# DATA GOVERNANCE & SOVEREIGNTY

A PRACTITIONER'S GUIDE TO THE ARCHITECTURE OF EXTRACTION

A **MADE** BRIEF

2026

**MADE**
Mass Atrocities in the Digital Era
**Yale** MACMILLAN CENTER
*Genocide Studies Program*

# Table of Contents

# Introduction

Control over data is of the utmost importance to people and countries around the world, given the immense power that data, and the control over it, carries. The rules and authority that decide who gets that control are what we refer to as data governance. This brief argues that data governance is the foundational layer of digital harm. Before AI bias, platform manipulation, or synthetic media can be meaningfully addressed, the issue of who owns, stores, processes, and profits from a population's data, and under whose legal framework, must be resolved. Regulating outputs without regulating inputs treats the symptoms while leaving the infrastructure of harm intact.

The brief also examines the largest regulatory regimes attempting to address this problem, the American, the European, and the Chinese, and finds that they differ in form but share a common failure: none resolves the underlying question of structural power over data, and none is as independent of the others as it appears. For atrocity prevention practitioners, the implication is that the downstream harms this field already monitors, from bias and homogenization to deception and manipulation, each depend on an underlying data architecture that remains largely unregulated.

# Key Terms

**Data governance** refers to the delegation of authority over data. In other words, who decides how information is collected, stored, processed, and deployed.

**Digital sovereignty** refers to the capacity of a state and its population to exercise meaningful control over the data, infrastructure, and governance norms that shape their digital lives.

**Data sovereignty** is the more specific question of who owns, stores, processes, and profits from a population's data, and under whose legal framework.

Both digital sovereignty and data sovereignty should be distinguished from cybersecurity (the protection of systems from intrusion) and internet freedom (individual access to information online), as the issue here is structural power over data governance itself. This brief examines the extractive logic that operates wherever data governance privileges the interests of those who control digital infrastructure over the interests of those who live within it.

# Enabled through Data Architecture

Every mechanism of harm examined in this brief depends on an underlying data architecture. Perhaps the clearest historical illustration is IBM's census tabulation machinery, which enabled the Nazi regime to identify, sort, and target Jewish populations with unprecedented efficiency. What made industrial-scale targeting operationally possible was the prior control the Nazi regime had over the population's data. Parameters such as religious affiliation, residential address, and even occupation were rendered legible and queryable by tabulating infrastructure. Then when prisoners were taken to concentration camps, the numbers tattooed on internees corresponded to individualized IBM punch cards. When a government or corporation controls granular data about a population's identity, movement, health, finances, and political expression, it holds the prerequisite infrastructure for targeted persecution, whether or not it currently intends to use it that way.

A contemporary counterpart is happening in Xinjiang. China's Integrated Joint Operations Platform aggregates data from police checkpoints, facial recognition cameras, banking records, and mandatory biometric collection, and flags Uyghurs and other Turkic Muslims for interrogation or detention on the basis of behaviors as ordinary as using a back door or fueling someone else's car. That infrastructure underpins a mass internment system that the UN human rights office concluded may constitute crimes against humanity.

The pattern is not strictly confined to authoritarian states: India's Aadhaar biometric ID system, introduced as voluntary, has become functionally mandatory, linking welfare, tax, and security databases into near-total legibility of citizens' movements and transactions. Separated by eighty years, IBM's punch card and the Indian and Chinese governance platform make the same point. Once a population's data is consolidated and queryable, the barriers to carrying out persecution collapse.

# The Harms of Unregulated Data

These cases show what consolidated data makes possible once a state decides to act. But harm does not wait for a government to weaponize a database. Unregulated data produces four recurring harms, ordered here from the broadest to the most individually targeted:

1. **Homogenization** occurs when systems optimized for dominant patterns flatten cultural expression and suppress minority dialects, identities, and perspectives. Because the optimization logic privileges majority patterns, the effect reinforces itself over time, and marginalized communities gradually become invisible to the systems that govern them, a precondition for targeting

2. **Deception** operates through "synthetic authority", the tendency of AI systems to produce credible-seeming content and shift trust from human expertise to automated systems. As fabricated material floods information channels, it produces the "liar's dividend". Once the public knows convincing fakes exist, perpetrators can dismiss authentic evidence as fabricated while fabricated realities circulate as true, so that atrocities that happened gain plausible deniability and atrocities that never happened gain evidence.

