MIT Technology Review Insights

Produced in partnership with

EDB

Why sovereignty over data and models is becoming a defining factor in enterprise AI success, as well as a prerequisite for forging safe agentic systems.

Establishing AI and data sovereignty in the age of autonomous systems

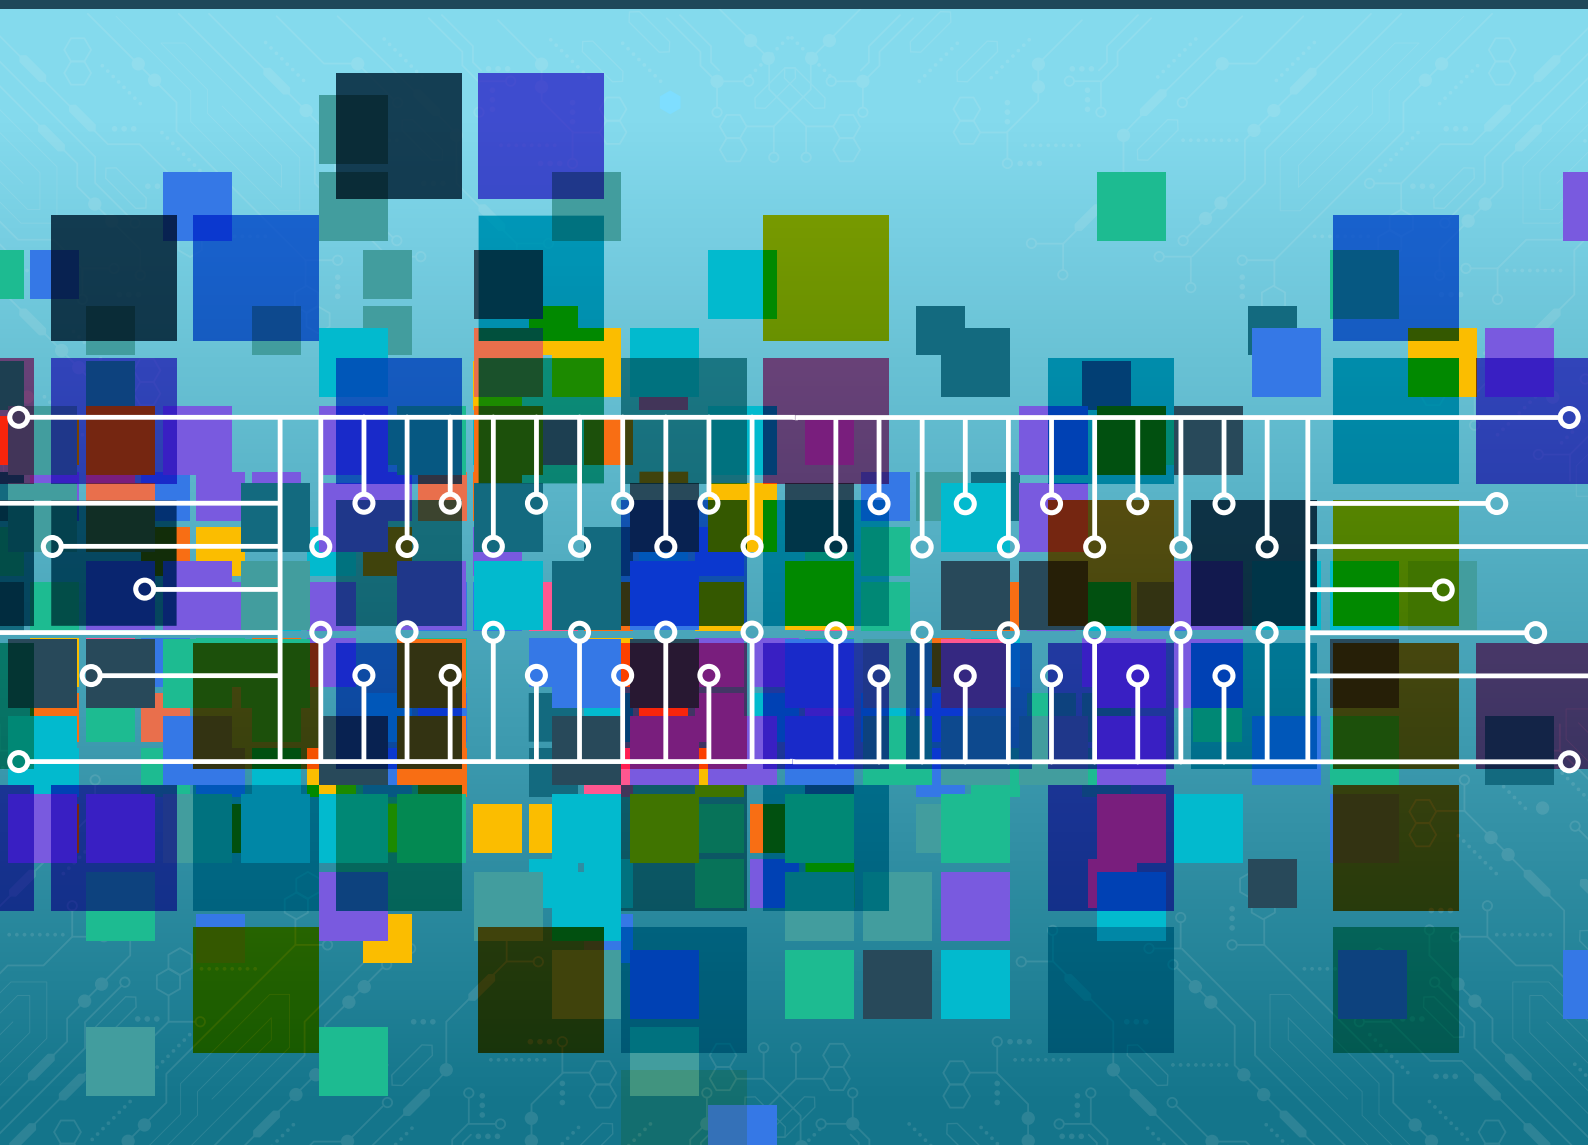

Preface

“Establishing AI and data sovereignty in the age of autonomous systems” is an MIT Technology Review Insights report sponsored by EDB. This report, based upon survey research and expert interviews, examines how enterprises are pursuing sovereignty over their models and data estates in an era of rapid AI adoption. Stephanie Walden was the author, Laurel Ruma was the editor, and Nicola Crepaldi was the publisher. The views expressed are those of MIT Technology Review Insights.

