EUROPEAN CONFERENCE

Conference Proceedings

XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

August 24–26, 2026
Prague, Czech Republic

PRIORITY AREAS OF SCIENTIFIC RESEARCH IN THE MODERN WORLD

Abstracts of XXXIV International Scientific and Practical Conference

Prague, Czech Republic
(August 24–26, 2026)

DATA SOVEREIGNTY AS A MARKETING DIFFERENTIATOR: THE RISE OF PRIVACY-FIRST POSITIONING

Prokopenko Serhii Olexandrovych
Doctor of Philosophy, Senior Lecturer
Simon Kuznets Kharkiv National Economic University, Department of Marketing

ABSTRACT

Over the past several years, data sovereignty has migrated from the vocabulary of regulators and IT architects to that of brand managers. What once denoted a purely technical or jurisdictional condition, where data physically resides, under which legal regime it is processed, and who may compel access to [3,8] it is now routinely deployed as a positioning: a promise addressed not to auditors but to customers.

This paper examines the emergence of "privacy-first" and "sovereign" branding as a marketing differentiator in the technology sector, with particular attention to the wave of European sovereign artificial intelligence and large language model (LLM) initiatives that have intensified public debate about who controls the infrastructure behind everyday digital services.

The paper argues that data sovereignty functions today as a differentiation strategy in the classical sense theorized by positioning and competitive-strategy scholarship, while simultaneously carrying risks familiar from other "virtue label" markets, most notably the risk of privacy-washing.

Keywords: data sovereignty, privacy-first marketing, brand positioning, sovereign AI, trust signaling.

1. INTRODUCTION

Digital markets have entered a period in which the question "where does my data go?" is no longer confined to compliance departments. It has become a question that ordinary consumers, corporate procurement officers, and public-sector buyers ask openly, and one that vendors increasingly answer in their advertising rather than only in their terms of service. This shift has been accelerated by a convergence of three developments: the maturation of data-protection regimes such as the General Data Protection Regulation (GDPR)[6] in the European Union; a series of high-profile disputes over the extraterritorial reach of foreign technology law; and, most recently, the rapid diffusion of large language models, which has made the question of who trains, hosts, and can inspect an organisation's data newly salient at the level of everyday software use rather than abstract infrastructure.

Within this environment, a distinct commercial narrative has taken shape across Europe: the case for "sovereign" artificial intelligence, built, hosted, and governed within a jurisdiction whose legal and political accountability is familiar to its users. This narrative is visible in the emergence of European-headquartered model developers positioning themselves explicitly against non-European incumbents, in national and regional initiatives to build local-language, locally hosted models, and in a broader

public conversation, carried by industry commentators, journalists, and policy analysts about the strategic and even security dimensions of who controls the algorithmic infrastructure a country or company depends on. This paper's central quesion is that this conversation is no longer only technical or geopolitical: it is being actively converted into a marketing asset. “Privacy-first” and “sovereign” have become positioning labels comparable in structural function to “organic,” “fair trade,” or “carbon neutral” in other consumer markets — statements that compress a complex, hard-to-verify property into a simple purchase signal.

We situates the phenomenon within positioning and competitive-differentiation theory. Then we trace the path from regulatory compliance to brand asset, using the European sovereign-AI discourse as its empirical anchor. Later we examines illustrative cases in which data locality and privacy have been foregrounded as commercial differentiators. Next section identifies the mechanisms through which sovereignty claims are signaled to buyers and consider the principal risk this strategy carries — the erosion of trust through unverifiable or exaggerated claims and draws a parallel with sustainability-label markets. Section 7 offers managerial implications, and later conclusions.

The empirical anchor for this discussion is not incidental. Sovereign artificial intelligence has moved, within a short span, from a niche concern of technology policy specialists to a recurring theme in mainstream business and political commentary across the European Union and neighboring states, including Ukraine, where public initiatives to build a national LLM have been explicitly framed in terms of historical, linguistic independence and data control This parallel discourse conducted simultaneously in policy, security, and consumer-facing registers is precisely what allows a single underlying property, jurisdictional control over data and models, to be marketed to a public-sector buyer as a matter of national resilience and to a private consumer as a matter of personal privacy, often within the same communication.