3. **Manipulation** happens when engagement-optimized algorithms privilege the most emotionally resonant version of a narrative, so false news spreads faster and further than true news. Manipulation requires no coordinated bad actor, because the system facilitates it on its own. The same logic that incentivizes outrage disproportionately amplifies dehumanizing narratives and conspiracy theories, since to the algorithm a rage-click and a curiosity-click are equivalent. Algorithmic radicalization accelerates the outcome. Where propaganda through print or radio historically took years to prime a population for violence, platform dynamics compress the process into months or days. SCL Group and Cambridge Analytica leveraged exactly this commercial infrastructure, deploying psychographic micro-targeting across more than 100 election campaigns in at least 30 countries (later leaks suggest operations in as many as 68 countries), including ethnically targeted voter suppression in Trinidad and Tobago and Nigeria. These dynamics have already fed mass violence. In Myanmar, anti-Rohingya and inflammatory content was repeatedly amplified on Facebook, where the platform had become synonymous with the internet itself, helping legitimize and accelerate real-world killing.

# The Harms of Unregulated Data

4. **Bias** emerges when AI systems trained on historical data reproduce structural inequality at the speed and scale of automation. The 2019 NIST Face Recognition Vendor Test found false positive rates ten to one hundred times higher for Black, East Asian, and Native American faces than for white faces; a risk-prediction algorithm applied to 200 million U.S. patients systematically underestimated the medical needs of Black patients because it used healthcare cost as a proxy for illness. When disparate-impact safeguards are removed, routine administrative categories become the categories of discrimination, or even of persecution.

The harms above compound upon themselves. Homogenization renders marginalized communities illegible to governance systems, deception collapses the public's ability to distinguish an atrocity from a fabrication of one, manipulation compresses the social distance between a population and the conditions under which violence becomes possible, and bias makes exclusion the default output of automated administration. Each of these is dangerous alone, but layered together on infrastructure the affected population does not control, they constitute the structural preconditions for mass violence.

# Efforts to Regulate Data

Regulatory frameworks allocate structural power over data, and most allocate it away from the populations the data describes. Nor do these efforts operate independently of one another. The U.S. policy apparatus that strips domestic safeguards simultaneously engineers the international frameworks that enable extraction abroad. Executive Order 14,179's promise of "global AI dominance" is just as much of a domestic posture as it is an international one, shaping how U.S. firms operate in Lagos and Nairobi just as it shapes how they operate in Detroit and El Paso. This brief calls the domestic pattern algorithmic ethnonationalism, in which state data infrastructure is oriented toward an exclusionary political project, and the international pattern digital colonialism, in which foreign firms and governments extract a population's data and control the infrastructure that population depends on. The conditions that produce one are the conditions that produce the other.

# The United States

Ethnonationalism, the ideology that limits full political membership to a dominant ethnic or cultural group and treats demographic diversity as a threat, has recurred throughout U.S. history, but today's backlash is equipped with computational infrastructure that earlier moments lacked.

As legal scholar Spencer Overton argues, amid rapid demographic change and cultural anxiety, federal AI policy has become a critical arena for that ideology. The federal government has assembled a vast data infrastructure across core state functions, and AI is now being layered onto it. DHS operates biometric and identity systems consolidating records on hundreds of millions of people, including the HART database; ICE contracts extensively with Palantir for targeting tools that aggregate disparate federal datasets into operational queries; the IRS uses algorithmic risk scoring to flag returns for audit, with documented disparate impact on lower-income Black households; the DOJ and FBI run facial recognition searches across state DMV photo repositories; and the GSA coordinates AI procurement government-wide.

