

MWC26: Taking Stock of Telco Security

- With evidence of AI risk increasing at both speed and scale, this year's Mobile World Congress saw Deutsche Telekom excel among its European telco peers with a vision for AI security that really met the moment.
- Telefonica also focused more than usual on trust and security this year. Despite its brand new "Trust in the Future" branding, Orange had very little to say or demonstrate on managing AI risk compared with its telco peers.

The state of AI beyond the confines of the Fira

The direction the telecom sector takes is increasingly driven by AI. Hence the general state of AI ethics, AI safety and AI security immediately before, during, and immediately after MWC is critical context for evaluating what was said about cybersecurity this year:

- **OpenClaw ran amok the week before MWC:** Days before MWC, Summer Yue, formerly of Google Deep Mind, now Director of Alignment at Meta Superintelligence Labs, took to 'X' to post screenshots of her OpenClaw autonomous AI agent going rogue. OpenClaw began deleting Yue's email inbox, despite her explicit instructions not to do so without confirming with her first. "I couldn't stop it from my phone", she posted. "I had to run to my Mac mini like I was defusing a bomb."
- **Anthropic resisted U.S government demands to shelve its ethics – so Open AI and xAI stepped in.** Two days before MWC opened its doors, the U.S and Israel launched their first missiles at Iranian targets. According to the Washington Post, "advanced AI technology is identifying targets in Iran and quickly prioritizing them, supporting the massive military operations." Among targets hit were the Shajarah Tayyebah girls' school in Minab. Most of the more than 160 killed were schoolgirls.

There followed a rupture in the Trump administration's relationship with Anthropic, the owner of Claude AI and a key supplier to the military. Around the time of those first strikes, reports emerged of Anthropic still declining to amend its terms to allow use of its LLMs for "any lawful use" as demanded by the government. Anthropic cited ethical concerns around mass surveillance and Claude's immaturity for autonomous weaponry use cases. The Trump administration responded by designating Anthropic a "supply chain risk to national security." Anthropic's competitors, OpenAI and xAI, exploited the fallout by hurriedly signing new deals with the military. On March 9th, Anthropic filed a lawsuit against the government. The day after, Microsoft filed an amicus brief supporting Anthropic.
- **AI agents are now hiring humans to do what they can't.** As reported by F-Secure on the last day of MWC, RentAHuman is a new platform that "lets AI agents hire real people for tasks they can't complete, like picking up packages or attending meetings. It promises flexible income in an automated future, but accepting tasks from AI agents means entering a system with little oversight and no ID verification."
- **The FCC Chairman called for "humble" approach to regulating AI.** During MWC, the FCC chairman, Brendan Carr, told Mobile World Live that "regulators need to be humble and understand the tremendous amount of upside that comes with [AI], and we need to be enabling and unleashing that."

RentAHuman is a new platform that lets AI agents hire real people for tasks they can't complete, but accepting tasks from AI agents means entering a system with little oversight and no ID verification.

There was no update to last year's 'C+' grade from the UK's NCSC

A year ago, the highlight of MWC 2025 was the 'C+' grade that Ollie Whitehouse, CTO of the UK's National Cyber Security Centre (NCSC) awarded the telecom sector for how it does cybersecurity in his talk at the GSMA's security summit. Whitehouse was in attendance again this year. He was sighted at a Sunday night dinner as well as on the exhibition floor on Monday. GSMA didn't invited him to speak again at this year's security summit, though. Perhaps last year was enough (perhaps it was more than enough).

Only Ollie Whitehouse would have awarded the sector a higher score this year, had he been invited to speak again. The scale of the breach impacts in South Korea in the last year might suggest not. So too would the freedom with which the big telcos continue to duck accountability for the Salt Typhoon breaches in the U.S. The disclosure of another sizable data breach on the Wednesday of MWC, this time by Telia Norway, will have pushed the odds of a 'B-' out still further. We'll probably never know.

Deutsche Telekom's "AI Agent-Ready" vision met the moment

In any case, like most trade fairs, MWC leans more towards the aspirational rather than the operational and last year's NCSC grade was most certainly levelled at the operational scorecard. On the much easier aspirational side of things it's fair to say that some of Europe's incumbent telcos really did up their game in terms of cybersecurity this year.