It is worth noting at the outset that the phenomenon under discussion is not confined to any single national market. While the empirical illustrations in this paper are drawn predominantly from the European Union and its immediate periphery, comparable dynamics are visible wherever a jurisdiction perceives its digital infrastructure as dependent on providers outside its legal reach: in discussions of data localization in India and Brazil, in Gulf-state sovereign-cloud initiatives [2], and in the broader international debate over export controls and access restrictions affecting frontier AI models. The European case is nonetheless a particularly instructive one for marketing scholarship because it combines an unusually mature data-protection legal regime, a large internal market capable of sustaining region-specific vendors, and a politically charged narrative of technological independence three conditions that, taken together, create ideal soil for a regulatory property to be converted into a brand asset at scale.

2. THEORETICAL BACKGROUND

Classical positioning theory holds that a brand succeeds by occupying a distinct, defensible place in the buyer's mind relative to competitors, rather than by being objectively superior on every dimension (Ries & Trout) [10]. A differentiator, in this framework, need not be the most important attribute of a product in absolute terms; it

need only be an attribute that a meaningful segment of buyers cares about and that competitors cannot easily replicate. Porter's (1985) treatment of competitive advantage [9] adds the requirement that a differentiator be difficult to imitate quickly, otherwise it collapses into table stakes within a single product cycle.

Data sovereignty satisfies both conditions unusually well for a technical property. First, salience has been manufactured largely by external events as regulatory enforcement actions, high-profile breaches, and geopolitical disputes over cross-border data access rather than by the marketing departments that now exploit it, which lends the claim a credibility that self-generated differentiators often lack. Second, genuine sovereignty is costly to replicate as it requires jurisdictional infrastructure, legal restructuring, and, in the case of AI systems, control over training data and model weights, none of which a competitor can assemble overnight. This combination — externally validated salience plus genuine switching cost explains why sovereignty claims have proven more durable as differentiators than many earlier "privacy" claims, which were frequently reducible to a checkbox in a settings menu.

A second relevant strand of theory concerns trust and control in data relationships. Zuboff's (2019) account of "surveillance capitalism" [12] framed the extraction of behavioral data as a structural feature of the dominant digital business model, one that consumers largely could not opt out of even when informed. Read against that background, sovereignty and privacy-first positioning can be understood as a market response to an accumulated deficit of perceived control: vendors are not merely adding a feature, but explicitly repudiating a business model that buyers have come to associate with incumbents.

A third theoretical resource worth noting is signaling theory as applied to markets with asymmetric information. Where a buyer cannot directly observe a seller's quality, the actual jurisdictional and architectural condition of a data or AI system, sellers have an incentive to send costly, hard-to-fake signals that separate genuine quality from cheap talk. Independent audits, published architecture disclosures, and binding contractual residency clauses function as such costly signals, whereas purely rhetorical claims of "sovereignty" do not, because they cost the seller nothing to make and nothing to retract. This distinction between costly and costless signals reappears throughout the remainder of the paper as the dividing line between durable and fragile sovereignty claims.

3. FROM COMPLIANCE INSTRUMENT TO BRAND ASSET

The GDPR and comparable regimes were designed as compliance instruments, not marketing copy; their vocabulary of lawful bases, data controllers, and cross-border transfer mechanisms was intended for regulators and legal counsel. Yet the last several years have seen this vocabulary translate almost directly into consumer-facing claims — "your data never leaves the EU," "GDPR-native by design," "hosted on European infrastructure" phrases that would have been unusual on a product landing page a decade ago and are now commonplace.

The debate over sovereign large language models has sharpened this translation. Public discussion of national and European LLM initiatives exemplified by France-based Mistral AI's positioning as a European counterweight to United States and

Chinese model providers, and by a wave of national-language, nationally hosted model projects across the region has consistently framed the question of AI infrastructure in terms that map directly onto brand differentiation: who trains the model, where the data resides, whose law governs disclosure requests, and whether the model's guardrails reflect local values and language rather than an imported default. Industry commentators tracking this space have increasingly argued that the choice of AI provider is becoming, for European public institutions and privacy-conscious businesses alike, as much a statement of values and control as a technical procurement decision. That framing is precisely the raw material from which a marketing differentiator is built: a decision criterion that buyers already care about, articulated in language a brand can adopt wholesale.

It is important to note that this translation does not require that the underlying sovereignty claim be false, only that it be simplified. A jurisdictional and architectural condition that is genuinely complex, involving subprocessors, backup regions, model-training pipelines, and the residual reach of foreign legal instruments even over locally incorporated entities is, in marketing use, compressed into a single reassuring label. That compression is what makes the claim usable as a differentiator, and it is also the source of the risks discussed in Section 6.

