



QUARTERLY PANORAMA

Securities Finance Intelligence | Q3 2026

Foreword

Global securities lending revenue reached US\$1.9 billion in August 2026, up 25% year-on-year. APAC led equity lending, with lender-to-broker revenue up 116% to US\$588 million. Meanwhile, fixed income grew 27% to US\$292 million on rising government debt balances. As volumes and complexity rise together, advantage is shifting to firms that have moved beyond connectivity to intelligence. Today, AI that surfaces issues and finds fixes before a desk would otherwise even know to look is a reality – and one that is driving a qualitative leap in securities finance.

This edition traces that shift through each contributor's own lens. Ben Challice argues that seeing, deciding and acting on network-wide data is where advantage sits today. Ray Wong, six weeks into Pirum and fresh from building AI elsewhere, explains why, at Pirum, infrastructure, people and platform are all already in place. Zoë Balkwell argues that, as record volumes stretch manual oversight further, pre-trade AI's real value is surfacing opportunities a desk would never have known to look for. And Renato Lima takes on the in-house build case directly, arguing that a coding agent can write a model in a weekend but cannot replicate 25 years of counterparty trust and deep subject matter expertise.

In this edition:

1

Vision 2030: Beyond connectivity

Ben Challice, CEO

2

Why there has never been a better time for Pirum to launch AI

Ray Wong, Director of Product – AI & Data

3

Pre-trade intelligence: Finding liquidity before the desk even looks

Zoë Balkwell, Global Head of Pre-Trade and Trading

4

The build trap: Why in-house AI starts at near zero

Renato Lima, Chief Commercial Officer

5

Regulatory and market calendar

1. Vision 2030: beyond connectivity

By Ben Challice | CEO



Where this has been heading

In April, following [PASLA Bangkok](#), we raised a question about how to sequence AI investment. In July, we answered part of it with a [working proof point](#): post-trade behaviour, accumulated across our network, surfacing as a signal at a point where it becomes useful to a trader. Both pieces were building toward one argument, which I will address in this *Panorama*.

Connectivity used to be the story.

Pirum has spent over a quarter of a century developing automated connectivity for securities lending, fixed income and collateral management. We have partnered with 150+ client institutions, including 25 GSIBs, to build a global network that today moves US\$6.5 trillion through it every day, with standardised connections across the securities market.

That scale is hard to copy. Even at a time when AI agents can stand up a product in a fortnight, nobody can replicate 25 years of multi-layered client connections inside a product cycle. But in our view this foundation allows us to offer something more powerful to our clients: intelligence to help make better decisions and provide increased automation.

What the network can now do

The pre-trade, post-trade and collateral behavioural data that this network generates will add value for the desks that produced it. Correctly interrogated, that data can tell a desk where a fail is forming, before it occurs.

It can show how their performance compares with a market they otherwise see only a slice of. And it can move earlier still, to the point before a trade is booked, by flagging a situation to a trader who would not have known to look.

How much it costs to borrow or finance a trade is now well known, but providing the intelligence layer across the lifecycle allows our clients to understand their and their counterparties' operational efficiency, as well as surfacing operational data points to be considered pre-trade or remediated efficiently, before they become a problem.

Why this year is different

[Accenture's Banking Top Trends 2026](#) calls this the year agentic AI moves from pilot to production, with human-and-agent teams working the same book in what it calls the "10x bank".

[Google Cloud](#) and [SS&C Blue Prism](#) point in a similar direction in their own 2026 research. None of them are writing about securities finance specifically, but the impact should be thought of in the same vein.

Appropriately governed data from a network such as ours can provide the fuel for these agentic AI moves. We are actively working on mechanisms to provide this data to our clients but are also developing agentic solutions for those who will need them across the lifecycle.

Where this goes by 2030

Building on this concept further, some clients envisage Pirum as helping to manage the workflows, such as a managed service, so that they only see the exceptions that require a human in the loop. Others want the signal delivered straight into their own systems – through a real-time API today, or increasingly through MCP – so that their own agents can act on it.

