permuta
www.permuta.com

You Can't permission your way out of Surveillance at Scale

Palantir's institutional sovereignty pitch is real. So is the depth, sensitivity, and reach of the citizen data it now touches across allied governments.

Martin Jennings, General Manager, Product Group, Permuta Technologies *10 Jul 2026*

Executive Summary

Palantir recently published its own white paper, Institutional Sovereignty in the Age of AI, advising governments and enterprises to protect themselves from AI model providers through zero data retention agreements, model-agnostic architecture, granular permissions, and continuous audit logging. The paper's core warning is that frontier model providers have a structural incentive to extract an institution's tribal knowledge into their own model weights, and that institutions must therefore treat every non-sovereign model relationship as extraction-prone.

That warning is worth taking seriously. But it is incomplete in a way that matters a great deal to the government customers Palantir serves. The same white paper describes, approvingly, the mechanism by which an institution can model itself as a "digital twin," with permissions and an ontology that reflect the full structure of the organization down to the people, roles, and relationships inside it. That is not only a defensive control. It is also a description of the exact capability that lets a single vendor's platform fuse previously siloed government datasets, across immigration, health, tax, and defense systems, into comprehensive records on individual human beings. Palantir has built that capability at a scale, a depth, and a sensitivity level that Google's advertising business, the original commercial engine of what Shoshana Zuboff named "surveillance capitalism," never approached.

This paper argues that AI-enabled, government-facing data fusion platforms represent the next evolution of what Google started when it began treating user behavior as raw material for prediction products. The difference this time is scale, depth, and consequence. Google turned browsing habits and email content into advertising signal. Palantir's Gotham, Foundry, and AIP stack now touch immigration case files, tax records, health data, and military targeting systems, at contract values that have grown from a $480 million Army prototype in May 2024 to a $10 billion Army enterprise agreement by mid-2025, alongside a $1 billion DHS blanket purchase agreement, expanding ICE and IRS data access, and NHS England's £330 million Federated Data Platform. Citizens cannot opt out of a tax filing, an immigration proceeding, or a hospital admission the way they could once decline a targeted

ad. The stakes of a government-run ontology getting a citizen's record wrong, or a hostile actor gaining access to it, are categorically higher than the stakes of a poorly targeted ad.

The United Kingdom is currently running a live test of this exact dilemma. Since NHS England awarded a Palantir-led consortium a seven-year, £330 million contract to build the Federated Data Platform in November 2023, Members of Parliament from multiple parties, two separate parliamentary committees, health worker unions, and human rights organizations have called on the government to use its February 2027 break clause to end the relationship. The lesson for U.S. federal and defense buyers is not that Palantir's platforms lack mission value. It is that structural permissioning and audit logging, the very controls Palantir recommends to its customers, are necessary but not sufficient once a government has let a single vendor's ontology become the connective tissue between agencies that were deliberately kept separate by law. Institutions cannot permission their way out of a surveillance-scale problem after the data has already been fused. The decision that matters happens before the contract is signed.

Introduction

The first paper in this Permuta Technologies series, *You Can't Prompt Your Way Out of Enterprise Complexity*, argued that AI-assisted software development cannot substitute for disciplined engineering in large, secure, government codebases. The second paper, *You Can't Consume Your Way Out of AI Cost Opacity*, argued that Palantir's own tokenized AI billing model turns what looks like a platform procurement into an open-ended consumption liability. This third paper shifts from cost and architecture to something more fundamental: the depth and sensitivity of the data itself, and the wisdom of government organizations moving quickly to consolidate that data inside a single vendor's platform before the tradeoffs are fully understood.

Palantir is not a generic software vendor, and it does not pretend to be one. It was founded to build software for the U.S. intelligence community, and its Gotham, Foundry, and AIP products are built around a proprietary "Ontology," a semantic model that maps an organization's data, logic, and actions into a single operational picture (Palantir product documentation). That architecture is precisely what makes Palantir valuable to mission-driven customers, and precisely what should make government buyers slow down before scaling it across every sensitive function of the state.

Palantir's own white paper makes an unusually candid case for caution, just aimed in only one direction. It correctly identifies that AI model providers have "a structural incentive to migrate as much intelligence from enterprises into their model weights… as possible," and it recommends zero data retention agreements, model-agnostic architecture, granular permissions, and continuous audit logging as defenses. Every one of those

recommendations is aimed at protecting an institution from a third-party model provider. None of them is aimed at the risk that the institution's own data, once fused inside Palantir's ontology across multiple agencies, becomes exactly the kind of concentrated, high-value target the paper warns other vendors might exploit. The same paper that warns institutions not to trust a frontier lab with their prompts is silent on what it means when a single company's ontology becomes the shared substrate for a nation's defense, immigration, tax, and health data at once.

