but more importantly staying close to practitioners. We'll continue partnering with clients on their AI journeys and building the framework capabilities they need to accelerate. This is a genuinely exciting time, and we'll look back on it as formative. The job is to stay innovative and flexible, meet the moment, and not compromise on the things that make an enterprise vendor worth trusting: security, governance, supportability, and maintainability of the solutions, backed by an effective team of people to work alongside our technology.

Two signals for 2026

What emerges across the conversation is a pair of arguments that deserve wider hearing.

The first is that interoperability, done seriously and done early, has quietly become an AI-readiness story: the clients moving fastest on AI are the ones who built a metadata-rich, governed desktop fabric for reasons that pre-dated the current wave.

The second is Myers' contrarian call on where AI spend should sit. If scheduled, ambient agents producing a single work product for many users are the right shape for a meaningful slice of enterprise work, as he argues, the industry's current tilt towards per-user assistants is skewed – a fitting note on which to mark the award.

Power smarter FX trading with trusted intelligence, when and where you need it

Take advantage of leading news coverage, commentary, analytics and a trusted data universe for smarter FX trading.

Discover more

LSEG FX

WINNER

MOST INNOVATIVE AI IN TRADING INITIATIVE

LSEG FX

Your complete FX connection

Only LSEG FX connects global FX markets with comprehensive, end-to-end solutions across the FX trading value chain.

We've been central to the FX community - driving market progress for over 50 years.

With our deep liquidity, global reach and scale, we're the essential partner for FX.

- Power smarter FX trading with trusted intelligence, when and where you need it. Take advantage of leading news coverage, commentary, analytics and a trusted FX data universe for smarter FX trading.
- Get more time to trade, with every step of the lifecycle covered. Simplify FX operations with connected, end-to-end solutions across the FX value chain, all from one trusted partner.
- Your gateway to global FX liquidity. Execute with greater clarity and control in global markets. Seize opportunity with access to deep liquidity.
- Unlock the opportunity in FX post trade. Increase efficiencies across the FX trade lifecycle and automate your monitoring, reporting and transparency requirements.

www.lseg.com/en/fx

LSEG FX

Power your FX trading with trusted data and analytics, deep liquidity, seamless execution, and unmatched operational and resource efficiency. All from one trusted partner with over 50 years at the forefront of the global FX market.
Only LSEG FX delivers end-to-end solutions across the global FX value chain, helping you trade smarter and more safely. We're the essential partner that connects the FX community.

www.lseg.com/en/fx

Alex Goraieb, Head of FX Pre- and Post-Trade Workspace Workflows, LSEG FX

HOW LSEG AND TRADEFEEDR ARE MAKING FX EXECUTION ANALYTICS A CONVERSATION

Alex Goraieb, Head of FX Pre- and Post-Trade Workspace Workflows, LSEG FX

When LSEG and Tradefeedr began talking about a partnership in 2023, neither side was primarily thinking about artificial intelligence. The conversation was about FX trade analytics - filling a gap in workflow, automating trade panel selection, democratising access to the kind of transaction cost analytics that had historically required a quant team and a working knowledge of Python. Then several things happened at once. LSEG launched Workspace as a single umbrella for its data, analytics and trading workflows. Its strategic partnership with Microsoft matured. And the capabilities of large language models moved from novelty to practical utility inside the enterprise.

The partnership that emerged - now recognised with the A-Team Innovation Award for Most Innovative AI in Trading Initiative - sits at the intersection of all three. Tradefeedr's unified data layer, a normalised API-delivered view of FX execution data covering more than a hundred institutional participants, has been integrated into Workspace and surfaced through Microsoft Teams. The result is a workflow in which a buy-side trader, a corporate treasurer or a bank sales desk can query normalised TCA data in natural language, generate on-demand reports without writing code, and share the output directly with a counterparty inside the chat tool both sides are already using.

In conversation with TradingTech Insight, Alex Goraieb, Head of FX Pre- and Post-Trade Workspace Workflows at LSEG FX, and Balraj Bassi, CEO of Tradefeedr, walk through what the partnership actually delivers today, how the architecture fits together, and why they believe this is only the first chapter of a multi-year, multi-asset transformation.