Normally, cybersecurity tends to feature at the margins of a telco's MWC exhibition booth. In past years quite a few haven't even featured cybersecurity at all. This year, Deutsche Telekom distinguished itself with a series of announcements, on-sight talks, and exhibition booth displays that genuinely met the moment in terms of articulating a vision for mitigating AI risk. These risks include the risk of AI agents themselves becoming what some are calling the 'new insider threat' of cybersecurity.

One reason DT stood out was that its messaging around AI security was simplified to make it easily accessible to non-security specialists. A display conveying DT's security framework for protecting against the number one new technology risk of our times was among the most inviting and engaging anywhere on the floor of the Fira.

Figure 1: Deutsche Telekom met the moment at MWC 2026



Source: HardenStance (pictured: Stefan Schneider, Deutsche Telekom's Head of AI Security)

One reason DT stood out was that its messaging around AI security was simplified to make it easily accessible to non-security specialists.

A sequence of 5 LCD panels at the centre of Deutsche Telekom's exhibition booth in Hall 3 depicted what DT framed as the 5 stages of AI security:

1. AI chaos (no security)
2. Development of an AI governance framework
3. Enforcement of an AI governance framework via a workflow automation platform capable of operating at extreme, unprecedented, scale and speed.
4. AI defence – network security operations
5. AI identity.

On each LCD panel representing one of the five phases, intuitive graphics depicted a variety of different AI agents in different shapes (rectangular as well as potentially malicious bug-shaped) and different colours (magenta pink for internal, white for external). They were depicted interacting dynamically with one another as well as with other resources in an IT environment. As each of the five phases of the AI security roadmap introduced new permissions, constraints and sanctions, some agents adjusted their movements and interactions, while others disappeared altogether.

Deutsche Telekom expects that governments will increasingly treat identity itself as a form of critical infrastructure, similar to water, power and telecom infrastructure.

Simplified storytelling underpinned by strong substance

DT's storytelling was highly simplified, but the technical substance underpinning the vision was solid too. The key messages were as follows:

- **Agentic AI needs to be secured as rigorously as any other software.** Out of the box there is no security built into the Model Context Protocol (MCP) and Agent-to-Agent (A2A) protocols. These are what AI agents have to use to access other resources and tools as well as communicate with one another; they need to be secured as rigorously as any other software.
- **Red teaming and other testing of LLMs is critical, including with AI agents to interrogate chatbots.** DT can test chatbots 20 times faster using AI agents. Humans can be removed from the loop entirely until an agent identifies a suspected compromise that requires human intervention. With AI, security monitoring is different from conventional approaches that monitor against technical metrics like anomalous protocol behaviours. Security monitoring of AI needs to monitor natural language content for attacks on LLMs that seek to exfiltrate sensitive information. With 6,000 honeypots deployed throughout its systems, DT has been collecting and leveraging threat data on how threat actors are targeting LLMs since last year.
- **AI agents, especially autonomous agents, must be assigned an identity and granted unique permissions.** When AI is enabling both human and machine identities to be replicated at scale and with very high quality, identity assignment and least privilege access rights are needed in the same way as they are for any other application, device, database or user in a Zero Trust framework.

At the level of an individual's own identity, DT shared the example of the actor, Matthew McConaughey. Earlier this year, he trademarked his image and voice to protect him from unauthorized use by AI platforms that risks damaging his reputation or personal brand. From the perspective of malicious threat actors exploiting a person's identity, the wealthy, famous and important are low hanging fruit for now. But as with all other types of cyber threats, this type of attack starts with these elites but then trickles down to all segments of society.

Drawing on the example of Denmark's leadership in passing a law that allows any Danish citizen to demand that a platform immediately remove any unauthorized use of their image and receive compensation, Deutsche Telekom expects other countries to follow suit. DT expects that governments will increasingly treat identity itself as a form of critical infrastructure - similar to water, power or telecom infrastructure.

DT believes any business or organization's brand is its' single most valuable asset. Boards must therefore engage in protecting their brand against potentially devastating damage by fake voice, images and videos.