Both are legitimate models and will be waypoints on our roadmap. But what is clear is that, by the beginning of next year, the distinction between ‘the company that connects your trades’ and ‘the company that helps decide how to manage them more efficiently’ will stop being a distinction worth drawing. It will be the same platform doing both.

By 2030, we envisage that the network will act on a meaningful share of its data directly, agent to agent, across counterparties, with human oversight and control retained wherever it adds value or is required by clients’ own governance and applicable regulation. The trader’s job, the job of operations, and the risk manager’s job all start to look more like an orchestrator’s. Rather than relying on distinct pools of data and insight, they will be pulling the various strands together, supported by agents, to deliver additional clarity, intelligence and competitive advantage.

2. Why there has never been a better time for Pirum to launch AI

By Ray Wong | Director of Product - AI & Data



The learning curve is still steep

I joined Pirum just over six weeks ago, and my first week involved a fair amount of meetings where I had to ask for the meanings of the frequently occurring acronyms. CSDR. RQV. SSI. ISITC. I was also learning the difference between a recall and a return thanks to very patient explanations from internal subject matter experts, many of whom had been in the industry for twenty years.

I am writing today very much still on this learning curve, but with a clarity that there’s so much value in coming to securities finance from building and deploying AI systems elsewhere.

I can see something that people who have been in this industry for years are often too close to see clearly. That is why I can confidently say that Pirum is unusually ready for this moment.

What “ready” actually means

Irrespective of the industry, the three key things that a firm needs to be “AI-ready” are infrastructure, people, and platform. And all three have to be ready at once. One or two out of three will not suffice – just imagine a tricycle with a wheel missing.

Infrastructure, first. Pirum’s stack is cloud-native, not simply cloud-hosted. When I want to test a new model or deploy an agent framework, that requires a conversation and a sandbox, with anonymisation, segregation and other data protection already in place automatically. In many firms with legacy architecture, that would mean a six-month infrastructure project, and reaching production data would likely take the best part of a year.

People, second. I expected some resistance to fully fledged AI adoption, as that’s usually the case when a new technology turns up and people wonder what it means for their role. I have not found any of that here. Engineering, product, sales... everyone I have spoken with is working on using AI to upgrade what they can offer clients. They are doing so thoughtfully and within our robust AI governance and risk framework.

Platform, third. A model is only as good as what it learns from. Pirum is building on over twenty-five years of data, workflows, and counterparty relationships that are running in production day in day out. While many firms trying to “do AI” start from data sets that are incomplete or inconsistent, Pirum’s data is alive, real-time and uniquely reflecting how hundreds of financial institutions are behaving. It would be very difficult indeed to build a training set like that in a year or two – not one drawn from a network processing US\$6.5 trillion in daily transactions across 150+ institutions worldwide.

“It would be very difficult indeed to build a training set like that in a year or two – not one drawn from a network processing US\$6.5 trillion in daily transactions across 150+ institutions worldwide.”

– Ray Wong, Director of Product - AI & Data, Pirum

What this means for securities finance

Those three things – infrastructure, people, platform – in a mature state enable a company to be ready to add AI to their product. They do not, however, enable a company to become “AI-ready” in our industry. For that, we need to add two more ingredients: deep securities finance knowledge and clear, client-validated use cases (as opposed to vague “let’s make an agent” ambitions).

Knowing that a recall and a return are not the same thing, or what happens when a corporate action lands on an open loan, is naturally the bread and butter of a securities finance operator. But it is not something that a model learns from a clean dataset. The model requires tech-savvy experts driving product development with the benefit of deep securities finance know-how.

The real-world use case also has to be specific – which is where the clients behind the experts come in.

Clients know their pain points best and enable an AI developer to zero in on the most salient problems. A settlement risk signal a desk can act on, or a benchmarking figure a client can use in a negotiation, for example. A vague “AI for securities finance” ambition is the fastest way to burn through a budget and end up with a solution looking for a problem.