Even nominally voluntary systems coerce. Declining the TSA's "optional" facial recognition scan means longer lines and secondary screening, normalizing surveillance through inconvenience. *Algorithmic ethnonationalism* occurs when this infrastructure is oriented toward an ethnonationalist political project. The policy question that follows is how to govern the specific data systems that determine who is rendered visible, sortable, and actionable to the state.

Recent policy shifts have removed the safeguards that would constrain that orientation. Executive Order 14,148 (January 2025) repealed prior AI safety and civil-rights protections and EO 14,179 from the same month reframed AI policy around global technological dominance, recasting equity considerations as opposed to innovation. Later federal efforts such as EO 14,319 ("Preventing Woke AI," July 2025) prohibited federal procurement of bias-mitigating AI systems, and EO 14,365 (December 2025) threatened funding penalties for states pursuing AI bias regulation. In parallel, A directive from the National Institute of Standards and Technology (NIST) removed references to AI fairness from its frameworks, while the Equal Employment Opportunity Commission (EEOC) and Department of Labor (DOL) withdrew guidance on algorithmic hiring bias. The combined effect has been to dismantle the mechanisms that made inclusion enforceable. With disparate impact review, fairness standards, and hiring bias guidance withdrawn, exclusion becomes the default output in domains such as hiring, surveillance, and public service delivery.

# The United States

The same apparatus operates outward, and because most companies with access to large-scale data are U.S. companies, U.S. trade policy functions as de facto global data policy. The Office of the U.S. Trade Representative has flagged data localization measures in Nigeria and Kenya as "barriers to digital trade" in its annual National Trade Estimate reports, framing local sovereignty over citizen data as a problem to be negotiated away. The African Growth and Opportunity Act conditions duty-free access to U.S. markets on "eliminating barriers to U.S. trade and investment," a category broad enough to encompass data governance frameworks, while the U.S. International Development Finance Corporation provides roughly $9.3 billion annually in development financing with no equivalent conditioning on digital regulatory practice. U.S. platforms also export an ideological frame of privacy as individual consent and data as personal property, assumptions that misfit contexts where data harms are collective rather than individual. The legal frameworks treat their extractive practices as default and data sovereignty as the exception.

# Europe

The European Union demonstrates that a government can be both willing and able to protect data. The General Data Protection Regulation (2016) reaches any company processing Europeans' data regardless of where that company sits, and backs that reach with penalties of up to 4% of global turnover. In practice, compliance means a company must have a lawful basis for collecting personal data at all, use it only for the purpose it was collected for, honor individuals' rights to access, correct, and delete their data, report breaches within 72 hours, and demonstrate adequate protections before moving Europeans' data out of the EU. The Irish privacy regulator's €1.2 billion fine against Meta in 2023 remains the largest data protection penalty issued anywhere. No other jurisdiction has forced global platforms to change practices at comparable scale.

But the European regime has two distinct limitations. First, its newer instruments bet on creating a transparency regime, rather than building structural privacy, in hopes of behavior change from companies. The Digital Services Act obliges the largest platforms on the continent to publish how their recommendation algorithms work, assess the systemic risks their services create, open their data to vetted researchers, explain content removal decisions to affected users, and submit to independent audits. Every one of these is an obligation to show rather than to change, documenting the extractive system while leaving untouched who owns it and who profits from the data it consumes.

Second, and more fundamentally, rights-based regulation stops at infrastructure ownership. In the case now more commonly known as Schrems II, the EU's highest court struck down the EU–US Privacy Shield because data held by U.S. companies remains reachable by U.S. surveillance law no matter where it is stored, the same reasoning that invalidated the Safe Harbor framework in 2015 and the same challenge now facing the 2023 EU–US Data Privacy Framework. The CLOUD Act makes the point explicit, since U.S. providers can be compelled to produce data held on European soil. Both rulings mark the same boundary. A court can strike down the legal mechanism for moving data. It cannot relocate the servers, and it cannot change who owns them.