TTI: LSEG and Tradefeedr announced the strategic partnership in early 2024, with the first Workspace integration following later that year. What was the original vision when the two teams first started talking, and how has the scope of what you're building together evolved since?

AG: LSEG has always aimed to deliver a seamless end-to-end trading experience to the FX community, across manual and API channels. Where we deliver best in class ourselves, we do; where we don't, we partner. We operate venue- and counterparty-agnostically, and we were looking for a trade analytics provider that would complement that approach. Tradefeedr fitted exactly.

In the background, we had just launched Workspace, whose vision was to bring everything together under one umbrella. The Microsoft partnership then turbocharged things - opening up channels for collaborative distribution and, through Teams and Microsoft's language models, making it far easier for customers to engage with the analytics as part of their workflow. Everything came together at once.

BB: TradeFeedr is an independent, open-banking-inspired data layer. We standardise TCA calculations across providers and give clients a single delivery mechanism - one data agreement and, wherever they trade, the data is delivered. Clients are pushing for more integrated workflows: they want rich data insights landing directly in their OMS and desk tooling. LSEG has the largest installed FX client base; we have a product that can normalise data for everybody. Combine those, and you give clients enterprise data access, workflows and AI, all in one place.

TTI: Can you talk about how LSEG's assets - Workspace, FXall, the FX reference data, the Teams surface and the global client footprint - come together with Tradefeedr's analytics network and APIs to create something neither firm could deliver on its own?

AG: Let's start with market reference data. We run a trading platform receiving streams from more than 40 liquidity providers, giving us a robust benchmark. Then there is FXall, the biggest bank-to-customer network. Putting the LSEG brand behind Tradefeedr as our preferred provider, and taking them to our 2,000-plus FXall buy-side takers, signalled to the banks that they needed to come with us. Global distribution matters too - we can reach regional and domestic banks that a smaller firm cannot easily reach. And analytics need to be part of the workflow. Dovetailing them at a critical decision-making point is where the sum becomes worth more than the parts.

BB: What TradeFeedr brings is a market-trusted, unified API that is connected everywhere - our advisory group tells you everything you need to know about market acceptance of our calculations. It is an API that can supercharge the already massive oil tanker that is LSEG. If a client is on Workspace but trades elsewhere, we bring that data in, in a trusted, normalised way. Sitting alongside it is the data lake containing all the pre-trade models. It's the combination of data access, pre-trade analytics and the backing of the community.

TTI: Your award submission highlights the ability to query analytics directly from Teams as central to the solution. What was the thinking behind making Teams the access point, and how does it change the workflow of pre-trade decision-making and post-trade review?

AG: Analytics have traditionally been either an application you query or something you code in Python. Applications are convenient but rigid; coding is flexible but requires quants. LLMs break the trade-off - you simply ask. The question is then why Teams? Behind every LLM query is an intent and a desired outcome, and the next step is almost always either to execute or to reach out to the counterparty. If you have an environment that brings all your counterparties together with a shared view of analytics, you are one click from the conversation that matters. Teams is already on everybody's desk. You are working in the customer's environment.

BB: Every large institution has Teams installed - it is the de facto standard. What the Teams experience does is two things. First, it lets LSEG federate clients across institutions, creating a community from existing Teams deployments rather than asking anyone to adopt a new chat platform. Second, it brings data into each client's own Teams instance, inside their own trust boundary, via AI and LLMs. That fundamentally reimagines how market participants talk to one another.

TTI: You've described a workflow that once took hours across multiple apps and stakeholder groups now becoming instantaneous. Could you walk our readers through a concrete example of that transformation?

BB: Take oversight and compliance on execution reporting on a buy-side desk. Previously, you'd pull reports from banks and data from platforms, try to piece it together, and have ad hoc conversations internally and with the banks. To do it systematically you needed to buy a system or have Python skills in-house. APIs and LLMs change both sides of the stack. On the ingest side, you get normalised, standardised data across wherever you trade, through a single API. On the consumption side, instead of dashboards and PDFs, you have a chat interface where you can query the data and encode your business practices into agents. It is the biggest change I have seen since voice to electronic trading. Collaboration starts from the fact that banks and clients can now see the same bilateral data.