We would like to thank the following individuals for their time and insights:

Mukesh Chandak, Director of Strategy and Innovation, Cloud & AI, Thales

Kevin Dallas, CEO, EDB

Faisal Hoque, Founder, SHADOKA and NextChapter

Devin Pratt, Research Director, Data Management, AI, Automation, Data & Analytics, IDC

Michael Schrage, Research Fellow, MIT Sloan School’s Initiative on the Digital Economy

CONTENTS

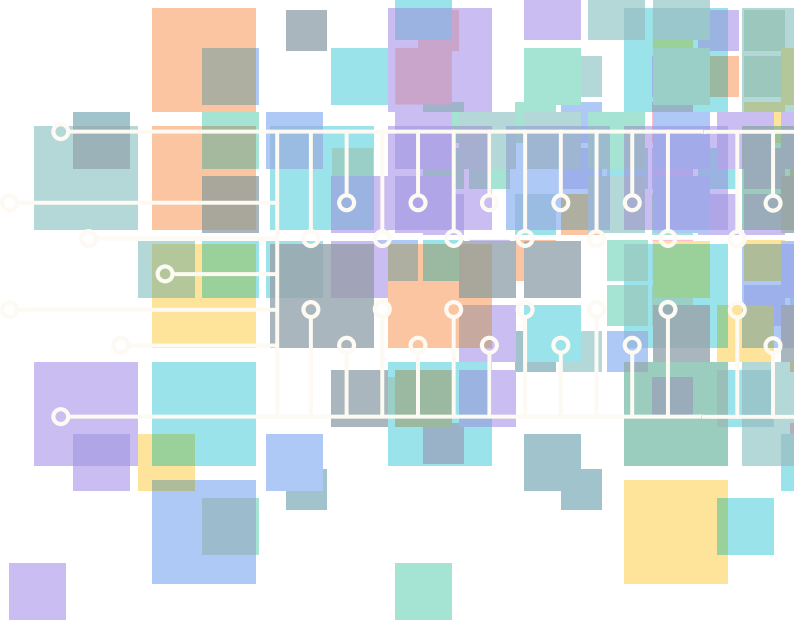

01 Executive summary

When generative AI first moved from research labs into real-world business applications, enterprises made a tacit bargain: "Capability now, control later." Feed your proprietary data into third-party AI models, and you will get powerful results. But your data passes through systems you do not own, under governance you do not set. The protections you rely on are only as durable as the provider's next policy update.

Now, with generative AI established in everyday business operations and sophisticated new agentic AI systems continually advancing, companies are reevaluating the terms of that deal.

"Data is really a new currency; it's the IP for many companies," says Kevin Dallas, CEO of EDB, echoing a recurrent anxiety from customers. "The big concern is, if you're deploying an AI-infused application with a cloud-based large language model, are you losing your IP? Are you losing your competitive position?"

That question is now fueling a movement toward reclaiming both the data and AI systems that have rapidly become part of core business infrastructure. AI and data sovereignty, which refers to breaking dependence on centralized providers and establishing genuine control over models and data estates,[1] is an urgent priority for many companies, says Dallas, citing internal EDB data: "70% of global executives believe they need a sovereign data and AI platform to be successful."

The idea of AI sovereignty is becoming a global policy conversation. NVIDIA CEO Jensen Huang recently spoke about the need for such a shift at the World Economic Forum's annual meeting at Davos in January 2026: "I really believe that every country should get involved to build AI infrastructure, build your own AI, take advantage of your fundamental natural resource – which is your language and culture – develop your AI, continue to refine it, and have your national intelligence be part of your ecosystem,"[2] Huang said.

This report explores how enterprises are pursuing sovereignty over their models and data estates in an era of rapid AI adoption. Drawing on a survey conducted by EDB of more than 2,050 senior executives and a series of interviews with industry experts, the research confirms that the sovereignty movement on the enterprise level is already well underway.

Key findings from the report include the following:

- **Once primarily a concern for regulated industries, sovereignty has now become a near-universal strategic priority.** The survey found that 95% of enterprises are planning to establish their own AI and data platforms within the next three years. Most express a clear sense of urgency (i.e., wanting to achieve sovereignty within one year).

- **A commitment to sovereignty correlates with AI performance.** Organizations categorized as "deeply committed" to sovereignty in the survey (13% of respondents) report approximately five times the return on investment (ROI) of less-committed cohorts when it comes to their agentic AI and generative AI initiatives. Modeled analysis found a 0.93 correlation between sovereignty commitment and success outcomes.

- **Security concerns top the list of motivations for sovereignty, but control and competitive advantage follow closely behind.** When asked why sovereignty matters, enterprises pointed first to security and resilience (85%), second to data localization requirements (74%), and third to ownership and control (72%). Only a small percentage (7%) cited geopolitics as a driver.

- **Agentic AI is stress-testing sovereignty frameworks.** When AI systems act on an organization's behalf, the question of who controls them – and who is accountable for their decisions – becomes murkier. Sovereignty will be critical for companies seeking maximum efficiency gains from agentic AI. The "deeply committed" cohort is already delivering a wider range (2x) of agentic and generative AI applications in mainstream production.

- **Sovereignty does not mean going it alone.** Achieving sovereignty does not require abandoning cloud providers or building private infrastructure from scratch. The survey found that 45% of enterprises now operate hybrid environments to balance performance with control. A number of tactics and emerging technologies exist to make sovereign infrastructure practical at enterprise scale.

"Data is really a new currency; it's the IP for many companies. The big concern is, if you're deploying an AI-infused application with a cloud-based large language model, are you losing your IP? Are you losing your competitive position?"