- **In the age of AI, trust mechanisms are broken because any identity can be replicated.** Thomas Tschersich, Chief Security Officer of Deutsche Telekom and CEO of Telekom Security, has coined the useful term that "for the first time in history, human beings can't trust what they see and hear with their eyes and ears anymore." Trust can no longer just be declared or promised; instead it must be verified or proven. The only way to do that accurately and at scale is by using cryptography for verification and authentication. That means embedding a cryptographic key in everything, including every AI agent.

At the level of an organization, cryptographic key management has traditionally been fragmented across different technologies and departments. With the increasing risk posed by threat actors using AI to exploit weaknesses and inconsistencies in identity management, DT is positioning itself as a leading trusted provider of key management services for the AI era.

DT already does this for some of its customers. It does this from the same platform, and with the same level of security, that it uses to manage the identity of its 200,000 employees. When organizations ask themselves who they trust to manage their cryptographic keys, DT positions the 'T' in its own brand as standing for 'Trust'. In Germany and throughout Europe, DT brings that message to customer discussions around digital sovereignty. It is working with a number of partners, including Palo Alto Networks, in evolving its platform roadmap.

While the positioning was mostly around DT's medium and long term AI security strategy, the company did unveil one new product on that roadmap which it plans to launch within the next six months or so. Magenta Security Mobile.ID is a secure chip that DT's own employees will start testing in the spring of this year, aiming for commercial launch in the summer. Employees will use Mobile.ID as an alternative to physical IDs, smart cards and physical keys. Use cases include opening a security barrier to an underground car park, opening the door to an office complex, logging onto their laptop, encrypting emails, and verifying themselves as a registered user of a service. In this initial phase, Mobile.ID will be supported on Samsung smartphones. The plan is to then extend it to other vendor partners. The Mobile.ID product team envisages embedding a type of watermark in pictures or videos generated by a Mobile.ID-enabled smartphone to protect against misuse of images for deepfakes.

DT was also able to highlight its commitment to safely deploying agentic AI in the telco network domain at MWC. In partnership with Google Cloud, DT demonstrated Multi-Agentive Intelligent Network Diagnostics & Remediation (MINDR), a multi-agentive AI system that enables autonomous diagnostics and operations across multi-domain telecom networks.

Telefonica went big on AI security too – but Orange didn't

Security and trust featured much more prominently than usual in the positioning of Orange and Telefonica this year.

Telefonica also shared a vision of AI security across secure development and security operations domains. There were some parallels between its presentation on 'Human and non-human identity in the age of AI' and Deutsche Telekom's identity-related vision. But from an AI security perspective, Telefonica's positioning was a little bit drowned out by too many other themes on the company's stand. Compared with DT's vision, Telefonica's was quite a bit more rooted in the recent past and the here and now and somewhat less centred on an AI agent-driven future. There also seemed to be less in the way of mapping of vision to Telefonica's actual roadmap features.

From an AI security perspective, Telefonica's positioning was a little bit drowned out by too many other themes on the company's stand.

Two weeks before MWC, Orange launched a new 5-year branding strategy which it calls 'Trust in the Future'. This also served as the theme of the Orange booth in Hall 3. In her MWC keynote, Orange CEO, Christel Heydemann, observed that "the rise of agentic AI capable of autonomous actions is a vertiginous prospect, full of amazing promises but also unknown threats." Amidst this acceleration, she said, one question has never been more critical: who should we trust? "Are we on the brink of a major disruption in the real economy, one that could cause a collapse in many company valuations? Do we want to create a digital world ruled by chaos and threats? Aren't we as telecom operators uniquely placed to change the scenario and regain control of our digital futures?"

Heydemann's soaring rhetoric certainly met the moment. No speechwriters did a better job all week. But at the Orange booth, evidence of what kind of trust the company can deliver in the AI age – and how – was almost completely lacking. There were plenty of revenue-generating and cost-saving use cases, many of them augmented by AI. But there was very little emphasis on AI risk or how Orange will mitigate it. Of the 11 demos, only three addressed any aspect of risk mitigation:

- **Network-augmented facilities**, illustrating the optimization of security operations at critical sites such as airports;
- **Family protection**, a unified platform of innovative services designed to protect families in their digital lives;
- **Branded calling**, which demonstrated Orange's support for the STIR/ SHAKEN standard for authenticating phone calls as mandated by the French government.