This translation from regulatory text to marketing copy has been reinforced by a parallel shift in buyer sophistication. Corporate procurement teams, particularly in regulated sectors such as finance, healthcare, and public administration, now routinely include data-residency and AI-governance questions in vendor due diligence checklists, which means that a sovereignty claim is no longer aimed only at end consumers but also at professional buyers who evaluate it against contractual and audit criteria. This professionalization of the buyer side has, somewhat paradoxically, made the marketing claim more valuable rather than less: a vendor that can substantiate its sovereignty claim under professional scrutiny gains a differentiator that survives due diligence, while a vendor that cannot loses the sale outright rather than merely losing some goodwill. The stakes attached to the claim have therefore risen on both sides of the transaction at once.

A related development worth noting is the growing role of sovereignty claims in business-to-business rather than business-to-consumer contexts. Cloud infrastructure providers, payment processors, and enterprise software vendors have increasingly built dedicated “sovereign cloud” or “sovereign AI” product lines, distinct from their global offerings, explicitly designed to satisfy public-sector and regulated-industry procurement requirements. These offerings often carry a price premium relative to the vendor's standard product, which is itself informative: buyers are demonstrably willing to pay more for a jurisdictional guarantee, confirming that the differentiator has real, measurable commercial value rather than functioning merely as costless reassurance. The existence of a price premium also creates a direct incentive for vendors to test how far a sovereignty label can be stretched before the underlying architecture must actually change, which is precisely the tension explored in Section 6.

4. ILLUSTRATIVE CASES

Several categories of firms illustrate how sovereignty and privacy have been operationalized as brand assets rather than background compliance. European AI developers building and marketing models trained and hosted within the European Union have positioned data locality and regulatory alignment as a headline distinction from larger non-European competitors, addressing both public-sector buyers with legal residency requirements and private buyers uneasy about dependence on foreign platforms. Privacy-focused consumer technology providers offering encrypted email, storage, and networking services have built their entire brand identity around the promise that even the provider itself cannot read user data — a claim that functions less as a feature list than as an origin story that distinguishes the firm from advertising-funded incumbents. National and regional public-sector LLM initiatives, while not commercial brands in the conventional sense, perform an analogous signaling function for the state: they communicate to citizens and institutions that a critical piece of digital infrastructure will not depend on a foreign private actor, a claim that has visible political as well as commercial value.

Across these cases, a common pattern emerges: the sovereignty or privacy claim is rarely presented alone. It is typically paired with a named threat — a foreign law, a dominant incumbent, a prior breach — against which the claim is implicitly or explicitly a defense. This pairing strengthens the positioning by giving the buyer a concrete alternative to reject, consistent with the comparative logic of positioning theory.

Further, less commercial but instructive case is the category of national LLM initiatives themselves. Public communication around these projects across the region has consistently emphasized the same cluster of claims found in commercial privacy-first branding: local-language competence, cultural and historical fluency absent from foreign models, and most centrally for this paper's argument assurance that sensitive data used in defense, healthcare, and financial applications will not leave the country's jurisdiction.

That a state, rather than a private firm, deploys this positioning vocabulary is itself evidence of how thoroughly the language of sovereignty-as-differentiator has spread beyond its original commercial context; the state is, in effect, competing for the trust of its own citizens and institutions against foreign AI providers using the identical rhetorical toolkit a private brand would use against a commercial rival.

5. MECHANISMS OF SIGNALLING

Firms operationalize sovereignty claims through several recurring mechanisms. Architectural disclosure involves publishing, in varying degrees of technical detail, where data is stored and processed, with the aim of making an invisible property visible and checkable. Third-party certification and audit reports derive their legitimacy from independent verification, functioning much like eco-labels in sustainability marketing. Contractual and jurisdictional guarantees, such as data-residency clauses and explicit statements about which courts have jurisdiction over disputes, translate the claim into a form that is enforceable by the buyer rather than merely promised by the seller. Finally, symbolic association with national or regional identity as flags, language,

references to local regulatory bodies signals sovereignty through cultural rather than technical means, a mechanism especially visible in the branding of national LLM projects[11].

These mechanisms differ sharply in verifiability. Architectural disclosure and independent audits are, in principle, checkable by a sufficiently motivated buyer or journalist. Symbolic and rhetorical signals are not, which creates the conditions for the divergence between claimed and actual sovereignty discussed next.