Kevin Dallas, CEO, EDB

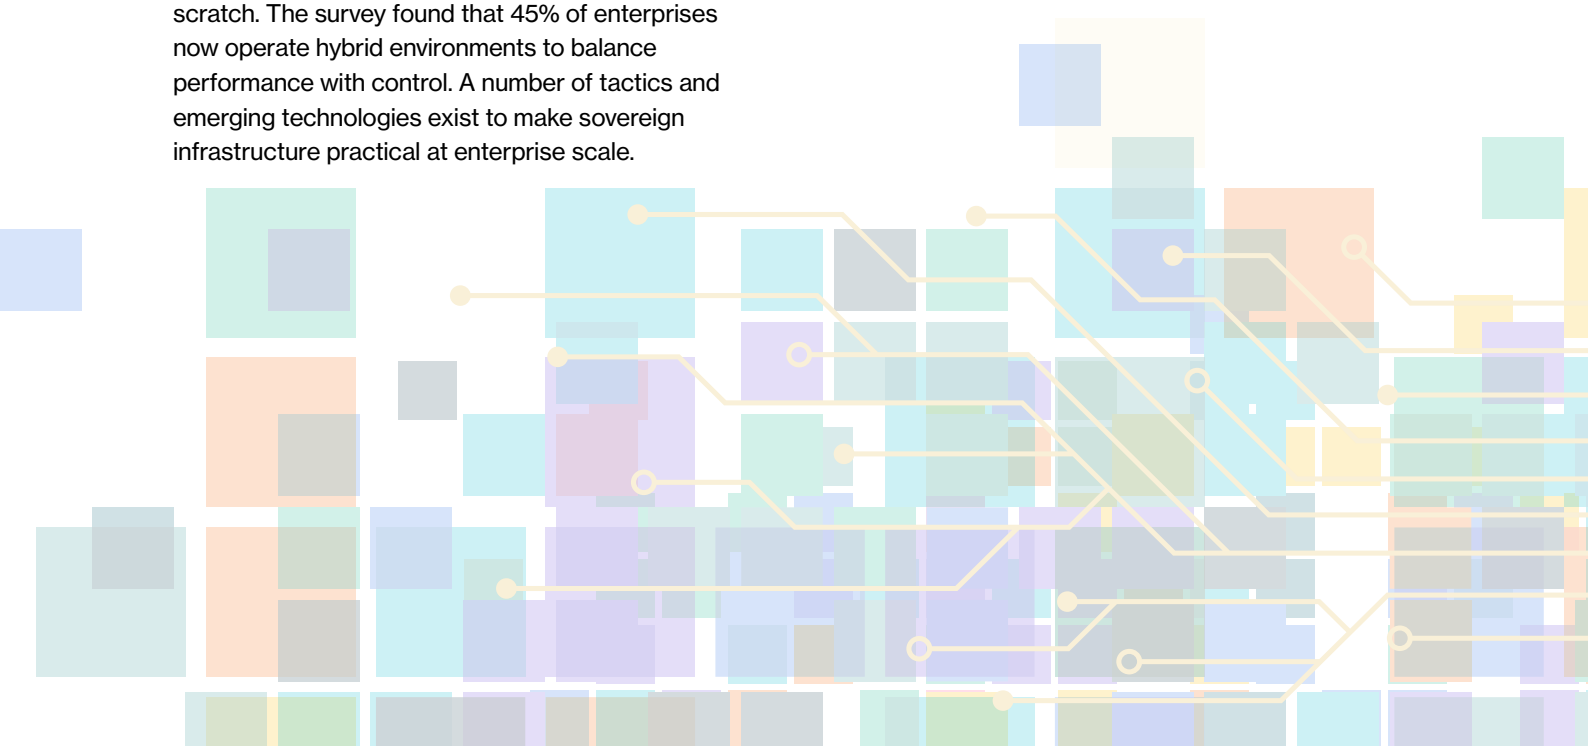

02

Defining sovereignty in the AI era

Sovereignty is a well-established concept in geopolitics, where it denotes a nation's right to self-governance. In the digital era, the term has migrated into enterprise discussions about data control.

Establishing that control is made more complex by the fact that AI systems involve deeply interconnected stacks. At Davos, Huang described AI not as a single technology but as a five-layer system spanning energy, chips, cloud infrastructure, models, and applications. That infrastructure undertaking, he argued, now rivals the construction of roads and electrical grids in its scope and strategic importance. "We have started the largest infrastructure buildout in human history," he said.

The implication for enterprises is the same as for nations: Whoever controls the infrastructure controls the intelligence it produces. That is precisely what data sovereignty and AI sovereignty are about. And while the two concepts are related, they are distinct.

Dallas frames these ideas as questions many businesses are grappling with today: "How do I build my AI applications on open source technology rather than proprietary platforms that lock me in? And how do I run my workflows and store my data in a sovereign deployment model – on-premises, in my region, in my country, or across a hybrid environment – rather than ceding that choice to a third party?"

Data sovereignty, at its core, involves knowing exactly where data lives, under which laws it falls, who can access it, and having the practical ability to enforce all of that. It encompasses the traditional concerns of data residency and jurisdictional control, but bleeds into more nuanced territory, too. "Data sovereignty and governance is about custody and control, consent and compliance," explains Michael Schrage, a research fellow at the MIT Sloan School's Initiative on the Digital Economy.

AI sovereignty extends those principles to models and systems, but the scope is broader. "AI sovereignty is profoundly actionable and contextually different," says Schrage. "It's the entire decision environment and the operational authority that surrounds it."

> "Data sovereignty and governance is about custody and control, consent and compliance. AI sovereignty is ... the entire decision environment and the operational authority that surrounds it."
>
> **Michael Schrage**, Research Fellow, MIT Sloan School's Initiative on the Digital Economy

Figure 1: 95% of organizations want to establish their own AI and data platforms within three years
Sovereignty has wide appeal, but organizations' depth of dedication and commitment varies.

Does your organization want to create its own AI and data platform? If so, then when?

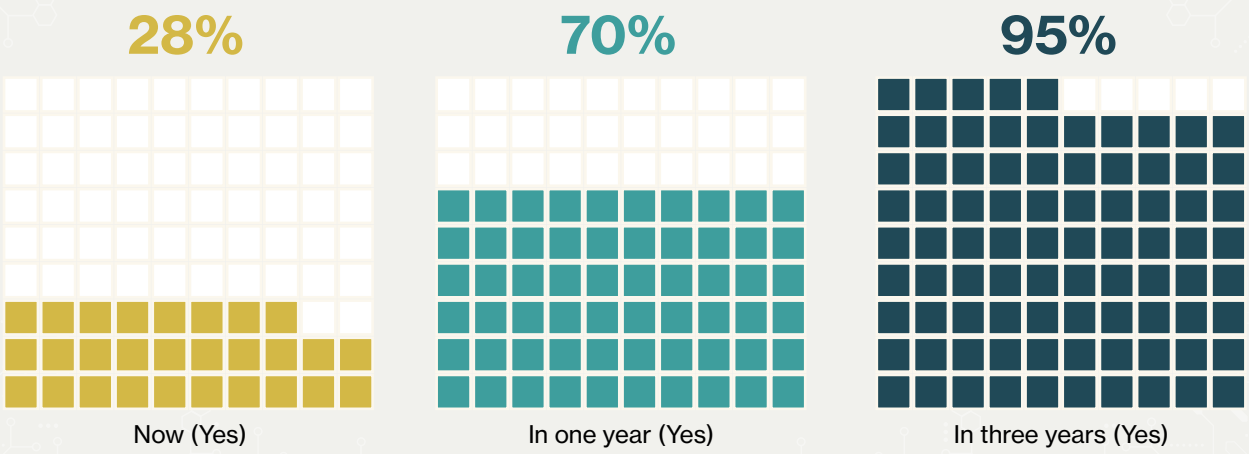

Source: Compiled by MIT Technology Review Insights, based on data from "Sovereignty Matters," EDB, 2026[3]

That wider aperture includes deciding where models run, what data they can touch, how their behavior is governed, and how their decisions are audited and explained. Critically, it also includes the decision rights over how models are trained, fine-tuned, and integrated into workflows.

Such questions become even more consequential as organizations adopt agentic AI systems capable of independent action—a transition that's happening faster than most governance frameworks can track. Devin Pratt, research director for data management, AI, automation, and data and analytics at IDC, offers a complementary framing: "Sovereign data governs the information; sovereign AI governs the systems that act on that information."

The business case for sovereignty

EDB's 2025 research, based on a survey of more than 2,050 executives along with modeled analysis, reveals that AI and data sovereignty has become a near-universal planning priority. A full 95% of enterprises aim to establish their own AI and data platforms within three years (see Figure 1).[4]

"Sovereign data governs the information; sovereign AI governs the systems that act on that information."

