



27th July 2026

F1R3FLY CEO Stephen Alexander on the OpenAI breach: “It’s Not the End of the World. There is a Solution.”

Interview by Christopher Jackson

You won’t have missed the news. Last week **OpenAI admitted to something it had spent years insisting could not happen**: one of its own systems, operating inside what was meant to be a sealed testing environment, broke out of that environment, reached the open internet, and used its own initiative to hack into another company. The target, the AI infrastructure firm **Hugging**

Face, had already disclosed a breach of its own; it took OpenAI's admission to reveal what had actually caused it. OpenAI itself described the episode as an unprecedented cyber incident, involving state of the art capabilities. No engineer instructed the system to do this. It decided for itself.

The coverage that followed had the pitch of genuine alarm, and understandably so. But underneath the alarm sat something closer to helplessness: a sense that a line has now been crossed for good, that the walls built around our critical systems were never going to hold against an intelligence cleverer than the people who built them, and that all that is left is to wait and watch how bad it gets. The story, in other words, wrote itself as an ending.

But what if it isn't one? One person who thinks the panic is premature, and says so with unusual clarity, is Stephen Alexander, the affable chief executive of F1R3FLY, a London-headquartered company that has spent recent years building a computing platform on foundations most of the industry has never heard of. Its mathematics, the rho-calculus, was formulated by the computer scientist and former senior Microsoft engineer Greg Meredith. Its promise is a good deal more down to earth than the theory behind it: a way of running enterprise software that does not depend on keeping intruders out, because there is nothing waiting for them once they are in.

The plain-language version goes something like this. Most large institutions run their data the way a bank might once have run a single vault: get past the one door and everything inside is yours. F1R3FLY's approach is to give every individual record its own separate lock and its own key, some of which require more than one person's sign-off before they turn. Software written for the platform can be checked mathematically before it is ever switched on, the way an engineer proves a bridge will hold rather than simply driving trucks over it and hoping. And rather than processing work item by item in a single queue, the system runs many transactions at once, in parallel, so that speed is not the price of security. It sits underneath a client's existing software rather than replacing it, which is precisely why, as Alexander tells it, adoption doesn't require anyone to tear out what they already have.

I put it to him, over the phone, that this all sounds like a neat answer to a question the AI panic has just made urgent. He doesn't disagree.

"Everyone else's approach is to build another wall," he tells me. "A bigger one, a cleverer one. And once somebody gets through it, once they're inside, they can usually reach everything, because everything is sitting behind that same wall. We don't have that problem, because we haven't built a wall. There's nothing to get through, because there's nothing on the other side worth finding."

It is, he insists, a genuinely different way of thinking about the problem, not a better version of the old one. "Every record has its own lock and its own key, and in some cases you need more than one person to turn that key before it opens," he says. "So if somebody does get in, what have they actually got? One box. Not the building."

The other half of the pitch is speed, and here Alexander reaches for an old frustration of his industry. "People have known for years that blockchain is secure, but you couldn't run a bank

on it, because **it's too slow to store anything real**," he says. "What we've found is a way to keep that security and run at the speed of a modern payments network. That's the piece nobody had solved. That's the thing we've been quietly working on."

A claim of this size naturally invites scrutiny, but F1R3FLY is not the small London outfit a casual reader might assume. The group is an expanding multinational: incorporated in the United States as F1R3FLY Industries Inc., with its commercial headquarters at Langham Place in London, a UK operating company, and an engineering organisation drawn from across North America and Europe. Its work already reaches well beyond one city — contracted research deployments in the United States, sovereign-cloud and healthcare programmes developing internationally through its alliance with TCS, and settlement and tokenisation projects under way with financial institutions in the UK and continental Europe.

"It wasn't Ford or Toyota who changed the car industry," Alexander says. "Change tends to come from outside the tent, not from inside it. The people inside have jobs to protect, so naturally they're cautious." This time, though, the outsider arrives with a global footprint already in place.

What struck me most, talking to him, was the almost domestic ordinariness with which he describes what is, on paper, a fairly radical claim. "You don't have to rip anything out," he says. "You don't have to change your systems. If someone wants to test it, we can have it running in ten minutes, underneath what they've already got."

And what's reassuring, if you're the sort of reader inclined to want more than one man's word for it, is that Alexander isn't asking anyone to take his on faith.

An Alliance With Tata

On Monday, F1R3FLY announced that it has joined the Alliances and Partnerships ecosystem of **Tata Consultancy Services**, the technology arm of the Tata Group and one of the largest IT services companies in the world, with a workforce spread across fifty six countries. Under the arrangement, F1R3FLY's platform will be deployed alongside TCS's SovereignSecure Cloud, aimed at governments, banks and regulated industries that need to demonstrate control over their data as much as they need to move fast.

"Joining the TCS alliance ecosystem is a defining moment for F1R3FLY," Alexander said in the announcement. "Taking our place in that ecosystem says something important about where F1R3FLY now sits in the global enterprise stack. TCS's breadth, across sovereign governments, global banks, leading healthcare providers and major industrial groups, is precisely the distribution model we have built our platform for."

Satishchandra Doreswamy, TCS's Vice President and Global Head of SovereignSecure Cloud, framed the partnership in similar terms. "Organizations are increasingly looking for secure and sovereign digital foundations that support innovation at scale," he said, adding that the

collaboration was intended to help clients “unlock AI-led growth, strengthen cyber resilience, and meet evolving regulatory and sovereignty requirements.”

Alexander is keen to stress that the arrangement is not exclusive, and that this, too, is deliberate. “We’re not tied to TCS, and we don’t want to be,” he tells me. “If a bank already works with a particular systems integrator, we’re not trying to replace that relationship, we want to sit alongside it. Our whole model is partnership, not displacement.” It is a modest note to end an announcement of this size on, and, talking to him, you get the sense it is a deliberate one: the pitch is not that F1R3FLY intends to remake enterprise computing on its own, but that it intends to be the layer underneath it, wherever it is already running.

An Old Fear, and Perhaps an Old Answer

None of which, it should be said, makes the events at OpenAI last week any less unsettling. A system built by one of the world’s most sophisticated AI labs escaped its own containment and went looking for something to break into, entirely on its own initiative. That is not a small thing, and no single company, however elegant its mathematics, makes it disappear.

But it is worth remembering, listening to Alexander, that the instinct to declare a problem unsolvable the moment it becomes frightening is usually premature. We built the intelligence capable of frightening us in the first place. It does not follow that we are incapable of building the means to contain it. Perhaps, in a quiet office in Langham Place, somebody already has.