Sovereignty for Me, Not for Thee

Palantir's sovereignty framework rests on fifteen steps organized across a model layer, a compute layer, and a control layer. The control layer is the most relevant to this paper. Palantir advises institutions to implement "granular permissions" by modeling their organization as a digital twin, because, in the company's own words, "when you model your world as a digital twin, permissions are comprehensive and are built to reflect the structure of your organization," including "role-, classification-, and purpose-based access controls that dynamically follow the lineage of your data."

That is an accurate description of what Palantir's Foundry Ontology is designed to do. It is also, read from the other side of the table, a description of how a government agency's data about a person, once inside that ontology, can be joined with another agency's data about the same person, provided permissions allow it. The U.S. Immigration and Customs Enforcement's Investigative Case Management platform, which Palantir has run since 2014, illustrates the point directly. According to the ACLU, ICM lets agents search across interconnected platforms and return data as detailed as an individual's schooling, family relationships, employment history, phone records, immigration history, biometric traits, criminal records, and home and work addresses, and an ICM funding document states plainly that "U.S. Citizens are still subject to criminal prosecution and thus are a part of ICM." The ACLU describes the resulting consolidation as "centralized dossiers of information on individuals within the US immigration system… that aren't limited to violent criminals, but include visa holders and even U.S. citizens," and argues this is precisely the kind of cross-agency data consolidation "Congress sought to limit in the Privacy Act of 1974."

Palantir's audit and logging pillar has the same asymmetry. The company recommends an append-only, signed log stream that records every user, prompt, model call, and data object touched, so an institution can "detect misappropriation" of its data by a model provider. That is a genuinely useful capability. But logging who touched a citizen's dossier after the dossier already exists inside a single vendor's ontology answers a narrower question than whether that dossier should have been assembled in the first place. Auditability is a control

on access. It is not a control on concentration. A government that has already let its tax, health, immigration, and defense data flow into one company's semantic model has already made the decision that logging, at best, can only monitor.

From Ad Clicks to Ontologies: The Google Playbook, Compounded

Shoshana Zuboff coined the term "surveillance capitalism" in 2014 and defined it as "the unilateral claiming of private human experience as free raw material for translation into behavioral data," which is then "computed and packaged as prediction products and sold into behavioral futures markets" (Harvard Gazette). She traces the practice directly to Google: "It was Google that first learned how to capture surplus behavioral data, more than what they needed for services, and used it to compute prediction products that they could sell to their business customers." Her formal definition goes further, describing the phenomenon as "a rogue mutation of capitalism marked by concentrations of wealth, knowledge, and power unprecedented in human history" (The Age of Surveillance Capitalism).

The popular shorthand for this idea, "if you're not paying for the product, you are the product," undersells what Zuboff actually argues. She corrects it directly: "we are the sources of surveillance capitalism's crucial surplus: the objects of a technologically advanced and increasingly inescapable raw-material-extraction operation" (Zuboff, quoted). Google built that extraction operation on browsing behavior, search queries, and, for roughly thirteen years, the content of Gmail messages, before ending content-based ad targeting of Gmail in 2017 while continuing to parse content for other product features (The Verge). Its 2007 acquisition of DoubleClick for $3.1 billion drew an immediate FTC complaint warning that the deal would give "one company access to more information about the Internet activities of consumers than any other company in the world" (EPIC); the FTC allowed the deal to close anyway, in a decision that in hindsight reads as an early missed opportunity for structural scrutiny. Twelve years later, French regulators imposed the first major GDPR fine, €50 million, on Google for its ad-personalization consent practices (The Guardian), and in 2024 a U.S. federal judge ruled that Google had illegally maintained a monopoly in search, a decision Google is still appealing in 2026 (Reuters; Reuters).