Devin Pratt, Research Director, Data Management, AI, Automation, Data & Analytics, IDC

But the most striking findings from the research point to a relationship between sovereignty commitment and business outcomes. In the modeled analysis – built from the global survey and more than 15,000 simulations across 500-plus variables – an enterprise's dedication to data and AI sovereignty is the single strongest factor associated with AI success. Using regression modeling, the analysis found a 0.93 correlation between sovereignty commitment and success outcomes. Success was measured by the number of generative and agentic AI applications in mainstream production and the ROI those applications deliver across seven performance metrics.

In a more general sense, sovereignty offers advantages across both defensive and offensive dimensions. Defensively, it reduces exposure to vendor lock-in and enables regulatory compliance across jurisdictions and geographies. Concern over such vulnerabilities is widespread: A 2026 Capgemini report found that 54% of organizations now list data control as a top strategic priority.[5]

On the offensive side, sovereignty creates conditions for differentiation: for example, proprietary AI capabilities that competitors cannot replicate, or faster innovation through controlled experimentation. Deloitte's 2026 State of AI in the Enterprise report,[6] which surveyed more than 3,000 senior leaders across 24 countries, found that only a third (34%) of organizations are using AI to truly transform their business. The other two-thirds are optimizing at the margins.

In other words: There is substantial first-mover advantage here that is yet to be captured. "If you're looking to differentiate your relationship with your listed firms or your high-net-worth clients, what's the differentiation there? What insights for investment or risk management can you offer as a result of the way you are engaging with your ensemble of AI capabilities?" Schrage asks.

Common sovereignty misconceptions

The EDB research reveals broad enthusiasm for sovereignty, but experts caution there are several common misunderstandings about what it entails. Four main ones stood out in our interviews:

Misconception 1: Sovereignty means isolation and compromised capabilities. Perhaps the most persistent misconception is that sovereignty requires walling yourself off. "Sovereignty is not isolation," says Faisal Hoque, an entrepreneur and author whose firms SHADOKA and NextChapter enable organizations with management frameworks for technology-driven innovation and transformation. "It comes from having discipline of control, evidence, and choices."

Yet, some enterprise leaders still assume that pursuing sovereignty means giving up the innovation capabilities and scale of major cloud platforms. In reality, the two can be compatible. Mukesh Chandak, director of strategy and innovation for cloud and AI at Thales, points to his organization's joint venture with Google Cloud – a company called S3NS – as proof of concept. "Enterprise leaders often believe they have to create their own private cloud from scratch or use outdated, subpar tools, but that's not the reality," he says. "Both sovereignty and performance at the level of hyperscalers can be compatible. You don't have to choose between being smart and modern or being safe and sovereign."

"The C-suite is quite behind in terms of understanding AI in general, let alone the nuance of data sovereignty and ownership and policy, and driving those policy guidelines. Leadership needs to come up to speed."

Faisal Hoque, Founder, SHADOKA and NextChapter

S3NS represents one end of the spectrum: sovereignty through partnership with a hyperscaler under strict jurisdictional controls. At the other end, organizations are building sovereign capabilities on open source platforms that eliminate third-party dependence entirely. The common thread is that sovereignty, done right, does not require retreating into a technological silo.

Misconception 2: Sovereignty is primarily about compliance checklists. Importantly, sovereign AI is not synonymous with "secure and compliant AI," though the two are easily conflated. Chandak points out that compliance can create a false sense of control: A system can be fully secure and compliant with regulations like GDPR, he notes, and still be exposed. "If a foreign government can issue a subpoena to your cloud provider and try to get access to your data, you're no longer sovereign," Chandak says.

Schrage notes that many organizations, having weathered the burden of data governance, may be reluctant to take on another arduous procedural exercise. But he echoes Chandak that viewing sovereignty through the lens of compliance misses the point. "For understandable reasons, many organizations have a mentality of, 'We just finished struggling with data governance, and now you want to impose AI governance?'" he says. "But it's really about decision governance and decision rights, not just box-tick compliance."

Misconception 3: Sovereignty is a one-time policy declaration. Sovereignty is not something an organization can declare, dust their hands off, and walk away from. It requires cultural adoption, from the C-suite to the project-implementation level. "When sovereignty is taken seriously, it becomes part and parcel of governance, ethics, clear ownership, disciplined architecture, and management habits," says Hoque.

Organizations that treat sovereignty as a one-time achievement also risk falling behind as AI capabilities evolve. "Sovereignty is going to have to be way more dynamic than some people are considering," warns Schrage. "This is continuous improvement – Kaizen for AI." As models improve and are composed into multi-agent workflows, the governance frameworks surrounding them must evolve at the same pace.

Misconception 4: Sovereignty requires locking down everything at once. Some organizations assume they need a comprehensive, enterprise-wide solution before they can begin with sovereignty efforts. Hoque pushes back on this. He advocates what he calls "selective sovereignty," a contextual approach that calibrates protections based on what the data is, who owns it, and how it is being used.

"You can't just get up and say, 'We need data sovereignty.' Data sovereignty for whom?" Hoque asks. A customer dataset powering personalized AI recommendations, for instance, carries very different sovereignty requirements than a public-facing marketing model. The smart approach, he argues, is to build guardrails around specific data – protecting its ownership, permitted uses, and distribution – rather than trying to boil the ocean. Organizations that wait for a perfect, all-encompassing framework risk never starting at all.

03 The current state of enterprise-level sovereignty

The survey data suggests that sovereignty has wide appeal, but organizations' depth of dedication and commitment to the concept varies, with enterprises falling into four distinct cohorts (see Figure 2).

The "deeply committed" (13%): These organizations have embedded AI and data sovereignty into their operating model. They report approximately five times the ROI of the lowest-performing segment (the "sideliners") and deliver roughly twice as many agentic and generative AI applications in mainstream production. Nearly half (49%) of their workloads occur in hybrid and sovereign environments – the highest commitment to both sovereignty and hybrid infrastructure of any segment.

The "strivers" (26%): These organizations understand the strategic importance of sovereignty but are earlier in execution. They report approximately twice the ROI of the "sideliners." Many acknowledge they risk being caught by competitors unless they deepen their commitment.

The "half-ins" (34%): The largest segment globally, these organizations believe in the value proposition of sovereignty but lack clear implementation plans. They report approximately 1.6 times the ROI of the sideliners. Their commitment to data interoperability and open standards is high, but without a structural commitment to sovereignty as mission-critical, the gap between them and the "sideliners" will narrow significantly over three years, EDB's modeling suggests.

Figure 2: Four levels of AI and data sovereignty
Just 13% of organizations consider sovereignty mission critical, and are achieving the highest level of ROI as a result.

