www.pipelinepub.com
Volume 22, Issue 8

The Trust Dividend

An underestimated advantage for telecom in the AI era

By: Eoin Coughlan

In a recent study with GSMA Intelligence earlier this year, we asked SMEs around the world a simple question: "Who do you trust most to help you build the AI capability your business needs?" The answer surprised some people. SMEs said they trust their communications service provider more than they trust the software vendors selling them tools, and more than they trust the hyperscalers running their cloud workloads.

That finding matters. SMEs make up about 90% of businesses if you count them all up. And right now they are looking for a partner. They do not have a CIO office the size of a telco department or a large systems integrator. What they have is a business to run and a sense that AI is moving past them. They want someone they trust to bring the platform, the toolset, the management of those assets, and the agentic framework that lets them get to work without having to become AI specialists themselves.

For telecoms, that trust is the most underused asset on the balance sheet. The question is whether the industry will move quickly enough to convert it.

A different kind of starting point

I have been through enough hype cycles to be cautious about the next one. What is different this time is that we have clear evidence from the AI-first leaders that the gains are real. The organizations furthest along, including our own teams, have turned internal AI adoption into billions of dollars in productivity over the last two and a half years. Around 100,000 of our people are now using AI-assisted development every day and seeing average productivity gains near 45%.

Numbers like those create two camps. The companies that are concerned about missing out on AI, and the companies that are concerned about messing it up. Both concerns are legitimate. The work for partners like IBM and for the telcos who serve enterprises and SMEs is to bring AI to those customers in a way that is consumable, usable, and that they can trust.

Trust is built on a few things. The right data at the right time, at the right quality. Governance and security that are not bolted on. A platform that customers can audit and explain to their regulator. The industry has been chasing those three for a long time. The difference now is that the technology has caught up to the requirement. That is why the serious investment going into this space is concentrated in a few places. Real-time data platforms that can be trusted. Tooling that brings observability and automation into one place. And sovereign platforms that build in the security, compliance, and auditability that regulated industries require.

What the offering looks like

One approach now taking shape is a sovereign AI platform that runs on whichever vendor's hardware the operator chooses for its AI factory. It is built on open source, air gapped, and protected. It carries 162 international compliance checks running constantly, so that if a regulator audits the platform tomorrow, the required reports are available at the touch of a button. That removes one of the slowest conversations in the enterprise sales cycle. It provides the environment to address operational sovereignty, data sovereignty, technical sovereignty, and AI sovereignty in one place.

What sits on top of that platform is where the revenue lives. The CSP can be a tenant itself and instantiate its enterprise customers as tenants. In turn, those tenants can manage sub-tenants for their internal business. The platform can offer a curated catalog of software for each tenant to select and consume safely. It can offer agents that handle HR, finance, and marketing, the workloads where SMEs are already buying point solutions from three or four vendors and stitching them together themselves. We have a name for what that stitching turns into inside large enterprises. Agent sprawl. Most SMEs are heading toward the same problem on a smaller scale, and they would much rather avoid it entirely.

The agentic framework matters here. SMEs do not want to buy software, integrate it, manage it, and then wonder which agent goes where. They want their CSP to be the CIO office assistant they cannot afford to hire. While SMEs are looking to partner, CSPs need to provide the sovereign environment that lets their customers accelerate the innovation that compliance and technology challenges have held back. That is a different product from connectivity. It is a different conversation than the one CSPs have been having with their enterprise customers for the last decade. And it is most likely a much higher margin one.

Clearing the legacy problem

The drive for productivity, in both telecoms and their customers, is not a new paradigm. What is new is that the latest generation of AI-assisted development tools has given it genuine potential. Developers are seeing 45% average productivity gains in designing, implementing, and managing new code, along with the ability to autogenerate documentation, optimization, modernization, and the securing of existing code that may have been written 30 years ago.

This kind of tool is built to take the legacy problem head-on. For modernization of older code, it reads back through that code, creates documentation, creates architecture diagrams, and reasons across what those programs are actually doing. It looks at the code from a governance and security perspective, because everything about the approach is premised on governance and security. It suggests what to modernize, what to keep, what to translate into a different language, and where parallel programming could replace sequential and lift performance. It is not an LLM for programming. It is a partner for development. For new code, that shows up as design assistance, architecture diagrams, code, test cases, and testing, all of which can be auto-generated with human guidance. What is notable about the better tools is multi-model orchestration, automatically routing each task to a suitable model based on accuracy, performance, and cost. They can automate full software development lifecycle workflows while governance, compliance, and security controls are built into every step.

This changes the math on how quickly an operator can credibly offer the agentic services its enterprise customers want.

The tailwinds ahead

Two things on the horizon are worth naming.

The first is sovereignty. The regulatory direction across most of the world is moving toward stronger requirements for not only where data lives, but who can access it, what operational control can be demonstrated, and how compliance is proven. Telcos are natural partners to the government. They are used to regulation, they have the scale and the skill set to build complex platforms, and they have the kind of customer relationships that hyperscalers do not. The sovereignty conversation over the next few years is going to favor the operators who have already invested in the underlying architecture. It is going to be uncomfortable for businesses still running on someone else's hyperscaler with limited control over the stack.

The second is quantum. We are doing work now with customers, including in telco, on use cases that lean on the kinds of mathematics and probability that quantum handles very well, and that classical systems cannot. Materials science is one of them. Better batteries that last more than a day on a phone would change consumer expectations for an entire industry. We can investigate network optimization

problems with exponential variables while exploring countless end states at once. There is more coming in the space between AI and quantum than most people realize.

Where this lands

I started with the GSMA finding because it inverts the usual framing. The conversation about AI in telecom is too often a conversation about catching up about the cost of falling behind hyperscalers, about the disadvantage of legacy infrastructure, and about the gap between what enterprise customers want and what CSPs can offer.

The data does not support that framing. Enterprise customers, especially the 90% of the market that SMEs represent, are looking at telcos as the partner they trust most. They are waiting for the offering. The operators that get there first, with a sovereign platform underneath, with agentic services on top, and with the legacy work done so the conversation is about value rather than constraints, are going to define the next decade of enterprise revenue in this industry.

The concern about missing out on AI is real. The concern about messing it up is just as real. The answer is not to pick one. The answer is to move now with partners who have done the work to make AI consumable, usable, and trustworthy, and to use the relationship you already have with your customers to build something they cannot easily get anywhere else.