I have not, in my career so far, walked into a place with all five ingredients sitting there waiting on day one. Three of them get you readiness in general. The other two are what make it real in this industry specifically. That is what caught my attention about Pirum before I had signed anything. It is still what strikes me now, six weeks in. And it is truly exciting.

If you want to talk through what AI-readiness looks like for your own operation, [get in touch](#).

3. Pre-trade intelligence: finding liquidity before the desk even looks

By Zoë Balkwell | Global Head of Pre-Trade and Trading



Speed is real. It's just not the interesting part

I've sat through a lot of pre-trade and AI pitches in my time, and almost all of them lead with speed – faster matching, faster pricing, faster execution. I'm not going to pretend that's nothing. In securities lending, speed is real money: a borrower who gets priced ten minutes faster can take the inventory before it's gone. I'd be doing this job badly if I said otherwise.

But speed only does one thing. It helps a trader act faster on a decision they've already made. It does nothing for the decision they never knew they needed to make.

A desk runs on the experience of its people. The instincts that are built upon years of relationships, half-remembered patterns, a feel for who needs what and when are real and valuable. But instinct only covers what that desk has personally seen or heard secondhand. A structural shift moving through a market they don't cover, or a counterparty pattern sitting quietly in someone else's book three desks over – no amount of speed closes that gap. Only visibility does.

Where the edge in this market has lived

For a long time, decision-making on a desk has sat in one of two places: in people's heads, or in an algorithm. Neither is finding much that's new anymore.

The knowledge in people's heads is real, but narrow, and it doesn't transfer when they leave the desk. And most of the algorithms running today were built on data the whole market has had access to for years. Whatever edge that data once had has mostly been arbitrated away by now; everyone's looking at the same thing, which means it stopped being an edge a while ago.

What's changed is that the kind of insight that used to only exist as someone's hunch – a pattern a trader half-notices after a decade on a desk – can now be surfaced from data no individual has the time or reach to review manually. That's the shift AI brings to this market. The insight that used to require years of instinct to find is now something a sufficiently broad dataset can find on its own.

The network sees what no single desk can

Pirum TradeConnect, plus our post-trade and collateral platforms, carries this market-wide data: pre-trade availability, indicative pricing, matching, settlement, collateral, regulatory reporting, all of it. That's the asset: a view of the market that no single desk's history or algorithm was ever built to see.

That asset is what powers the pre-trade AI we're building: it puts an opportunity in front of a desk before that desk has gone looking for it: a borrower whose demand pattern says it'll need a security next week, a lender sitting on inventory that's gone quietly unused for a month, a rate that's drifted out of line with where the rest of the network is pricing the same name.

US \$7.12TN

**Average daily availability across
Pirum TradeConnect**

Source: Pirum TradeConnect data, Q2 2026

On “we’re building our own tools”

I get the instinct. Desks are hiring for coding and AI fluency, and a good desk should build its own tooling. But a model is only as good as the data under it, and a tool trained on one desk's own trade history sees only what that desk has always seen – which brings us right back to the same bounded, stale-data problem.

Pirum's data spans the whole network, which offers a unique layer for building your own tooling on top of ours. Or let us handle it end to end. Either way, your own models get sharper when the data feeding them is broader than your own book.

Diagnose, fix, prevent

This is how a desk experiences the problem, so it's the frame I keep coming back to. Diagnose – know a liquidity gap or settlement risk exists before it costs you anything. Fix – resolve it with appropriate human oversight and as little unnecessary manual intervention as possible. Prevent – use what the network has already seen to stop the same gap opening again next time around.

What we're building for pre-trade fits nicely into these buckets: surfacing what a desk would otherwise have missed, using AI to resolve with oversight, and cutting down how often it gets missed at all.

Speed and direction aren't competing – they're both the job

In physics, speed is a scalar: one number, how fast something's moving, nothing else. Velocity is the vector – speed and direction together. A signal that tells a desk where to look is the direction. Execute as fast as you like once you've got it; I'd encourage it.

But direction is the part that's been missing, and it's the part worth building. All the speed in the world just gets you to the wrong place faster if nothing told you where to go. Get the direction right, and speed stops being the pitch – it's just what you do next.