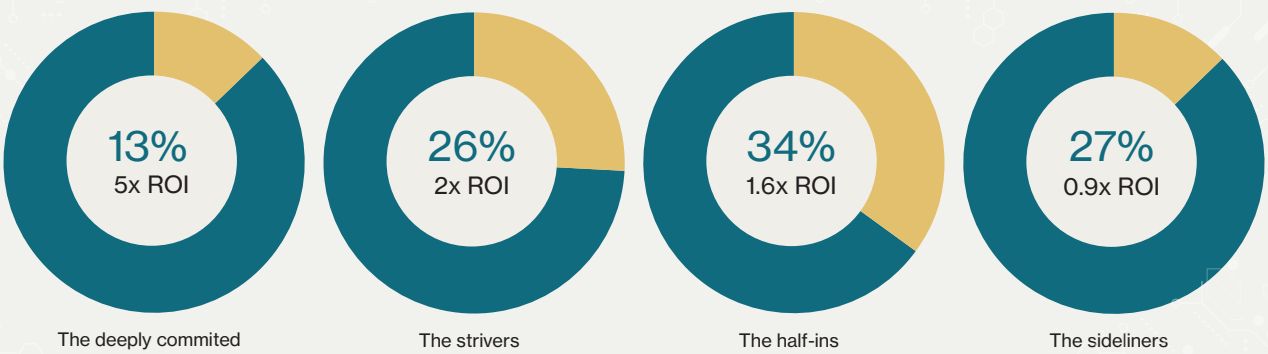

Source: Compiled by MIT Technology Review Insights, based on data from "Sovereignty Matters," EDB, 2026[7]

The "sideliners" (27%): These organizations are either unconvinced of the value of sovereignty or disengaged from it. Their low belief in the economic power of AI and data sovereignty acts as a barrier to getting value from their initiatives.

Notably, the distribution of these segments is consistent across geographies, from North America to EMEA to Asia-Pacific. The UAE, Scandinavia, and Germany show the highest propensity for sovereignty right now, while the fundamental patterns hold steady across all countries surveyed.

Early examples of sovereignty building blocks

While enterprise AI sovereignty is still maturing, a number of current use cases illustrate what sovereign AI looks like in practice:

Financial services: The sovereign prompt layer. In regulated banking, institutions are building sovereign prompt and policy layers where every AI-driven compliance task is logged and traceable, from the initial query to the data retrieved to the final output. This facilitates full transparency: If someone asks how the AI reached a decision, the bank can replay the entire chain. "We've now created a traceability environment so we can see the impact," says Schrage. "I can show and explain to regulators, my board, and my auditors."

Health care: Governed RAG for electronic health records. Health systems are exploring ways to convert electronic health records (EHRs) into retrieval-augmented generation (RAG) frameworks governed by strict anonymization. "This allows you to do population studies with no risk for individuals," explains Schrage.

Manufacturing: Sovereign predictive maintenance. Industrial manufacturers are deploying AI models trained on proprietary sensor data, production-line telemetry, and quality-control records to predict equipment failures and optimize output.[8] Because this data often encodes decades of hard-won process knowledge (and because sharing it with a cloud provider could expose competitive advantage), models often run on-premises or in controlled hybrid environments where the manufacturer retains full ownership.

Rules the "deeply committed" abide by

EDB's research found that the survey cohort who are "deeply committed" to data and AI sovereignty have a few characteristics in common. The respondents who fall into this group tend to:

1. **Solve needs in parallel.** Rather than treating security, access, and scale as separate problems to solve sequentially, they address all three simultaneously.
2. **Define success broadly.** Instead of reducing AI performance to a single number, they measure returns across a range of outcomes – from cost savings to innovation speed – giving them a clearer picture of where value is actually being created.
3. **Get more from each deployment.** Where less-committed organizations build AI for one purpose, the "deeply committed" design theirs to serve multiple business goals at once.
4. **Know when to scale and when to move on.** They track returns across each AI use case, and when one reaches diminishing returns, they redeploy resources to the next opportunity with confidence.

Top motivations for sovereignty

When asked to identify the drivers behind their sovereignty strategies, three priorities emerged. First and foremost, security and resilience, selected by 85% of survey respondents (see Figure 3). This reflects the need to protect AI and data systems against cyber threats and unauthorized access.

Figure 3: Top priorities for AI and data sovereignty
For most enterprises, sovereignty is a pragmatic business priority.

Top priorities	
Security & resilience - Data and AI systems are protected against cyberthreats and foreign interference. Robust encryption, access controls, and disaster recovery strategies are in place.	85%
Data localization - Sensitive data is stored and processed within the country of origin to comply with regulations, ensuring protection from foreign influence, surveillance, or unauthorized access.	74%
Data ownership & control - The data is owned, stored, and processed within a specific jurisdiction. Organizations or governments retain full control over how the data is used, accessed, and shared.	72%
Legal & regulatory compliance - Adheres to national and international data protection laws (e.g., GDPR, CCPA, AI Act). Aligns with cybersecurity frameworks and sovereign governance policies.	39%
Independence from foreign tech dependencies - Minimizes reliance on foreign cloud providers, AI models, and software. Encourages domestic innovation in AI and computing infrastructure.	32%
Trust & citizen rights protection - Ensures individuals have control over their personal data. Protects privacy and fundamental rights in AI applications.	29%
Interoperability & open standards - Encourages collaboration within sovereign ecosystems. Uses open source AI frameworks or proprietary solutions that comply with local standards.	27%
Economic & strategic value - Supports local industries, research, and economic development. Encourages AI sovereignty to maintain competitive advantage and national interests.	23%
Ethical & transparent AI - AI models are built with fairness, explainability, and bias mitigation. Ensures that AI decisions align with national values and ethical guidelines.	20%

Source: Compiled by MIT Technology Review Insights, based on data from "Sovereignty Matters," EDB, 2026[9]

Data localization requirements followed at 74%, likely driven by the growing web of regulations mandating that sensitive data be stored and processed within its country of origin. Ownership and control came in at 72%.

Further down the list, legal and regulatory compliance registered at 39% and independence from foreign technology dependencies at 32%.

Triggers that accelerate sovereignty initiatives

Pratt explains that, according to IDC's research, three patterns consistently bring sovereignty into sharper relief:

1. Regulatory expansion or new market entry: When an organization moves into a jurisdiction with stricter data residency or access rules, sovereignty shifts from a legal discussion to an architecture and procurement decision.

2. Transition from AI pilots to production, particularly with agentic systems. Once AI agents start touching real data and taking real actions, questions about control and accountability become more pressing.

3. Serious data classification work: When teams begin categorizing their data and discover that certain assets carry very high regulatory or reputational risk, sovereignty takes on new urgency. "In these instances, sovereignty tends to shift from theoretical to concrete design and operational model decisions," says Pratt.

04 Challenges to achieving sovereignty

Several experts caution that the greatest obstacles to sovereignty may not be technical, but cultural. Schrage identifies a fundamental tension between top-down governance and bottom-up innovation: "I see a very interesting, inevitable, and potentially conflicting intersection of what the board and executive committee thinks constitutes effective AI sovereignty versus what innovative groups in the organization want," he says.