Hardly any evidence of competence in AI security on Orange's booth

In terms of AI risk mitigation, however, Orange appeared to have almost nothing to say. Orange Cyberdefense didn't appear to have any dedicated presence on the booth. With its long telecoms heritage, it can't be likely that this particular French emperor has no clothes at all. In fact, in a Mobile Europe webinar on 'Becoming an AI-native telco' held in January, Philippe Ensarguet, Orange VP, Software Engineering, enthused about the many opportunities and challenges with AI. Among several agentic AI use cases Orange has been working on, he even singled out a multi-agent 5G security operations use case as "the best one." Orange's R&D, engineering and operations teams are bound to recognize the importance of securing AI agents in the development environment too.

For some not very obvious reason, the company chose not to showcase any of these capabilities at MWC. For some not very obvious reason, supporting evidence of competence in AI security was almost completely lacking. Despite the new 'Trust in the Future' branding, Orange somehow contrived to make a rookie error at MWC. The company asked for trust without providing a way to verify.

"MWC26: Taking Stock of Telco Security", Copyright: Patrick Donegan, HardenStance Ltd, 2026

HardenStance's online Telecom Threat Intelligence Summit

The screenshot shows a website banner for the 'TTIS Telecom Threat Intelligence Summit 2026'. The banner has a dark blue background with white and yellow text. At the top left is the logo 'TTIS Telecom Threat Intelligence Summit 2026'. To the right are navigation links: 'HOME', 'ABOUT', 'SPEAKERS', 'SPONSORS', 'FAQS', and 'EVENT MANAGEMENT BY St Albans Web Design'. Below the navigation, it says '2-DAY VIRTUAL EVENT' in yellow. The main title 'THE TELECOM THREAT INTELLIGENCE SUMMIT 2026' is in large white letters. Below the title is a subtitle: 'An event dedicated to improving cyber security outcomes for the telecom sector and its customers through better use of cyber threat intelligence.' At the bottom right, it says 'June 9th and 10th 2026'.

[Register here](#) for HardenStance's 2026 Telecom Threat Intelligence Summit

For some not very obvious reason, supporting evidence of competence in AI security was almost completely lacking.

More Information

- **Virtual Event:** Register for HardenStance's two-day "[Telecom Threat Intelligence Summit 2026](#)", taking place on June 9th and 10th 2026.
- **Briefing:** "[Telco Strategies for Consumer Security 2026](#)" (February 2026)
- **Briefing** "[MWC 25: Taking Stock of Telco Security](#)" (March 2025)
- **Briefing:** "[Telco Strategies for Consumer Security 2025](#)" (January 2025)
- **Briefing** "[MWC 24: Taking Stock of Telco Security](#)" (March 2024)

About HardenStance

HardenStance provides trusted research, analysis and insight in IT and telecom security. HardenStance is a leader in custom cyber security research and leading publisher of cyber security reports. HardenStance is also a strong advocate of industry collaboration in cyber security. HardenStance openly supports the work of key industry associations, organizations and SDOs including NetSecOPEN, AMTSO, The Cyber Threat Alliance, The GSM Association, OASIS, ETSI and TM Forum. HardenStance is also a formally recognized Cyber Threat Alliance 'Champion'. www.hardenstance.com.

To receive new public domain HardenStance reports as soon as they are released, register here (there are only five fields): [Registration Link](#).

Contact: Founder & Principal Analyst patrick.donegan@hardenstance.com

HardenStance Disclaimer

HardenStance Ltd has used its best efforts in collecting and preparing this report. HardenStance Ltd does not warrant the accuracy, completeness, currentness, noninfringement, merchantability or fitness for a particular purpose of any material covered by this report.

HardenStance Ltd shall not be liable for losses or injury caused in whole or part by HardenStance Ltd's negligence or by contingencies beyond HardenStance Ltd's control in compiling, preparing or disseminating this report, or for any decision made or action taken by user of this report in reliance on such information, or for any consequential, special, indirect or similar damages (including lost profits), even if HardenStance Ltd was advised of the possibility of the same.

The user of this report agrees that there is zero liability of HardenStance Ltd and its employees arising out of any kind of legal claim (whether in contract, tort or otherwise) arising in relation to the contents of this report.