A further mechanism worth distinguishing is negative signaling: a claim built by explicit contrast with a named category of competitor, rather than by positive description of one's own architecture. Statements such as “unlike other providers, we do not train on your data” or “your prompts never leave the country” derive their persuasive force from an implied comparison that the buyer must complete unaided. This mechanism is rhetorically powerful precisely because it requires no independent evidence about the speaker's own practices, only a plausible accusation against an unnamed or lightly named competitor which makes it, in signaling terms, among the least costly and therefore least reliable of the mechanisms surveyed here, even though it is frequently the most prominent in advertising copy.

A final signaling mechanism deserving separate mention is temporal signaling: the practice of merging a sovereignty commitment to a specific past event — a regulatory investigation, a competitor's breach, a change in a foreign government's surveillance law rather than presenting it as a timeless feature.

Anchoring the claim to a datable event lends it a narrative structure that buyers find easier to recall and repeat to others, which partly explains why sovereignty branding so often takes the form of a founding story (“we were built after event X because we believed Y”) rather than a neutral specification sheet.

This narrative framing borrows directly from origin-story techniques long used in consumer branding generally, but it acquires particular force in the sovereignty context because the triggering event is typically one the buyer already remembers independently, which corroborates the brand's account rather than asking the buyer to take it on faith.

6. RISKS: PRIVACY-WASHING AND THE LIMITS OF THE CLAIM

The sustainability marketing literature offers a useful cautionary parallel. “Greenwashing” [5] became a recognized risk once environmental claims proliferated faster than the verification infrastructure needed to substantiate them, producing a gap between labeled and actual environmental performance that ultimately undermined consumer trust in the label category as a whole, not merely in individual offending firms. Data sovereignty claims are structurally exposed to the same dynamic. A vendor may host primary data within a given jurisdiction while relying on subprocessors, backup infrastructure, or model-training pipelines located elsewhere; may describe a service as “sovereign” while remaining a subsidiary of a foreign parent company subject to that parent's home-country legal obligations; or may market local hosting for one product line while a shared back-end analytics or AI feature routes data outside the claimed jurisdiction.

This risk is compounded in the AI context by the genuine technical difficulty of verifying training-data provenance and model governance from the outside. A model can be described as "trained on European data" or "governed under European law" in ways that are difficult for an ordinary buyer, or even a sophisticated procurement officer, to audit directly. Where the underlying claim resists independent verification, the differentiator risks becoming a matter of assertion rather than demonstrated fact, and the category as a whole becomes vulnerable to the same erosion of credibility that has periodically affected other virtue-signaling label markets. For firms and public bodies genuinely investing in sovereign infrastructure, this creates a collective-action problem: bad-faith or exaggerated claims by some actors raise the skepticism cost borne by all.

The parallel with sustainability labeling also suggests how the market may self-correct over time. In environmental marketing, the proliferation of unverifiable claims eventually produced regulatory intervention — mandatory substantiation requirements, standardized labels, and enforcement actions against misleading claims, alongside market-driven responses such as third-party certification schemes that firms adopted voluntarily to differentiate themselves from less scrupulous competitors. There is early evidence of an analogous pattern forming around data-sovereignty claims, visible in the emergence of specialized certification schemes for cloud and AI sovereignty and in regulatory proposals that would require vendors to substantiate residency and governance claims in a manner comparable to environmental disclosure requirements. Firms that adopt costly, verifiable signaling mechanisms early are likely to benefit disproportionately once such standardization arrives, since their existing practices will already satisfy requirements that late adopters will have to build from scratch under greater public scrutiny.

It is also worth considering the demand side of this market failure. Individual consumers rarely have the expertise or the incentive to audit a vendor's subprocessor list or model-training pipeline, which means that the disciplining function typically performed by informed buyers in competitive markets. In this category, audit performed instead by a small number of specialised intermediaries: investigative journalists, security researchers, competing vendors with an incentive to expose rivals, and, increasingly, regulators equipped with audit powers the ordinary consumer lacks. The credibility of the entire "privacy-first" category is therefore disproportionately dependent on the continued activity of these intermediaries, a dependency that marketing scholarship on credence goods [4] — products whose quality cannot be verified even after purchase has long identified as a structural vulnerability rather than an incidental one.