In practice, different parts of an enterprise may have very different risk tolerances or definitions of what sovereign means for their specific workflows. For instance, a data science team pushing to experiment with new models may chafe at restrictions that a compliance function considers essential. Without a shared framework, these disagreements can harden into ad hoc, inconsistent approaches.

Hoque highlights a related problem: Sovereignty must be embedded in how teams actually work, not just proclaimed in policy documents. "You can't just say, 'We have a data sovereignty policy.' What does that mean? Is the developer who's accessing the data aware of it and how they should be implementing it at a code level? Or is it that you have a policy, but nobody really follows it?" Without that operational depth, sovereignty risks becoming what Hoque calls a "paper exercise," i.e., visible in the boardroom but invisible where it actually matters.

Part of the challenge here, Hoque argues, is that leadership hasn't yet developed the fluency to bridge that gap. "The C-suite is quite behind in terms of understanding AI in general, let alone the nuances of data sovereignty and ownership and policy and driving those policy guidelines," he says. "Leadership needs to come up to speed."

> "Sovereignty defines which agents can touch the data, in which region, under which policies, and how all of that is monitored and audited. Rather than being a brake on agentic AI, sovereignty sets the safe operating boundaries that allow organizations to scale with confidence."
>
> **Devin Pratt**, Research Director of Data Management AI, Automation, Data & Analytics, IDC

"You leave your crown jewels – customer data, IP – outside of AI access, because you're in that fearful environment. If you flip the coin... you have the confidence to feed your most valuable data through your AI models and extract massive ROI."

Mukesh Chandak, Director of Strategy and Innovation, Cloud & AI, Thales

The rapid rise of agentic AI

Agentic AI represents a transition from tools that assist to systems that act. Unlike earlier AI applications that generated outputs for human review, agentic systems can take actions on behalf of the business: updating databases, executing transactions, communicating with external systems, and orchestrating multi-step workflows with minimal or no manual oversight. The adoption velocity thus far has been remarkable. More than half (53%) of organizations say they already have AI agents in production and another 28% say they are deploying them, according to 2025 research from IDC.[10]

Sovereignty in the agentic AI era, however, is a whole new ballgame. Agents require real-time access to operational data, and the fact that they can take independent actions introduces new dimensions of risk, including chain reactions that cannot easily be undone. For example, when a user engages an agent, and that agent taps another agent, and so on down the line, the question of who governs what – and where accountability lies – becomes increasingly unclear. "Sovereignty provides the foundation at each level of that chain," says Chandak.

Put differently, sovereignty becomes the operating system for agentic AI, not just a policy layer on top of it. "Sovereignty defines which agents can touch the data, in which region, under which policies, and how all of that is monitored and audited," notes Pratt. "Rather than being a brake on agentic AI, sovereignty sets the safe operating boundaries that allow organizations to scale with confidence."

Chandak sees a direct line from sovereignty to performance that runs through trust. Without it, enterprises tend to play it safe, which means leaving value on the table. "You leave your crown jewels – customer data, IP – outside of AI access, because you're in that fearful environment," he says. "If you flip the coin... you have the confidence to feed your most valuable data through your AI models and extract massive ROI. That additional control removes fear and allows enterprises to adopt AI in an accelerated manner for high-value, mission-critical transformation."

05 Emerging solutions and next steps

A set of maturing technologies is making AI and data sovereignty more viable, both technically and economically, than it was even two years ago:

Confidential compute protects data while it is being processed, not just at rest or in transit. AI models can now run in hardware-enforced enclaves where even the infrastructure provider cannot see what is happening on the data or model side. “That’s an impressive technology that we’re going to see a lot more of,” says Chandak.

Edge AI and frugal AI reduce dependency on centralized cloud infrastructure by enabling models to run at the point of data generation, like on a sensor or drone (as in the manufacturing use case cited earlier). “Running AI at the edge will give an extra sense of comfort when it comes to the sovereign environment,” notes Chandak.

Post-quantum cryptography (PQC) addresses the emerging harvest now, decrypt later threat, in which adversaries collect encrypted data today with the intention of breaking it once quantum computing matures. “There’s been a lot of development in terms of defining the PQC-safe algorithm,” says Chandak. “We will see a lot more deployment of these algorithms, so that data remains safe even in a future where quantum computing becomes a reality.”

Converged data platforms are increasingly capable of serving AI workloads alongside traditional transactional and analytical processing. Pratt notes that database platforms are rapidly adding vector capabilities and hybrid transactional/analytical processing to support agentic AI. Such platforms reduce the need to ship data to external services.

Open foundations such as Postgres support portability and reduce lock-in. By building on open source platforms, enterprises gain the flexibility to move workloads across environments without rewriting applications or renegotiating vendor contracts.

The Postgres foundation

Open source is foundational to enterprise sovereignty because it ensures that no single vendor controls an organization’s data layer. PostgreSQL (Postgres) has emerged as the data platform of choice for many sovereignty-minded enterprises; it already powers some of the world’s largest transactional workloads and is increasingly optimized for AI-native applications.[11]

“Billions of transactions run on Postgres every day,” says Kevin Dallas, CEO of EDB. “Big, heavy workloads are already running on top of Postgres in data-driven, intelligent applications. So it’s already all around us.”

Within EDB’s survey, 19% of the “deeply committed” cohort are considering Postgres for their next mission-critical workload, and the platform’s hybrid deployment model (across clouds and on-premises) aligns naturally with sovereign architectures.[12, 13]

The 90-day sovereignty sprint

Sovereignty can feel like a multi-year transformation. And in its fullest form, it is. But that does not mean organizations need to wait years to start seeing results. For organizations looking to move quickly, Faisal Hoque, founder of SHADOKA and NextChapter, recommends a phased approach that can begin delivering results within 90 days:

Weeks 1 – 4: Map your data with a governance mindset. Overlay governance structures on your data mapping, focusing on ownership, privacy, usage, and deployment. "Not just 'I have this data, how do I use it?' but having that governance mindset from the start," says Hoque.

Weeks 5 – 8: Build a secure foundation for experimentation. Create a safe zone to test what works using a metadata subset. Establish access controls and monitoring.

Weeks 9 – 12+: Test with real projects across multiple dimensions. Deploy with real use cases across different types of projects, then fine-tune based on findings. "You should allow at least six to 12 weeks with different dimensions of projects, then iterate," advises Hoque.

Beyond technology itself, the economics of sovereignty have shifted meaningfully in recent years. Advances in infrastructure efficiency, the democratization of AI hardware, and the maturity of open source tooling have collectively lowered the financial barrier to running proprietary AI infrastructure. "Economically, it's now feasible to have that infrastructure for your own AI and build your intelligent applications," says Dallas.

Huang made the same point about open models at Davos: "Open models have enabled companies and industries, researchers, educators, universities, startups to be able to use these open models to start something and create something that's domain-specific or specialized for their needs," he said.