AG: Before, sharing the output with a counterparty meant organising a meeting and sending reports. Now it is one query and a click on a Meet Now button. And it is highly customised: you might be mid-conversation and want to change the question - "show me my spread, reject rate and cost-signal ratio over the last two months" - and it doesn't matter whether a pre-built report exists. A couple of months ago we were on a call with a senior salesperson at one of the biggest global liquidity providers. He looked at the workflow and said: "This is a game changer. I can generate reports mid-conversation with the customer without writing a line of code and without relying on my quants." That last point is the key one.

TTI: How is the ability to share execution analytics with counterparties directly through Teams changing the quality of the conversation between liquidity consumers and liquidity providers?

BB: It changes everything. The old way was: download a PDF, schedule a meeting, go through it with a pencil. Now you can call on the analytics on demand, so the frequency goes up. You can record decisions, reference them later, and build an electronic audit trail that LLMs themselves can work with. It is going to digitise pretty much the whole interaction.

AG: Transparency is at the core - an open, data-supported conversation. This is not "maybe"; it is fact, in real time. The metrics are easy to generate and share, and the outcome should be better execution on both sides: better service for the customer, a bigger share of wallet for the bank. Better decisions, quicker engagement, fewer delays.

TTI: Can we dig into the AI architecture? Balraj, you have talked to TradingTech Insight before about Tradefeedr's LLM-agnostic approach and the orchestration layer that sits between the language model and the trading data. How does that architecture work in the LSEG context, and what does it mean for clients in terms of flexibility, accuracy and governance?

AG: We each have our own AI strategy, and the two are meeting. The driver is the same - use AI and LLMs to augment the trader or salesperson. But LSEG is a conservative institution. We can't build without care, because our customers need to trust us. We have a strong AI Governance Forum covering security, accuracy, reliability and compliance. Tradefeedr can deliver solutions directly, and does; and Tradefeedr can deliver solutions with us. For more conservative customers, we provide the same services with an LSEG filter applied.

Architecturally, our implementation has been rubber-stamped against LSEG's responsible AI principles, which align with the NIST AI Risk Management Framework, the Microsoft responsible AI standards and the relevant ISO guidance. Our data, analytics and trading workflows are delivered through an MCP connectivity framework - the same approach Balraj uses - alongside third parties. The aspiration is that whether a client wants data, analytics, deep research, trade analytics or a combination, and whether they want it as a human or an agent, we deliver it through a single interface.

BB: An MCP server isn't AI - it enables someone's AI strategy. It takes a powerful API and makes it consumable by AI services. We have no intention of building our own LLMs; every large institution may have its own internal strategy, so we hand them the MCP server. LSEG goes a step further - with its Microsoft partnership and scale, there are plenty of customers who don't want to manage their own LLMs at all, and for them using LSEG for the whole stack makes sense.

TTI: Without naming names where that's sensitive, what can you share about how the solution is being used in the market today, and what clients are telling you they can now do that they couldn't before?
AG: Hundreds of customers are using the analytics. The move now is from analytics to Teams. The eye-raiser is how easy it is to perform extremely complex analytics simply by asking. One large UK asset manager told us it was easier to outsource to us than to implement internally - they had asked their IT department for an LLM capability and, by the time it went live, it was out of date. Internal teams tend to be too ring-fenced; the quickest way up the AI adoption curve is to outsource. And it isn't only asset managers - regional banks are using it, a couple of large corporates are using it. The big banks tend to prefer their own internal tooling, but as soon as the communication dimension comes in, they still need a common collaboration platform.

TTI: How do you see the LSEG/Tradefeedr partnership evolving from here - both in deepening what's possible in FX and in extending the model to other asset classes? What's on the roadmap that you're most excited about?
AG: When we started, we weren't even talking about AI - it was about FX trade analytics and filling a gap in the workflow. Then AI arrived, Microsoft came in alongside, and we have almost changed tack from where we began. It is AI-first now. You are not going to keep building a new app every time someone wants a new field; for report writing, intelligence and agentic workflows, it is AI-first.