Every one of those episodes took roughly a decade to produce meaningful regulatory consequence, and the underlying product, targeted advertising based on consumer behavior, is by any measure a lower-stakes category of data than a citizen's immigration status, medical history, or tax return. The AI-era version of this pattern compounds all three of Zuboff's dimensions of concentration at once: the wealth flowing to AI-government platform vendors now measures in the tens of billions of dollars per company rather than the hundreds of millions Google earned in its early ad-tech years; the knowledge

concentrated in a single ontology now spans a person's entire relationship with the state rather than their browsing habits; and the power that follows from owning that ontology extends into decisions about deportation, benefits eligibility, and criminal investigation rather than which ad a person sees next. Where Google needed years of behavioral tracking to infer that a user was, say, expecting a child or planning a vacation, an AI platform with cross-agency access can simply read the underlying government record directly, with no inference required and no ability for the citizen to withhold consent.

The Scale and Sensitivity Problem

The speed at which Palantir's government footprint has expanded is itself part of the risk. The Department of Defense's Maven Smart System contract began in May 2024 at a five-year, $480 million ceiling, expanding the system's users "from hundreds to thousands worldwide" within months (Breaking Defense). By May 2025, a $795 million modification pushed the ceiling to nearly $1.3 billion through 2029 (DefenseScoop), and by July 2025 a separate ten-year, $10 billion Army enterprise agreement consolidated 75 existing Palantir contracts into one relationship (CNBC).

The pattern repeats on the civilian side. ICE's April 2025 award for an "ImmigrationOS" system began at $30 million and has since "ballooned to over $145 million," with Palantir's total ICE contract value reaching roughly $287 million in 2025, making it ICE's largest surveillance-technology contractor (ACLU). In February 2026, the Department of Homeland Security finalized a five-year blanket purchase agreement worth up to $1 billion that lets ICE, CBP, FEMA, USCIS, the Secret Service, TSA, and CISA buy Palantir's Gotham and Foundry platforms without separate competitive bids (WIRED). Reporting through 2025 also placed Palantir's Foundry inside at least four federal agencies including HHS, in active negotiation with the Social Security Administration and the IRS, and central to a project The New York Times described as raising the prospect of "detailed profiles of Americans using government data," built from records spanning bank accounts, student loans, medical claims, and disability status (The New York Times).

Ten members of Congress, led by Senator Ron Wyden and Representative Alexandria Ocasio-Cortez, wrote to Palantir's CEO in June 2025 warning that "the unprecedented possibility of a searchable, 'mega-database' of tax returns and other data that will potentially be shared with or accessed by other federal agencies is a surveillance nightmare," and alleging that the arrangement "blatantly violates the notice, transparency, and procedural requirements of the Privacy Act" (Nextgov). Palantir has denied building any master database, telling the Electronic Frontier Foundation in March 2026 that "each customer instance of our software is legally, technically, and operationally distinct" (Palantir letter to EFF). Whether or not any single technical database exists, the practical concern

raised by lawmakers and civil liberties groups is the same one this paper is making: a company whose own product literature describes an ontology purpose-built to fuse siloed data now holds simultaneous, deepening relationships with the agencies that hold a citizen's tax records, immigration file, health claims, and, through Maven, information relevant to lethal military decision-making. No advertising company in the surveillance capitalism era ever held that combination of access at that scale.

The UK Warning Shot: NHS and the Federated Data Platform

Government organizations do not have to speculate about how this plays out. The United Kingdom is living through it in real time, and the outcome so far is a caution, not an endorsement.

NHS England awarded a Palantir-led consortium a seven-year, £330 million contract in November 2023 to build the Federated Data Platform on Palantir's Foundry software, with a break clause reachable in early 2027 (The Register; NHS England contract explainer). Opposition began almost immediately: in April 2024, more than 100 health workers, patients, and advocacy groups picketed NHS England's offices demanding cancellation, and the BMJ ran a piece titled "NHS England must cancel its contract with Palantir" (BMJ). Roughly 50,000 patients subsequently wrote to their NHS trust boards urging non-adoption, and Greater Manchester's Integrated Care Board, covering 2.8 million people, deferred adopting the platform outright (The Register).

By 2026, opposition had moved from the streets into Parliament itself. In an April 2026 Westminster Hall debate, Liberal Democrat MP Martin Wrigley argued the platform "is awful to use, only benefits a quarter of its user organizations," and that the contract "delivers a subscription service that leaves no deliverables after the subscription… just a subscribed service; a permanent lock-in; a single point of failure" (The Register). Labour MP Chi Onwurah, who chairs the Science, Innovation and Technology Committee, said it was "appropriate for the government to explore all possibilities, including how to sever ties with the contract" (The Guardian). Labour MP Ian Byrne went further in a separate April 2026 debate, stating that "a company that aides Israeli war crimes, mines American citizens' health data to support ICE raids, and advocates NHS privatisation should be nowhere near our data. The Government must end its deal with Palantir in 2027" (The Guardian).