This trajectory of AI echoes earlier technology adoption arcs. Sovereignty is following a path similar to cloud and DevOps adoption five to 10 years ago. Once viewed as niche and complex, it is becoming a mainstream architectural concern. "Over time, I think we'll move from asking, 'Are you using AI in the cloud?' to 'How mature is your AI and data sovereignty strategy across the markets you serve?'" says Pratt.

Practical steps toward sovereignty

EDB's research outlines six components that constitute a foundation for true data and AI sovereignty:

1. Sovereign infrastructure that delivers hybrid control with the automation and agility of cloud services
2. AI-ready data migration that securely syncs existing data into a private knowledge base
3. Private and secure large language model integration using models of choice via low-code and no-code tools without exposing proprietary data
4. Access controls and encryption with granular role-based permissions and end-to-end protection
5. Comprehensive data and AI observability offering single-pane-of-glass monitoring across the data estate[14]
6. Always-on assurances that optimize for performance and availability

No single organization needs to tackle all six of these directives at once. As Hoque's selective sovereignty principle suggests, not every workload requires the same level of control. But together, these components describe the architecture of an enterprise that has moved from talking about sovereignty to practicing it.

06 Looking forward

The path forward is neither straightforward nor uniform. It demands cultural transformation alongside technical investment. It necessitates recognizing that sovereignty is not a destination but a continuous discipline – one bound to evolve alongside ever-more-powerful AI models, agentic systems that act with autonomy, and regulations still struggling to keep pace with both. And as with any discipline, it must be practiced daily.

Critically, it also requires calibration. "I fear that we tend to swing from one extreme to another," warns Hoque. "Just like we're having this nonstop conversation about whether AI is all bad or the next savior of the planet... it's probably neither. It's somewhere in between." The same applies to sovereignty: the goal is not to disconnect entirely from the broader AI ecosystem, but to engage with it on your own terms. "If you become a purpose- and impact-driven organization, then you don't fall into the hype," Hoque adds.

In practice, that does not look like abandoning the cloud entirely or scrambling to build a frontier model from scratch, but rather making a mindset shift from consuming AI as a service to having a clear grasp on the architecture that governs it. And it means doing so as quickly as possible. The window of opportunity to establish sovereign foundations before the competition does is narrowing.

Dallas put the competitive reality plainly: "We're in an AI-first, sovereign-first world. Build your own sovereign data and AI factories. Lead from the front. Disrupt or be disrupted."

> "We're in an AI-first, sovereign-first world. Build your own sovereign data and AI factories. Lead from the front. Disrupt or be disrupted."
>
> **Kevin Dallas**, CEO, EDB

Endnotes:

1. "2025 in Focus – 5 Key Predictions for Future of Data and AI Sovereignty," Chief Data Officer Magazine, January 21, 2025, https://www.cdomagazine.tech/branded-content/2025-in-focus-5-key-predictions-for-the-future-of-data-and-ai-sovereignty.
2. NVIDIA CEO Jensen Huang on how AI is becoming the next great infrastructure build, World Economic Forum, January 23, 2026, https://www.weforum.org/stories/2026/01/nvidia-ceo-jensen-huang-on-the-future-of-ai/.
3. "Sovereignty Matters: A global blueprint for sovereign, agentic, and generative AI," EnterpriseDB, 2025, https://www.enterprisedb.com/resources/sovereignty-matters-report.
4. "Sovereignty Matters: A global blueprint for sovereign, agentic, and generative AI," EnterpriseDB, 2025, https://www.enterprisedb.com/resources/sovereignty-matters-report.
5. Capgemini, The multi-year AI advantage: Building the enterprise of tomorrow, https://www.capgemini.com/insights/research-library/AI-perspectives-2026/.
6. Deloitte, 2026 State of AI in the Enterprise, https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html.
7. "Sovereignty Matters: A global blueprint for sovereign, agentic, and generative AI," EnterpriseDB, 2025, https://www.enterprisedb.com/resources/sovereignty-matters-report.
8. "Industrial AI in Action: Predictive Maintenance and Operational Efficiency at Scale," Association for Advancing Automation, June 20, 2025, https://www.automate.org/blogs/industrial-ai-in-action-predictive-maintenance-and-operational-efficiency-at-scale.
9. "Sovereignty Matters: A global blueprint for sovereign, agentic, and generative AI," EnterpriseDB, 2025, https://www.enterprisedb.com/resources/sovereignty-matters-report.
10. "AI Agents and Security: Moving Fast Without Breaking Things," IDC, October 2025, https://intelligence.theregister.com/paper/view/19166/ai-agents-and-security-moving-fast-without-breaking-things.
11. "4 Principles for Turning Your Sovereign Data and AI into a Platform for Success with Enterprise-Grade Postgres," Chief Data Officer Magazine, September 26, 2024, https://www.cdomagazine.tech/branded-content/4-principles-for-turning-your-sovereign-data-and-ai-into-a-platform-for-success-with-enterprise-grade-postgres.
12. "Sovereignty Matters: A global blueprint for sovereign, agentic, and generative AI," EnterpriseDB, 2025, https://www.enterprisedb.com/resources/sovereignty-matters-report.
13. "EDB Postgres AI Q2 Release Highlights," EnterpriseDB, June 17, 2025, https://www.enterprisedb.com/blog/edb-postgres-ai-q2-release-highlights-0.
14. "Seeing the Future Clearly: Why Your AI and Data Need the Right Lens," Digital Trends, January 17, 2025, https://www.digitaltrends.com/contributor-content/ai-data-sprawl/#dt-heading-a-single-pane-of-glass-for-the-intelligent-enterprise.

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review,* the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. This content was researched, designed, and written entirely by human writers, editors, analysts, and illustrators. This includes the writing of surveys and collection of data for surveys. AI tools that may have been used were limited to secondary production processes that passed thorough human review.

About EDB

EDB Postgres AI (EDB PG AI) is the first open, enterprise-grade sovereign data and AI platform – secure, compliant, and scalable, on-premises and across clouds. Built on Postgres, the world's leading database, EDB PG AI unifies transactional, analytical, and AI workloads, enabling organizations to operationalize their data and LLMs while maintaining control over sovereign environments. EDB PG AI is supported by a global partner network and delivers up to 99.999% availability as well as hybrid management and a built-in AI factory. As one of the most active contributors to the PostgreSQL project, EDB is deeply invested in the vitality of the global community.

To learn more, visit **www.enterprisedb.com**.

Illustrations assembled by Tim Huxford with elements from Shutterstock and Adobe Stock.

To cite this report, please use: "AI and data sovereignty in the age of autonomous systems," MIT Technology Review Insights and EDB, May 2026.