The next direction is the move from FXall to our other platforms, including advanced dealing, liquidity aggregation and advanced trade execution. Then comes asset class expansion. Fixed income is next: expect an announcement very imminently on an expansion into fixed income with Tradefeedr analytics. It is no secret that LSEG has a considerable stake in Tradeweb. The approach is asset-, venue- and counterparty-agnostic - trade analytics through our AI agent, via Workspace or Copilot within Teams.

BB: What clients want is simple: automated, trusted, stable workflows. Trades come in, pre-trade insights are put into the blotter, trades are executed, exceptions are managed. On top of that, traditional dashboard reporting will become agentic, because every client wants something slightly different, and the best way to deliver

that at scale is to encode business logic in LLM calls. Given the size of LSEG's client base, we are at the beginning of a five-year cycle, not the middle.

Fixed income is particularly interesting because it is even more fragmented than FX, and that is precisely why AI is the right way to approach it - normalise the data through APIs, then let clients interact with it through the AI lens. You no longer need five years of Python programming experience to write customised queries on vast amounts of data. If you get all three things - capabilities, trust and scale - you will meaningfully change the market.

The longer view

What makes the LSEG/Tradefeedr partnership a deserving winner of this year's award is not just the FX trade analytics it delivers, impressive though that is. It is the shape of the workflow the two firms have built around it. Natural language access to normalised execution data, inside a chat environment both buy-side and sell-side already use, with counterparties federated into shared conversations and an electronic audit trail running underneath - this is the kind of infrastructure shift that tends only to become visible in hindsight. Goraieb's point that neither side was talking about AI when the partnership began, and Bassi's insistence that the market is at the beginning of a five-year cycle rather than the middle of one, reinforce the same message: the pieces are only now starting to click into place.

The imminent fixed income expansion - backed by Tradefeedr's recently announced equities and futures partnership with BMLL, and by LSEG's considerable stake in Tradeweb - suggests this is a template that both firms believe travels. Whether the conversational, agentic model of execution analytics proves as natural a fit for the more fragmented world of fixed income as it has for FX is a real question, and one worth watching closely. But the direction of travel is now unmistakable. The TCA dashboard and the PDF report, for so long the default outputs of execution analytics, are starting to look like artefacts of an earlier era.

WINNER

BEST AI FOR OPERATIONAL RESILIENCE

Smartstream

Smart Agents is Smartstream's agentic AI solution for the autonomous back-office. It plans, reasons and executes across end-to-end workflows, serving data directly to users instead of forcing them to chase it across disconnected systems. Built on decades of reconciliation domain expertise, Smart Agents transforms exception-heavy processes into intelligent, self-optimising operations. The impact is transformational. Today, up to 70 percent of back-office effort is spent chasing exceptions, held back by fragmented systems, manual investigations and endless email threads. In pilot deployments, Smart Agents has cut investigation time per user by 70 percent, reducing exception handling from 14 minutes of manual work to just 30 seconds, with 30 to 60 percent faster resolution and 20 to 40 percent fewer escalations. In one benchmark, 500 exceptions that typically consume 116 hours of team effort were resolved in a few hours under fully autonomous operations. Smart Agents interpret exceptions, retrieve and orchestrate data, execute end to end, and collaborate with humans only where governance or judgment is required. Every action is fully logged for auditability. Already deployed across bank reconciliations, cash breaks, settlement exceptions and AML/KYC investigations, Smart Agents integrate natively with the Smart solution suite and third-party systems.

Learn more **https://smart.stream/solutions/smart-agents/**

smartstream

Smartstream unlocks trusted, intelligent data insights – clearly, consistently, and without compromise.
Empowering leading global financial institutions and enterprises with innovative solutions that deliver accurate, timely data insights to streamline operations, reduce costs, and meet regulatory demands with confidence.
By harnessing data at the core of customer operations, Smartstream drives smarter, faster outcomes across reconciliations, liquidity, collateral, corporate actions, fees, and reference data solutions.

smart.stream

Yogesh Shenai, Senior Product Manager

SMARTSTREAM CELEBRATES WINNING A-TEAM INNOVATION AWARDS' BEST AI FOR OPERATIONAL RESILIENCE PRIZE

Yogesh Shenai, Senior Product Manager

Smartstream has been named Best AI for Operational Resilience provider in the A-Team Innovation Awards 2026.