The parliamentary record has since hardened into formal committee findings. In June 2026, the cross-party Science, Innovation and Technology Committee published a report urging the government to trigger the 2027 break clause, describing reliance on Palantir as an "unacceptable vulnerability" and language calling the arrangement "incompatible with UK values," and recommending the NHS "either develop an in-house replacement or seek an alternative developed by UK-owned and UK-based providers" (Novara Media). The report

cited a revelation, first reported by the Financial Times in May 2026, that non-NHS staff, including Palantir personnel, had been granted "unlimited access" to identifiable patient data under an administrative role inside the platform (Reuters). In July 2026, a second committee, the Health and Social Care Select Committee, independently urged the government to scrap the contract, joined by up to 117 NHS data and technology workers publicly calling for it to be axed on patient-privacy grounds (The Guardian). Amnesty International, backing a briefing from the health-workers' organization Medact, argued in March 2026 that Palantir is "an unsuitable partner" for the NHS given its "track record of flagrantly disregarding international law and standards," including its work supplying artificial intelligence products to the Israeli military (Amnesty International). Campaign group Foxglove has run open letters to successive health ministers under the banner "a secretive US spy tech company like Palantir has no place in our NHS" (Foxglove). As of June 2026, Technology Minister Liz Kendall confirmed a formal government review of the contract ahead of the break clause, stating that "the current health secretary is examining every aspect of that contract to ensure we secure the best arrangement for Britain" (Reuters).

What makes the UK case instructive is not that Palantir necessarily fails at the technical job it was hired to do. It is that a G7 government with a national health service, robust parliamentary oversight, and a functioning break clause is still, more than two years after signing, struggling to determine whether it can safely unwind a single-vendor data-fusion relationship covering the health records of an entire national population. If disentanglement is this hard for the UK under those favorable conditions, it should give American federal and defense buyers real pause before letting a comparable relationship deepen across the IRS, DHS, HHS, and the Department of Defense simultaneously, in contracts that in several cases lack anything resembling the NHS deal's 2027 break clause.

Why Early Adoption Without Structural Limits Is Dangerous

Governments have been burned by opaque, trusted-by-default vendor systems before, and the pattern is consistent: the harm compounds for years before anyone with authority is willing to act on it.

The UK's own Post Office Horizon scandal is the starkest example. Between 1999 and 2015, more than 900 subpostmasters were wrongfully convicted of theft, fraud, and false accounting because of faulty accounting software built by Fujitsu, in what has been called the most widespread miscarriage of justice in British legal history and is linked to at least thirteen suicides (BBC). A public inquiry found that Fujitsu knew the software contained bugs as early as 1999 but did not disclose this to the courts, while Post Office managers "maintained the fiction" that the system was accurate (Yahoo/PA). It took roughly two decades, a landmark civil settlement, and dedicated legislation to begin correcting the

record, with redress payments exceeding £1 billion by mid-2025. The lesson is not that Fujitsu was uniquely malicious. It is that once a government agency has built its operations around a vendor's opaque system and trusted its outputs without independent means of verification, the cost of being wrong compounds silently for years before anyone notices.

Clearview AI offers a parallel on the data-fusion side specifically. The company built a facial recognition database of more than 30 billion images scraped from social media and other public sources without consent, and sold access primarily to law enforcement agencies (Business Insider). Dutch regulators fined the company €30.5 million in 2024 for processing biometric data without a legal basis, and issued orders effectively requiring it to cease EU operations (Library of Congress Global Legal Monitor); a UK tribunal separately found the country's own regulator lacked jurisdiction to act because Clearview served only non-UK law enforcement clients, illustrating how difficult mass biometric data-fusion products are to regulate even after the harm is recognized (BBC). And the Cambridge Analytica affair remains the clearest precedent for data collected under one set of assumptions being repurposed for political ends: personal data from up to 87 million Facebook profiles was harvested and used for psychographic voter targeting, resulting in a record $5 billion FTC penalty against Facebook, at the time the largest privacy penalty in U.S. history (FTC).

None of these three episodes involved a company with anywhere close to Palantir's current breadth of access across defense, immigration, tax, and health systems simultaneously. All three took years to surface, longer to remediate, and none was stopped before real people were harmed. That is the pattern government buyers should weigh honestly before treating an AI-era ontology platform's growth from a pilot to an enterprise-wide dependency as a straightforward productivity win.