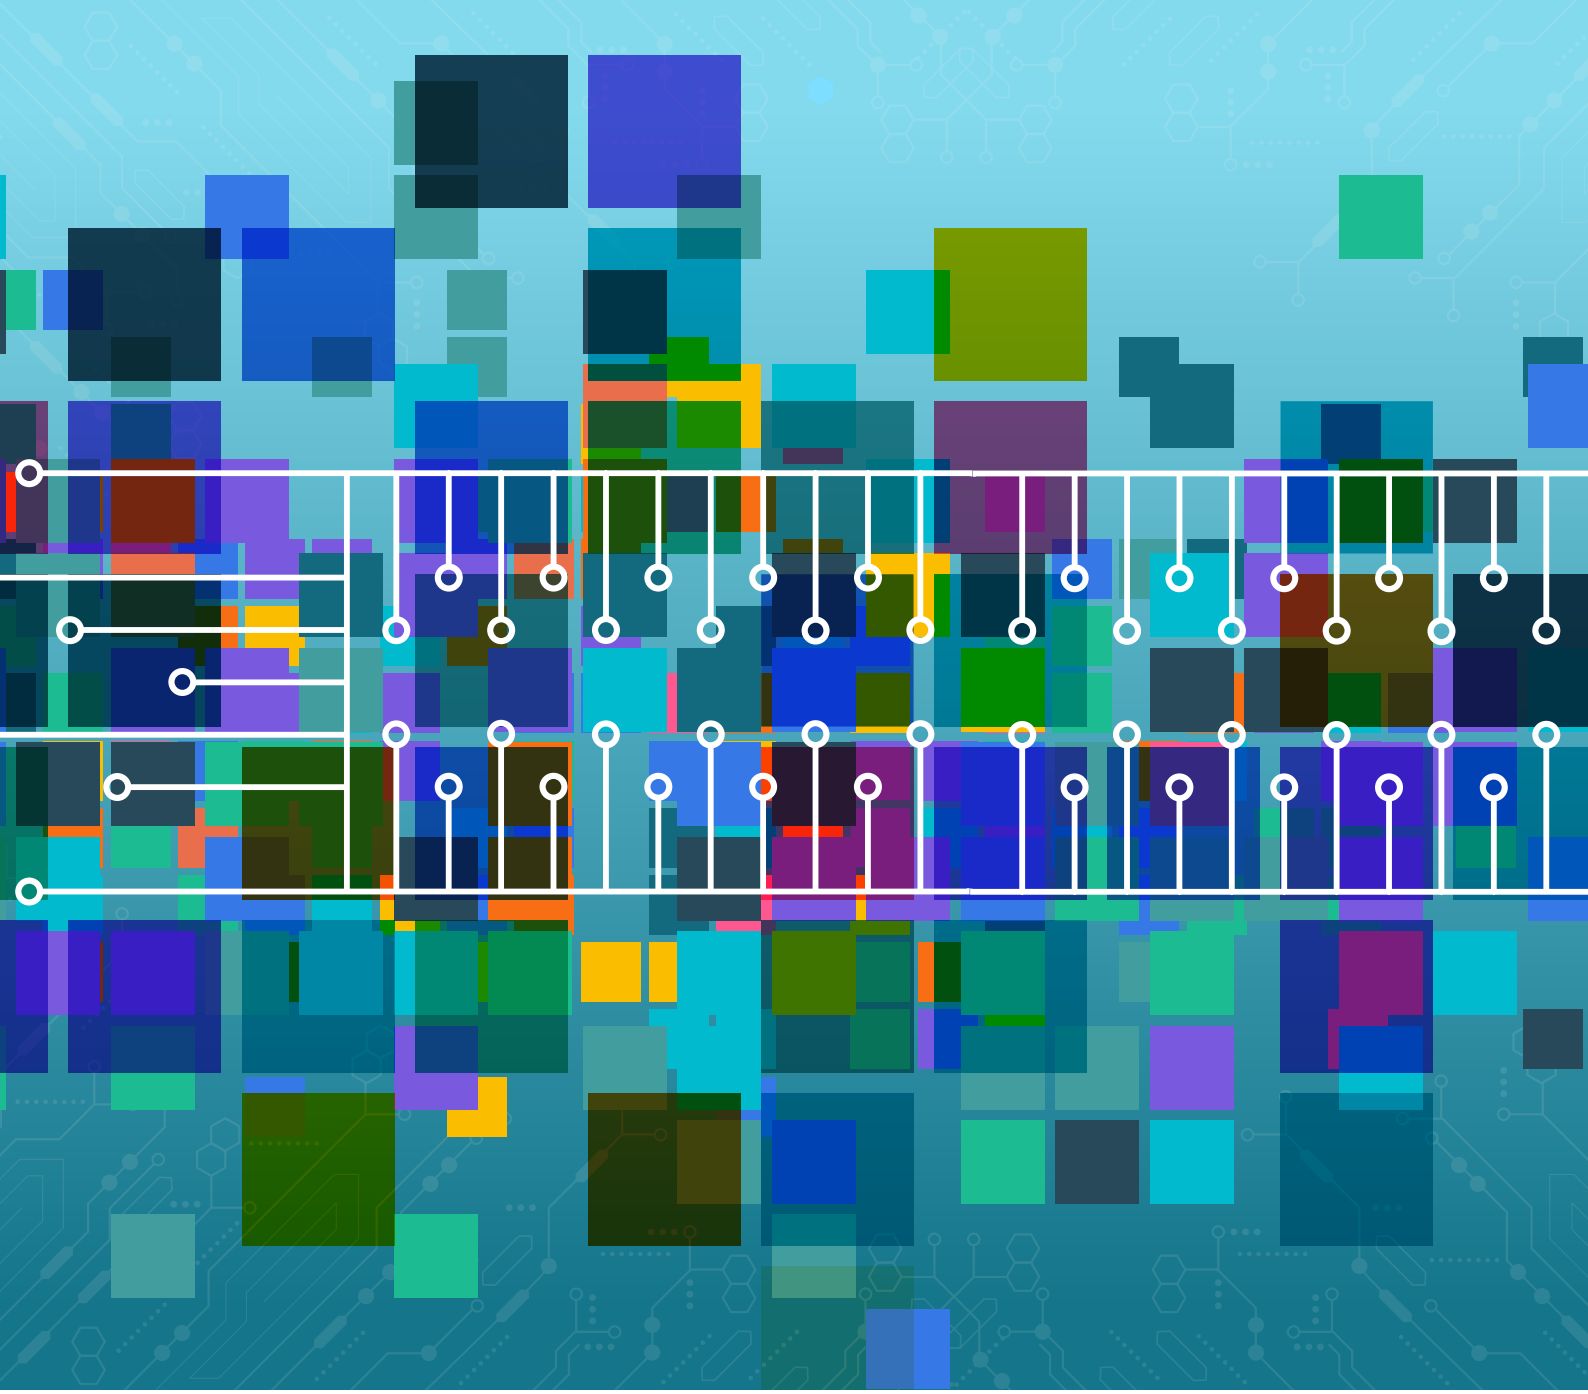

MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com