The the data solutions specialist won the prestigious award for its Smart Agents for Investigations innovation, which introduces an AIdriven approach to exception handling in reconciliations and backoffice operations.

Data Management Insight spoke to Yogesh Shenai, Senior Product Manager at Smartstream about the product and the company's plans.

Data Management Insight: Hello Yogesh. Congratulation on your success. What does winning this mean to Smartstream?
Yogesh Shenai: Winning the A-Team Innovation Award 2026 for Best AI for Operational Resilience is a proud moment for Smartstream and a strong endorsement of the work our teams have put into both Smart Agents and Air. It recognises that we are not just talking about AI in financial operations, we are delivering measurable outcomes that change how banks run their back office.

Smartstream has spent decades building deep domain expertise in reconciliations, exception management and data automation. Air, our cloud-native data automation and intelligence platform, was already showing clients what is possible when AI and machine learning are applied to data ingestion, cleansing, enrichment and matching. Smart Agents takes that foundation a step further by replacing manual investigation with autonomous, intelligent workflows. Together, they form the backbone of a more resilient, scalable, and intelligent operating model.

The award validates the direction we have taken. It tells our customers, our people and the wider market that the autonomous back office is no longer a concept, it is happening now, and Smartstream is leading the way.

DMI: How do Smart Agents for Investigations work and how can clients benefit from using them?
YS: Smart Agents for Investigations follow a structured execution model that mirrors how an experienced analyst works. They interpret and classify each exception, retrieve and validate data across multiple sources including databases, client applications, market data and industry services, then autonomously execute the resolution actions across connected systems. Every step is logged, so there is full auditability, and a human is brought in only where governance or judgement is required.

The benefit to clients is straightforward. Pilot results show a 70 per cent reduction in investigation time per user, per break, cutting exception handling from 14 minutes manually to around 30 seconds. In one benchmark, manually processing 500 exceptions, which usually takes 116 hours of team effort, was reduced to just a few hours under fully autonomous operations. Clients see 30-to60 per cent faster resolution times, 20-to40 per cent fewer escalations, lower risk, and analysts freed up to focus on oversight and higher-value work.

Smart Agents do not just resolve the break, they clean the data once they find the root cause, so the same exceptions stop coming back.

DMI: How does Smart Agents for Investigations integrate with Smartstream's other services and products?

YS: Smart Agents are built to work natively across the Smart solutions suite, starting with Smart Reconciliations. They sit within the wider Smartstream ecosystem and draw on privileged access to our proprietary product intelligence, private APIs, and decades of reconciliation knowledge. That is what gives them a level of precision and actionability that generic AI tools cannot match.

Air is a significant part of the wider Smartstream picture in its own right. As our data automation and intelligence platform, Air cleanses, enriches and matches the data that flows into back-office operations, improving data quality at source so fewer breaks arise in the first place. As Smart Agents extend further across the Smartstream portfolio, the combination will give clients a fully connected view from data ingestion through to autonomous exception resolution.

DMI: Does this require a rip and replace process or can it be integrated with existing tech?

YS: There is no rip and replace. Smart Agents are designed to deploy into a client's existing operational environment with minimal configuration, using current data flows, reconciliations and system connections as the starting point. They connect into third-party and internal systems using protocols such as MCP, so clients keep what they have and add agentic intelligence on top. This is a deliberate choice. We know financial institutions cannot afford long, disruptive technology programmes. Smart Agents deliver measurable automation benefits from day one, without the lengthy setup. The same is true of Air, which uses a no-code integration layer to connect to over 300 systems including Oracle, Salesforce and Amazon S3, with no need to disrupt existing workflows or rely on IT.

DMI: How do the agents perform during extreme market volatility?

YS: Volatility is exactly where Smart Agents prove their value. The platform is built with elastic, on-demand scalability, so it can handle sudden spikes in volume without adding headcount or operational overhead. Multi-agent orchestration coordinates triage, investigation, resolution, alerts and analytics across a network of specialised agents, working like a co-ordinated workforce that scales up and down as the workload demands.