The dependence is structural as well, with U.S. hyperscalers (the handful of cloud providers such as Amazon, Microsoft, and Google that operate data center networks at global scale) holding roughly 85% of the European cloud market, and the EU's flagship sovereign-cloud initiative, GAIA-X, has struggled to gain traction.

Europe therefore marks the ceiling of the regulatory strategy, because the strongest data protection law in the world cannot manufacture sovereignty over infrastructure it does not own. Where data physically sits, and who owns the systems it sits on, is where the sovereignty question is actually decided- which is the subject of the next section - is not addressed

Spain's experience with data regulation illustrates the gap between recognition and capacity. Foreign Minister José Manuel Albares has stated that digital sovereignty is a top issue for Spain, which is actively working to combat disinformation and electoral interference and to reduce dependence on non-EU companies. However, this commitment is stalled at the same checkpoint the EU is hitting as a whole. The infrastructure Spain would regulate belongs to the U.S. hyperscalers that dominate the European cloud market, so the disinformation flows it wants to combat and the dependencies it wants to unwind are governed by decisions made in Silicon Valley under U.S. law. The rhetorical commitment is in place. The structural capacity to fulfill it sits offshore, beyond the reach of Spanish law. Even inside the world's most developed data-protection regime, populations experience the consequences of norms they had no vote in importing.

# Global North-Based Regulation: The Physical Location of Infrastructure

Regulatory regimes of whatever origin share the Western assumption that companies have a meaningful presence in the jurisdictions where they operate. Take for example a Coca-Cola bottling plant in Tanzania that occupies local land, employs local workers, and is subject to Tanzanian law as a matter of practical necessity; **a social media platform** serving millions of Tanzanian users from servers abroad **operates under no equivalent constraint**, and its compliance is voluntary, contingent on its own assessment of reputational risk. Data generated by African populations is predominantly routed to data centers in Europe, the United States, or China, where it is stored, processed, and monetized beyond the reach of the governments nominally responsible for the people who generated it. When a country's voter registration database sits on foreign servers, or its health ministry depends on a foreign cloud provider, sovereignty over population data becomes more formal than real.

# The dependency extends into core state functions.

Outsourcing sovereignty refers to the shift in responsibility for governmental functions to private contractors, which moves decision-making power from public to private hands. However, it is worth considering that when that contractor is a foreign entity—the power shifts beyond the reach of domestic law entirely. In 2025, Kenya signed a five-year, $2.5 billion health agreement with the United States that critics feared would give the U.S. access to health databases containing sensitive personal medical records. In 2017, Kenya contracted French firm OT-Morpho (now IDEMIA) for electronic voter identification and results transmission; on election day, configuration failures and suspected hackers produced manipulated results naming the incumbent the winner. The opposition called the election rigged; OT-Morpho denied tampering; and when lawyers demanded the servers be opened for inspection, the electoral commission refused, citing French privacy law and the servers' location abroad. The same year, Rwanda's utilities regulator fined MTN Rwanda $8.5 million for hosting its IT services in Uganda and forced their repatriation, a complex, expensive, and disruptive undertaking that most African states lack the infrastructure to replicate.

## Infrastructure dependency is also weaponizable during political crisis.

Infrastructure dependency is also weaponizable during political crisis. Nigeria disrupted internet access nationwide during the 2024 #EndBadGovernance protests to slow the spread of information; Ethiopia cut social media and communications in the Amhara region from 2023 until July 2024 during fighting with Fano fighters; India, which conducts more internet shutdowns than any other country, has made recurring blackouts in Kashmir a routine instrument of control; Sudan experienced a nationwide shutdown in February 2024 amid civil war, leaving millions unable to contact family, receive money, or reach humanitarian services; and Tanzania shut down the internet following its 2025 elections to silence dissent.

# The China Variable

The Chinese model of digital diplomacy is less of an alternative than it is a separate, but parallel approach to the U.S. model.