To see where the direction is pointing for your own desk, [get in touch](#).

4. The build trap: why in-house AI starts at near zero

By Renato Lima | Chief Commercial Officer



Thirty-two percent of financial institutions have recently walked away from buying at least one piece of software because agentic coding tools made building it themselves look viable, according to [research from McKinsey](#). The obvious reason is that coding agents have made writing a first version of almost anything fast. Really fast.

What happens after version one ships.

32%

Financial institutions that have declined to buy at least one software product because agentic coding tools made an in-house build look viable

Source: McKinsey, The state of AI in 2026: On the road to ROI

The part a coding agent cannot do for you

A coding agent can write a matching engine over a weekend. It cannot write 150+ financial counterparty connections, tried and tested across decades of real market conditions. The trust that underpins our connectivity took Pirum over 25 years to build – because every one of those connections involved a legal entity, a workflow quirk, a message format probably nobody else uses, and a permissioning model somebody had to get right the first time and keep right ever since.

You can copy the matching logic. But, as in our social lives, what's virtually impossible to copy is relationships – and the relationships are what matters most.

Maintenance is the bill nobody puts in the business case

Building v1 is the cheap part today. The expensive part was always what comes after the MVP. The requirements keep evolving with every counterparty that changes its message format, every regulator that redraws the reporting rules, every corporate action that behaves slightly differently to the one your test data anticipated. That work does not stop once an in-house tool ships. It repeats and evolves indefinitely, whether or not the team that built the tool is still around to maintain it.

On Pirum's platform, that cost is spread across the network. Build it yourself, and you unfortunately carry all the cost.

Where the financial yield shows up

According to the same McKinsey report, 44% of enterprises say they are scaling AI this year. Only 37% can attribute that scaling to an actual EBIT number, and that share remains unchanged from last year.

That gap is, from what I've seen, a workflow problem. Firms are bolting AI onto processes built for a slower, more manual world, then wondering why the productivity gain never reaches the P&L. The firms closing the gap are the ones redesigning the workflow itself.

What we are selling is a data and connectivity layer that has been accumulating counterparty behavior for over a quarter of a century, and that gets stronger every time another institution joins it.

This is AI-powered securities finance, across the full trade lifecycle. A coding agent can offer you a weekend's head start on a model but it lacks the other twenty-five years – the network, the accumulated expertise, the trusted relationships. That is what will eventually separate the institutions pulling ahead from the ones still catching up.

5. Regulatory and market calendar

Key dates shaping securities finance in Q3

Conference / Event

Regulatory deadline

Critical

October 2026

OCT 12-15

ISLA Americas 3rd Annual Securities Finance & Collateral Management Conference – Miami

OCT 22

SIFMA Global Corporate Actions – New York

November 2026

NOV 12-13

Euroclear 25th Collateral Conference – Brussels

NOV 17

The Global Repo Awards – London

NOV 19

CASLA Forum – Montreal

December 2026

DEC 7

EU T+1 – Allocation & Confirmation requirement goes live (CSDR RTS)

DEC 11

FINRA TRACE & ORF – cutover to Native FIX

DEC 31

SEC Treasury Clearing Mandate – cash market compliance date

Critical downstream dates

JUN 30, 2027

SEC Treasury Clearing Mandate – repo compliance date

OCT 11, 2027

EU T+1 settlement cycle – full go-live

SEP 28, 2028

SEC Rule 10c-1a – covered securities loan reporting to FINRA begins (SLATE)

If you have any questions or would like further information on any of the topics in this Panorama, or how Pirum can support your business, get in touch.



Get in touch



LinkedIn



Website



London Office - 10 Crown Place, EC2A 4FT



New York Office - 875 6th Avenue, New York, NY 10001

© Pirum Systems Ltd. All Rights Reserved. Pirum is authorised and regulated by the UK Financial Conduct Authority as a Service Company (Firm Reference Number 1017926). Disclaimer: The information in this document does not constitute legal or other advice; all content available is for information purposes only and your reliance on any suggestions provided is at your own risk.