In practical terms, this means clients can rely on Smart Agents to keep pace when markets move sharply, when settlement volumes surge, or when corporate actions create a backlog. The autonomous workflow continues to plan, decide

and execute, while human teams stay focused on oversight and the moments that genuinely require judgement. That is what operational resilience looks like in practice.

DMI: How does the system distinguish between a legitimate new resolution pattern and a one-off manual error by an analyst?
YS: This is where domain expertise really matters. Smart Agents do not learn from a single action in isolation. They learn from patterns that are validated against the operational playbooks Smartstream has built up over decades, alongside a client's own standard operating procedures and historical outcomes. A new resolution pattern needs to repeat, hold up against the underlying data, and align with the agent's goal-driven behaviour before it is treated as a reusable approach.

Where there is uncertainty, the maker-checker workflow and human-in-the-loop controls bring an analyst in to confirm. A one-off manual error therefore stays a one-off, it does not become a learned behaviour. Continuous learning loops convert genuine analyst expertise into reusable playbooks that improve automation efficiency over time, while governance keeps the system accurate and trustworthy.

DMI: How does it generate a human-readable audit trail that satisfies internal compliance and external regulatory scrutiny?
YS: Auditability is built into the platform from the ground up, it is not bolted on afterwards. Every action a Smart Agent takes is logged with full explainability. That includes the data it retrieved, the source it came from, the reasoning it applied, the decision it made, and the action it executed. The output is a complete, transparent audit trail that aligns with regulatory expectations and supports internal compliance reviews.

Because the trail is human-readable, compliance teams, internal audit and external regulators can follow each step in plain language, see why a decision was made, and trace the supporting data. Combined with Air's secure environment, which assists with DORA and PCI compliance, clients get a level of governance and transparency that gives them the confidence to move further along the autonomous journey.

WINNER: MOST INNOVATIVE USE OF GENERATIVE AI

Canoe Intelligence

Canoe Intelligence is the platform for smarter alts management. We redefine alternative investment intelligence with AI-driven software that directly addresses the core challenges of private markets. Our technology empowers institutions, LPs, and wealth managers to future-proof their alts infrastructure, modernizing systems and providing a scalable foundation for long-term growth. By automating manual data processing with AI-native precision, Canoe helps clients reduce operational costs and risks, mitigating errors.

canoeintelligence.com

Jason Eiswerth, CEO and team

WINNER: MOST INNOVATIVE USE OF OPEN SOURCE/CLOUD TECHNOLOGIES

Connamara Technologies

Connamara Technologies is the company enabling the most exciting and innovative markets in the world. Its flagship product, EP3®, is a new breed of marketplace technology that seamlessly integrates all exchange, clearing, and market surveillance functions into a single, robust platform. It is adaptable, scalable, and quick-to-market. Engineered for the evolving needs of the next generation of trading venues, EP3 is shaping the future of financial markets.

www.connamara.tech

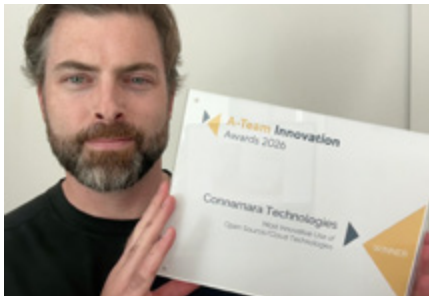

Daniel Davis, Chief Revenue Officer

WINNER: MOST INNOVATIVE CLIENT ONBOARDING AND LIFECYCLE MANAGEMENT SOLUTION

Delta Capita

Delta Capita is a global leader in capital markets consulting, managed services, and technology innovation. Karbon is our CLM technology platform, built on 25+ years of hands-on CLM, KYC, and onboarding expertise. Our practitioners span Advisory, Managed Services, and Technology, working side by side to help firms cut costs, strengthen compliance, and run a smarter client lifecycle. Karbon is designed by the people who run CLM every day, so you can move faster, manage risk with confidence, and deliver better client experiences.

