

AI Spectrum INDIA | FREE WHITEPAPER · WORLD ENVIRONMENT DAY 2026
Carbon Cognition
 The hidden cost — and real promise — of AI on a finite Earth.
 Download free →



Home > Trustworthy AI / Cybersecurity > Industry Upwind

Industry Upwind

Why Sovereign AI Is Emerging As A Strategic Priority For European Enterprises

The challenge of managing Shadow AI, and the importance of building governance frameworks that align technology innovation with long-term societal and business priorities

05 June, 2026 | Interaction | By Shraddha Warde|shraddha.warde@mmactiv.com Share



JEFF SHAPRIO
Managing Director of Europe
for HaystackID

As Europe accelerates its push toward AI governance and digital sovereignty, organisations are facing mounting pressure to balance innovation with compliance, cybersecurity, and data privacy. From the evolving requirements of the EU AI Act and GDPR to the operational challenges posed by NIS2 and DORA, enterprises are increasingly seeking AI solutions that are transparent, auditable, and locally governed.

Against this backdrop, HaystackID is expanding its footprint across Europe to help organisations navigate complex cross-border regulatory environments while deploying AI responsibly at scale. In this interview with AI Spectrum, Jeff Shapiro, Managing Director of Europe for HaystackID, discusses the growing demand for "Sovereign AI," the rise

of AI-driven cybersecurity and compliance risks, and how the company is leveraging its proprietary GenAI capabilities to support defensible workflows in highly regulated sectors. The conversation also explores the future of responsible AI in Europe, the challenge of managing Shadow AI, and the importance of building governance frameworks that align technology innovation with long-term societal and business priorities.

What market trends or legal shifts are driving HaystackID's expansion across Europe?

I started my career in legal practice in Virginia before moving to London to lead teams within a Magic Circle law firm and a Big Four consultancy. I have spent my career navigating the very real cultural and legal friction between US and European data mentalities. While I am not a practising lawyer here at HaystackID, my background helps me understand that when a General Counsel faces a cross-border regulatory investigation, they need local, on-the-ground support that understands transatlantic tensions.

We have long served our European clients with a local, on-the-ground team supported by an international operational model. A significant driver of our local expansion is the increasing demand from our clients and the market for 'Sovereign AI'—the necessity to deploy artificial intelligence within strict, localised boundaries to maintain control over proprietary information within the strict regulatory frameworks across Europe. Additionally, we are seeing the disruption AI brings to the workplace, which, coupled with recent workforce restructurings, has led to a significant increase in Data Subject Access Requests (DSARs). This places a substantial burden on corporations to respond within strict GDPR timelines.

We are also responding to the need across Europe to track and prepare for critical regulatory milestones such as with the upcoming EU AI Act (various timelines from December 2026 to August 2028) and how this integrates into existing regulatory frameworks as well as the EU e-Evidence Regulation (applicable from August 2026) which will require companies in a criminal law context to respond to cross-border data production orders in as little as eight hours, demanding a very high degree of data readiness.

What are the biggest legal or regulatory risks addressed by HaystackID's technology and services?

The risks vary depending on the executive sitting across the table. For a General Counsel overseeing litigation or regulatory investigations, one of the most pressing risks is maintaining defensibility when deploying artificial intelligence such as Generative AI, keeping costs proportionate with outside counsel, consultancies, and technology providers, and ensuring compliance with cross-border data transfer laws.

Conversely, for CISOs and CTOs, the threat landscape is expanding. Threat actors are now actively using GenAI to execute increasingly convincing phishing campaigns. We are even seeing threat actors weaponising DSARs to maliciously map an organisation's sensitive data footprint prior to launching an attack. This risk is compounded by the rise of bring your own device (BYOD) and 'Shadow AI'—unsanctioned GenAI tools used by employees—which creates an ungoverned data footprint outside corporate control.

Corporations are also frequently exposed through third-party technology supply chains. This is one of the reasons why we own our GenAI engine. We have direct architectural control and strict data localisation which enables us to know who has access to the environment and keep the data ring-fenced within the required jurisdiction. This helps mitigate the risk of data leakage and reduces the vulnerabilities inherent in relying on external software and APIs.

How is HaystackID ensuring its AI-powered workflows remain compliant, auditable, and defensible under frameworks like GDPR, the EU AI Act, NIS2, and DORA?

To support compliance, it is important to balance competing legal frameworks. GDPR grants individuals the "right to be forgotten", whereas NIS2 and DORA require detailed, long-term forensic logs to demonstrate operational resilience. Our technology can help resolve this tension through pseudonymisation and redaction so clients can retain structural forensic records while honouring privacy mandates.

However, technology is only one part of the equation. Outside of the office, I am a passionate amateur garden designer. In garden design, you cannot just drop plants into a space without understanding their relationship to each other, the soil, climate, and the long-term ecosystem. Deploying AI across an enterprise is remarkably similar. You cannot just deploy an algorithm into a corporate workflow; you have to cultivate the surrounding processes, governance frameworks, and human oversight so the technology can actually take root and thrive sustainably.

This is why we have established strict international compliance frameworks at our company level. This structural setup helps ensure our people understand the regulatory environment and know what they can and cannot do. Frameworks like the EU AI Act introduce rigorous transparency requirements, and although obligations for certain high-risk AI systems have been delayed until December 2027, preparation is required now. Defensibility requires both safe tools and highly trained professionals guiding them.

How does the acquisition of eDiscovery AI help HaystackID operationalise/apply AI at scale in highly regulated environments?