7. MANAGERIAL IMPLICATIONS

For firms considering data sovereignty as a positioning strategy, three implications follow from the analysis above. First, verifiability should be treated as a design requirement of the marketing claim itself, not an afterthought: claims paired with independently auditable evidence — published architecture diagrams, third-party certification, contractual residency guarantees — are more defensible and more durable than rhetorical or symbolic claims alone.

Second, the scope of the claim should be stated precisely rather than in the abstract; a claim limited to "primary customer data processed within the EU" is more defensible under scrutiny than an unqualified "fully sovereign" claim that a single subprocessor or backup routine can falsify. Third, firms operating in the AI sector should anticipate that buyers, journalists, and regulators will increasingly extend sovereignty scrutiny from data storage to model training and governance, and should prepare disclosure practices accordingly before, rather than after, the claim is publicly challenged.

A fourth implication concerns internal alignment between marketing and engineering functions. Because sovereignty and privacy claims are unusually easy to falsify through a single overlooked subprocessor, backup routine, or shared analytics pipeline, the organizations most exposed to reputational risk are precisely those in which marketing teams issue claims that engineering and legal teams have not formally reviewed and signed off against the actual architecture. Establishing a standing internal review process in effect, an internal audit trail for every sovereignty-related claim before it is published is a comparatively low-cost safeguard against the kind of public correction that, once it occurs, tends to attach to the entire sovereignty-marketing category rather than remaining contained to the offending firm.

8. CONCLUSION

Data sovereignty has completed a migration familiar from other regulatory-to-commercial transitions: a property first defined by law and infrastructure has become a differentiator defined by brand strategy. The European debate over sovereign artificial intelligence and national LLM initiatives has served as both a cause and evidence of this migration, supplying firms with a ready-made vocabulary for jurisdictional control that maps cleanly onto positioning theory's requirements of salience and inimitability. The strategy is commercially powerful precisely because the underlying property is genuinely costly to replicate — but that same complexity makes it difficult for buyers to verify the claim, exposing the category to the privacy-washing risk that has previously undermined other virtue-label markets.

Firms and public institutions that pair sovereignty claims with precise scope and independently verifiable evidence are best positioned to capture the differentiator's commercial value while avoiding the credibility erosion that awaits the category if verification continues to lag behind the marketing built upon it.

References

1. Autolitano, S., & Pawlowska, A. (2021). Europe's quest for digital sovereignty: GAIA-X as a case study (IAI Papers 21). Istituto Affari Internazionali. https://www.iai.it/en/publications/c03/europes-quest-digital-sovereignty-gaia-x-case-study

2. Couture, S., & Toupin, S. (2019). What does the notion of "sovereignty" mean when referring to the digital? New Media & Society, 21(10), 2305-2322. https://doi.org/10.1177/1461444819865984

3. Darby, M. R., & Karni, E. (1973). Free competition and the optimal amount of fraud. Journal of Law and Economics, 16(1), 67–88. https://doi.org/10.1086/466756

4. de Freitas Netto, S. V., Sobral, M. F. F., Ribeiro, A. R. B., & Soares, G. R. da L. (2020). Concepts and forms of greenwashing: A systematic review. Environmental Sciences Europe, 32(1), Article 19. https://doi.org/10.1186/s12302-020-0300-3

5. European Parliament and Council. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). Official Journal of the European Union, L119, 1–88.

6. Hummel, P., Braun, M., Tretter, M., & Dabrock, P. (2021). Data sovereignty: A review. Big Data & Society, 8(1). https://doi.org/10.1177/2053951720982012

7. Porter, M. E. (1985). Competitive advantage: Creating and sustaining superior performance. Free Press.

8. Ries, A., & Trout, J. (2001). Positioning: The battle for your mind (20th anniversary ed.). McGraw-Hill. (Original work published 1981)

9. Anthropic. (2026). Claude Sonnet 5 (Aug 2026 version) [Large language model]. https://claude.ai

10. Gatrell, C., Muzio, D., Post, C., & Wickert, C. (2024). Here, there and everywhere: On the responsible use of artificial intelligence (AI) in management research and the peer-review process. Journal of Management Studies, 61(3), 739–751. https://doi.org/10.1111/joms.13045

11. Spence, M. (1973). Job market signaling. The Quarterly Journal of Economics, 87(3), 355–374. https://doi.org/10.2307/1882010

12. Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.

Scientific publications

MATERIALS

The XXXIV International Scientific and Practical Conference
«Priority areas of scientific research in the modern world»

Prague, Czech Republic
(August 24–26, 2026)