karbon.deltacapita.com

Tracey Allen, Global Head CLM Service Delivery

WINNER: MOST INNOVATIVE DATA QUALITY INITIATIVE

FINBOURNE Technology

FINBOURNE Technology provides cloud-native financial infrastructure for asset managers, asset owners, and asset servicers. Our platform unifies enterprise data management, investment management, and fund servicing operations with complete lineage and bi-temporal accuracy, eliminating the fragmented systems and inflexible legacy architecture that create operational risk and slow teams down. Seamless ecosystem integration means every team works from the same trusted, real-time data. Modernise incrementally, retire technical debt strategically, and make your data intelligence-ready.

finbourne.com

Gareth Evans, Chief Product Officer

WINNER: MOST INNOVATIVE AML / FINANCIAL CRIME COMPLIANCE INITIATIVE

FinScan

FinScan® offers advanced anti-money laundering (AML) and sanctions compliance technology solutions, providing a data-first, risk-based screening approach to ensure unparalleled accuracy and efficiency in identifying and reducing hidden threats, accelerating AML screening workflows, and optimizing team productivity. FinScan's best-of-breed screening platform includes customer onboarding, match-ready sanctions, payments and adverse media screening, and data quality tools, and integrates with leading watchlist and regtech providers to AML compliance end to end.

www.finscan.com

Deborah Overdeput, Chief Operating Officer

WINNER: MOST INNOVATIVE DATA LINEAGE SOLUTION

Rocket Software

Rocket Software is a global technology leader in modernisation and a partner of choice that empowers the world's leading businesses on their modernisation journeys, spanning core systems to the cloud. Trusted by over 12,500 customers and 750 partners, and with more than 3,200 global employees, Rocket Software enables customers to maximise their data, applications, and infrastructure to deliver critical services that power our modern world. Follow Rocket Software on **LinkedIn** and **X**.

www.RocketSoftware.com

Ray Sullivan, Vice President, Product Management

WINNER: MOST INNOVATIVE ENTITY DATA HIERARCHIES APPROACH

One Entity
powered by S&P Global 'Counterparty Manager'

S&P Global
Market Intelligence

At S&P Global Market Intelligence, we understand the importance of accurate, deep and insightful information. Our team of experts delivers unrivaled insights and leading data and technology solutions, partnering with customers to expand their perspective, operate with confidence, and make decisions with conviction. S&P Global Market Intelligence is a division of S&P Global (NYSE: SPGI). S&P Global is the world's foremost provider of credit ratings, benchmarks, analytics and workflow solutions in the global capital, commodity and automotive markets. With every one of our offerings, we help many of the world's leading organizations navigate the economic landscape so they can plan for tomorrow, today.

www.spglobal.com/marketintelligence

WINNER: MOST INNOVATIVE AI IN REGULATORY COMPLIANCE INITIATIVE

Saifr

The Saifr Team celebrates their award

Saifr redefines how compliance operates with advanced AI technology, the right data, and deep industry expertise. Incubated within Fidelity Labs, Saifr harnesses the power of AI agents to help address the limitations and inefficiencies within traditional compliance frameworks—helping firms save time, reduce costs, improve accuracy, and safeguard from regulatory and reputational risks. Saifr's AI-powered risk prevention and management solutions include capabilities for marketing compliance review, adverse media monitoring, and electronic communications surveillance.

saifr.ai

saifr

WINNER: MOST INNOVATIVE KYC INVESTIGATION & DUE DILIGENCE

smartKYC

Hugo Chamberlain, Chief Commercial Officer

smartKYC's technology drives faster, better and more cost-effective KYC at every stage of the relationship – liberating human effort to focus on decision-making rather than laborious research. smartKYC fuses artificial intelligence with linguistic and cultural sensitivity and deep domain knowledge to set new standards for KYC quality, whilst transforming productivity and hardwiring compliance conformance.

smartkyc.com

smartKYC

WINNER: MOST INNOVATIVE REGULATORY REPORTING SOLUTION

TradeHeader

TradeHeader is the definitive authority on financial data standardisation. We are the lead architects who co-create the industry's blueprints. With a 20-year track record of driving the development of global standards, including CDM, FpML, FIX, and ISO 20022 - we operate upstream of the vendor layer. Known as the "Teacher to the Market," we have trained the world's leading Tier-1 banks, central banks, and asset managers to master the data foundations they operate on.