In many African contexts, it serves as another vehicle for extractive digital governance. Foreign firms arrive with a similar logic of data capture, monetization, and infrastructural dependence, even if under a different legal and geopolitical framework. Large-scale infrastructure financing, such as ZTE's landmark deal with Ethiopia's state telecommunications provider, built the material backbone, and Chinese firms have supplied surveillance systems to numerous African states under Belt and Road–era partnerships, exporting components of an apparatus refined domestically in Xinjiang. Even Myanmar's junta already relies on Chinese-built surveillance networks. Thus, framing the issue as a U.S.–China contest therefore obscures more than it reveals.

The problem is architectural, because once data governance is subordinated to extraction, origin matters less than operational logic. India's MOSIP identity platform now underpins national ID systems from Sri Lanka to Ethiopia and the Togolese Republic, and African firms themselves may adopt the same extractive assumptions in pursuit of scale and profit. The practitioner's analytical task is to assess whether the governing model erodes local control, concentrates decision-making power externally, and transforms populations into raw material for optimization, regardless of which flag is attached to the platform.

# How Structural Power over Data Governance Creates Atrocity Risk

Sovereignty erosion locks the four harms of homogenization, deception, manipulation, and bias in place. When a population's data is stored, processed, and governed outside domestic legal frameworks, that population cannot mitigate homogenization, deception, manipulation, or bias, because mitigation requires exactly the structural power over data governance that has been extracted. Loss of data sovereignty is loss of self-governance, and with it, loss of the capacity for self-protection.

# How to Establish Data Sovereignty

Establishing data sovereignty requires assembling four elements, each of them visible in the counter-models this brief has already documented.

## Legal authority

The starting point is a data protection law that treats protection over data as the norm. Nigeria's 2023 Data Protection Act prohibits cross-border transfers unless adequate protections are demonstrated, reversing the burden that most legal regimes place on the population. The African Union Data Policy Framework offers a template, and indigenous data sovereignty frameworks add the element most Western statutes omit, which is rights held collectively by the community the data describes rather than only individually by consenting users.

## Physical infrastructure

Law without infrastructure produces the Kenyan election problem of formal authority over data that physically sits beyond reach. Senegal's SENIX exchange point keeps local traffic onshore; Côte d'Ivoire's DataConnect Africa provides sovereign cloud capacity. But localization is complex, expensive, and disruptive, and most states lack the infrastructure to do it alone. This is why the regional route, with shared infrastructure and collective negotiation, matters. Bargaining over digital terms as a bloc through AfCFTA converts many small markets, each individually dispensable to a platform, into one market too large to ignore.

# How to Establish Data Sovereignty

### Enforcement capacity

A statute and a data center still require a regulator willing and able to act. Kenya's Office of the Data Protection Commissioner ordered Worldcoin to halt its iris-scan collection and was ignored until a more powerful ministry, like the Ministry of Interior, who had physical leverage, intervened. The lesson from this episode is two-fold: Independent regulators can act faster than legislatures, but without resourcing and political backing their orders are advisory.

### Collective governance, and a warning

Data sovereignty for the state is not the same as data sovereignty for the population. A government that localizes its citizens' data while shutting down their internet during protests, as this brief's own case studies show, has concentrated power rather than distributed it. The test practitioners should apply mirrors the one this brief proposes for foreign providers. What matters is whether the governance of that data shifts structural power toward the population it describes, and keeping data inside the border does not by itself accomplish that. Indigenous data governance models are the clearest articulation of that standard, which is why they anchor the recommendations that follow.

# Counter-Model: Indigenous Data Sovereignty

*Indigenous sovereignty* refers to the right of indigenous groups to self-determination and ownership of their own land, knowledge, and heritage, as opposed to external exploitation of those resources. Indigenous data sovereignty frameworks assert that this right extends to authority over the collection, ownership, and application of community data, a right reaffirmed in the UN Declaration on the Rights of Indigenous Peoples. Where this brief documents state and corporate claims to community data, indigenous modes of governance offer the clearest alternative, treating data as collective patrimony that remains with the communities it comes from.