Our acquisition of eDiscovery AI isn't just about reviewing legal documents and data for traditional litigation and regulatory investigations. We are using eDiscovery AI to help fill client needs across a much broader spectrum of data challenges. This includes cyber breach responses, M&A due diligence, information governance, data identification for retention and minimisation, and complex data privacy workflows.

Owning the GenAI engine provides distinct structural advantages for all of these use cases. It supports strict data localisation, allowing us to deploy the model directly within the jurisdiction our clients require. It also gives us direct control over the core infrastructure. For compliance teams, this provides a highly practical benefit: it significantly shrinks the client's sub-processor list, reducing the friction of auditing the vendors your vendors use.

However, deploying AI at scale is never just a technology exercise; it requires a deep operational partnership. Having a highly advanced GenAI engine means little if you do not have an experienced development team along with legal technologists and project managers who can communicate the mechanics of the workflow clearly and build bespoke, defensible procedures. We focus heavily on this human layer, ensuring our people can translate complex AI processes into practical, transparent, and collaborative strategies for our clients.

What do you see as the primary challenges or opportunities for deploying responsible AI systems in Europe over the next few years, especially in the legal/regulatory technology space?

The primary challenge is the tension between the velocity of AI development and the measured pace of European regulation. Managing "Shadow AI" is an ongoing challenge for enterprise security. Furthermore, responsible AI in Europe requires consideration of ESG factors. The significant computing power required for large language models (LLMs) carries a heavy carbon footprint, meaning sustainable, 'green' computing must be a core component of corporate responsibility here.

When we talk about the future of responsible AI in Europe, it is easy to get lost in the compliance jargon. At its core, it is about aligning AI with humanity and where we want to go as a society to help ensure that the benefits of the technology outweigh its risks. When I look at the children growing up today and reflect on my own childhood, I think deeply about the digital landscape their generation will inherit. Responsible AI is not just about avoiding regulatory fines; it is about building a digital ecosystem where privacy, human oversight, and transparent data practices are the default.

Pragmatically, having a clear, compliant data framework is becoming a fundamental baseline for market entry. It supports business continuity, helps protect intellectual property, and contributes to the operational stability required to navigate a regulatory audit or a cyber incident while maintaining the company's core operations.

AVTRON
POWER SOLUTIONS

Simplifying Data Center Readiness with Liquid-Cooled Load Banks

Download the whitepaper

The graphic features the Avtron logo at the top left. Below it, the title 'Simplifying Data Center Readiness with Liquid-Cooled Load Banks' is displayed in bold black text. At the bottom, there is a dark blue button with the text 'Download the whitepaper' in white. The background shows two workers in high-visibility vests and hard hats standing in a data center aisle.

Microsoft

Accelerating progress to 2030

Explore our annual environmental sustainability report →

The graphic features the Microsoft logo at the top left. Below it, the text 'Accelerating progress to 2030' is displayed in white on a green speech bubble background. At the bottom, there is a dark blue button with the text 'Explore our annual environmental sustainability report' in white and a right-pointing arrow. The background shows a rocky coastline with a glowing blue line representing progress or a path.

AI Tools Directory

NEWS

Claude Launches Opus 4.8, Bringing Smarter Automation And Lower AI Operating Costs

📅 01 June, 2026

NEWS

Google Launches AI Threat Defence To Counter AI-Powered Cyberattacks

📅 28 May, 2026

NEWS

Vestmark Launches AI-Powered Portfolio Intelligence Solution Vestmark Pulse

📅 13 May, 2026

NEWS

OpenAI Launches Advanced Account Security For ChatGPT Users Facing Elevated Cyber Risks

📅 05 May, 2026

Funding Tracker



NEWS

Anthropic Expands AI Compute Capacity Through New SpaceX Partnership

📅 07 May, 2026



NEWS

Anthropic, Blackstone And Goldman Sachs Launch New AI Services Company


📅 05 May, 2026

NEWS

IBM, Oracle Expand Partnership With New AI And Hybrid Cloud Innovations

📅 05 May, 2026

AI Watchlist



NEWS

Bayer Expands AI-Powered FarmRise Platform To Reach 5 Million Indian Farmers


📅 14 May, 2026



INTERACTION

The Future Of Mobility Will Be Software-Defined, AI-Native, And Open Source


📅 13 May, 2026



NEWS

EU Researchers Develop AI Tool To Turn Disaster News Into Actionable Risk Intelligence

📅 05 May, 2026



NEWS

SAP To Acquire Prior Labs To Establish A Globally Leading Frontier AI Lab In Europe

📅 05 May, 2026

NEWS

IBM, Oracle Expand 40-Year Partnership With New AI And Hybrid Cloud Innovations

📅 05 May, 2026

NEWS

China Makes AI-Powered Robots Core Of National Strategy

📅 05 May, 2026

Contact Us



📍 Ashirwad Bungalow, First floor,
36/A/2, S.No. 270, Pallod Farms,
Near Bank of Baroda, Baner Road,
Pune, Maharashtra, India 411045

✉ communications@mmactiv.com

☎ +91-9579069369

Enquiry

Send message

Subscribe To AISpectrum India Newsletter

AI Spectrum *INDIA*

The Business of AI and Industry Transformation

Enter your email

Subscribe

[About Us](#)

[Editorial Calendar](#)

[Terms & Conditions](#)

[Privacy Policy](#)

[Disclaimer](#)

Copyright © 2026 MM Activ Sci-Tech Communications. All Rights Reserved.