Marc Gratacos - Managing Partner
Anna Muñoz - Senior Consultant & DRR Project Manager
Laura García - Consultant

www.tradeheader.com

WINNER: MOST INNOVATIVE TRADE SURVEILLANCE SOLUTION

Trillium Surveyor

Trillium Surveyor delivers powerful, intuitive trade surveillance and best execution analytics. It's designed to help compliance and trading leaders elevate oversight, meet evolving regulatory expectations, and drive real performance gains. Used by leading financial institutions worldwide, Surveyor turns complex trade data into actionable insights with advanced manipulation detection and execution analysis. With automation, transparency, and speed at its core, Surveyor helps institutions transform compliance into a competitive edge.

David Vins, VP Engineering & Technology;
Lisa Balter Saacks, President;
Melissa Watras, Director of Product

www.TrilliumSurveyor.com

WINNER: MOST INNOVATIVE HOSTED/MANAGED ENTERPRISE DATA MANAGEMENT

viaNexus

viaNexus is a next-generation financial data platform built for the API, AI and agent-driven world. It enables data to be onboarded, normalized, and delivered through a unified, high-performance infrastructure. Operating as a marketplace, a white-label solution for exchanges, and a private enterprise platform, viaNexus supports flexible entitlements and seamless API access. It empowers data providers and consumers to efficiently distribute, monetize, and operationalize financial data at scale.

Tim Baker, Co-Founder & CEO

vianexus.com

A-Team Innovation
Awards 2026

More A-Team Innovation Awards 2026 winners

MOST INNOVATIVE MARKET DATA SOLUTION
Fenics Market Data

MOST INNOVATIVE DATA ORCHESTRATION FOR REGULATORY REPORTING
Accedia

MOST INNOVATIVE ALTERNATIVE DATA SOLUTION
Alkymi

MOST INNOVATIVE USE OF AI IN AN ENTERPRISE DATA MANAGEMENT INITIATIVE
Amberdata

MOST INNOVATIVE SUPPLY CHAIN RISK MANAGEMENT / KNOW YOUR VENDOR INITIATIVE
AML Partners

MOST INNOVATIVE DATA-DRIVEN TRANSFORMATION PROJECT
Andaria

MOST INNOVATIVE TEAM
AQX Technologies

MOST INNOVATIVE DATA ARCHITECTURE PROJECT
Arcesium

MOST INNOVATIVE USE OF DISTRIBUTED LEDGER TECHNOLOGY/ BLOCKCHAIN
Archax

MOST INNOVATIVE CYBER SECURITY INITIATIVE
Blackwired

MOST INNOVATIVE CLOUD-BASED TRADING ANALYTICS INITIATIVE
BMLL

MOST INNOVATIVE IMPLEMENTATION OF KNOWLEDGE GRAPH TECHNOLOGIES
Cognaize

MOST INNOVATIVE REGULATORY COMPLIANCE MONITORING / MANAGEMENT
ComplyCube

MOST INNOVATIVE CLOUD-BASED MARKET DATA DELIVERY
DiffusionData

MOST INNOVATIVE SOLUTION FOR FRONT-OFFICE/TRADING
Exegy

MOST INNOVATIVE ALTERNATIVE DATA SOLUTION FOR TRADING & ANALYTICS
ExtractAlpha

MOST INNOVATIVE ESG REGULATORY REPORTING SOLUTION
Gist Impact

MOST INNOVATIVE SMART TRADER DESKTOP SOLUTION
GlobalLink

MOST INNOVATIVE PROJECT
Harmoney

MOST INNOVATIVE API FRAMEWORK/MICROSERVICES ARCHITECTURE FOR TRADING
Marvelsoft

MOST INNOVATIVE ESG DATA SOLUTION
NeoXam

MOST INNOVATIVE UNSTRUCTURED DATA MANAGEMENT PROJECT
ResponsivAI

MOST INNOVATIVE SANCTIONS AND PEPS DATA SOLUTION
Ripjar

MOST INNOVATIVE FINANCIAL DATA SECURITY SOLUTION
Sedicii

MOST INNOVATIVE DATA GOVERNANCE INITIATIVE
ZLTech

A-TEAMGROUP

For Marketing: **www.a-teamgroup.com**

For Content: **www.a-teaminsight.com**