In practice, the African Union Data Policy Framework, which positions data as a strategic asset, offers a meaningful starting point, and several states have already translated its principles into policy. Nigeria's 2023 Data Protection Act prohibits cross-border data transfers by default; Senegal's SENIX internet exchange point keeps local traffic onshore; Côte d'Ivoire's DataConnect Africa launched a sovereign cloud platform in 2024. Enforcement matters too. When Worldcoin offered Kenyans 8,000 shillings each in exchange for iris scans in 2023, Kenya's Office of the Data Protection Commissioner ordered a halt to the predatory collection, an order Worldcoin ignored until the Ministry of Interior suspended its operations.

# Warning Signs & Recommendations

The central argument of this brief is that data governance is the prerequisite to regulating any technology built on top of it. The recommendations below follow directly from that premise, organizing what monitoring, advocacy, and capacity-building look like when data infrastructure is treated as the foundational layer. They are not exhaustive, and they should be read alongside existing frameworks on platform accountability, AI governance, and atrocity prevention rather than as a substitute for them. Their function is to keep data infrastructure visible as a site of intervention, particularly where harms get framed as downstream problems of "AI bias" or "platform manipulation" without attention to the data architectures that make them possible.

**Monitoring:** Monitoring means treating data infrastructure as an early warning indicator in its own right. The task is to track who is acquiring control over a population's data before that control is exercised.

- Track bilateral data agreements and trade provisions affecting data sovereignty in at-risk regions
- Monitor biometric data collection by foreign entities under information asymmetry
- Document when state functions become dependent on foreign data infrastructure
- Track AI procurement in government, especially immigration, policing, and benefits
- Monitor rollbacks of algorithmic accountability requirements
- Identify dominant platforms in at-risk regions and languages, and their moderation capacity

# Warning Signs & Recommendations

**Advocacy:** Advocacy entails pushing the governance of data toward the populations it describes, in legislation, in trade negotiations, and in platform obligations.

- Advocate for regional data governance frameworks, with AfCFTA as a model for negotiating digital terms as a bloc
- Push for data localization, portability provisions, and researcher access to platform data
- Advocate for interoperability mandates so users and states are not locked into extractive infrastructure

**Capacity-Building:** Capacity building looks like closing the gap between formal authority and practical control, so that states and communities can exercise the sovereignty their laws already claim.

- Support domestic data processing infrastructure
- Build civil society capacity to monitor and contest data sovereignty violations
- Support indigenous data governance initiatives that assert collective rights over community data
- Help communities understand how data extraction and bilateral agreements affect their rights

Unprotected data sovereignty is a standing vulnerability. The power structure of the world is shifting quickly, and control over populations' data is shifting with it, toward whichever states and firms can extract it fastest. In that environment, the populations least able to contest extraction face greater risk than ever, because every downstream harm the atrocity prevention field monitors, from predictive policing to platform-driven incitement, runs on data someone else controls. The recommendations above are written so that practitioners can see those risks while they are still forming, at the level of infrastructure, where prevention is still possible.

# About MADE

The Mass Atrocities in the Digital Era (MADE) initiative is a first-of-its-kind program within Yale University's Genocide Studies Program, created to examine how digital technologies are transforming the landscape of mass violence and protection. MADE studies how information systems shape the commission, prevention, and accountability of atrocities, while cultivating a new generation of scholars and practitioners equipped to navigate these challenges. The initiative's 2025 programming tackles these issues through four core modules — digital authoritarianism and extremism, platform dynamics and violence, militarization of technology, and digital infrastructures and identity — each producing practitioner-oriented primers that translate complex research into accessible, action-focused tools for early detection, prevention, and accountability.

**MADE**

Mass Atrocities in the Digital Era

**Yale** MACMILLAN CENTER

*Genocide Studies Program*