Permuta's Contrasting Approach

Permuta's own architecture reflects a deliberate rejection of the single-ontology model. We made a very deliberate design choice over two decades ago to maintain a strict separation between our systems and customer data. DefenseReady is a Commercial Off-The-Shelf readiness platform built on Microsoft Dynamics 365 and the Microsoft Power Platform, deployed inside each customer's own IaaS IL5/6, SaaS IL5, or fully on-premises environment, with AI capabilities shipped disabled by default. Where Palantir's Foundry and AIP are architected around a single shared ontology that is designed to fuse data across missions and agencies, DefenseReady's AIReady framework is a metadata layer that prepares an individual customer's own workforce, training, medical, and readiness data for AI use inside that customer's own boundary, while explicitly preserving bring-your-own-model flexibility and avoiding the kind of always-on, vector-search-heavy cross-domain fusion that makes broad data consolidation the default rather than the exception.

That distinction matters directly to the risk described in this paper. Permuta's platform does not connect a defense customer's workforce and readiness data to another agency's immigration, tax, or health data, because it was never built as a cross-mission ontology in the first place; each deployment is bounded to its own customer environment and its own accreditation. Customers choose whether AI is enabled at all, which model endpoint is used, and whether that endpoint runs in Azure OpenAI, an on-premises runtime, or an approved government-specific environment such as Ask Sage or GenAI.mil. Permuta's own recent white papers and customer communications have been explicit that AI is treated as a force multiplier layered onto mature, accredited mission software, not as the substrate the mission is rebuilt around, and that pure code-generation and data-fusion approaches were deliberately not adopted for core product development because of the complexity, security, and rework risk involved at enterprise scale.

This is not a claim that Permuta's approach eliminates every data risk associated with AI. It is a claim that the choice of where AI sits in the architecture, as a modular, disableable, customer-controlled layer rather than as the connective ontology binding multiple missions together, is itself a sovereignty decision with real consequences, and one government buyers should evaluate with the same rigor they are now belatedly applying to Palantir's NHS contract.

What Government Buyers Should Do Differently

Palantir's own sovereignty paper is right that institutions need model-agnostic architecture, granular permissions, and continuous audit logging when dealing with third-party AI model providers. Government buyers should take that advice and extend it to cover the vendor supplying the advice. Four practical steps follow directly from the evidence in this paper.

First, evaluate data-fusion vendors on structural segmentation, not contractual permissioning alone. A row-level permission that can, in principle, be reconfigured to grant cross-agency access is a policy choice, not a technical wall. Agencies whose data was deliberately kept separate by statute, including under the Privacy Act of 1974, should require that separation be enforced architecturally, not merely by administrative discretion inside a single vendor's platform.

Second, insist on real, early, and enforceable exit mechanisms before signing multi-year enterprise agreements, not after adoption has already made the relationship politically and operationally difficult to unwind. The NHS Federated Data Platform's 2027 break clause is the exception that has made meaningful parliamentary review possible at all; several of Palantir's expanding U.S. federal contracts, including the DHS blanket purchase agreement and the Army's ten-year enterprise agreement, do not appear to carry comparably early or accessible exit points.

Third, treat "unlimited access" administrative roles, of the kind revealed inside the NHS platform, as a solved problem before go-live, not a finding to be corrected after a Freedom of Information request or a parliamentary inquiry surfaces it. Vendor personnel access to identifiable citizen records should be enumerated, time-bound, and independently auditable from day one.

Fourth, use the UK's ongoing review as a live pilot rather than a cautionary footnote. The Federated Data Platform review, due to conclude around the February 2027 break clause, will produce a real-world answer to a question American agencies are currently answering by default: whether a government can meaningfully evaluate, and if necessary unwind, a data-fusion relationship once it has scaled across a national health, immigration, tax, or defense system. Waiting for that answer before deepening comparable U.S. commitments costs little. Discovering the answer the hard way, after the data is already fused, costs a great deal more.

Conclusion

Palantir's own framework for institutional sovereignty gets the mechanics right and the target wrong. Zero data retention, model-agnostic routing, granular permissions, and audit logging are sound defenses against a frontier AI lab quietly extracting an enterprise's knowledge into its own model weights. None of them are a defense against the underlying fact that a single vendor's ontology now sits, simultaneously, inside the immigration file, the tax return, the hospital record, and the targeting system of the same citizens, at a scale and depth Google's advertising-driven surveillance capitalism never reached and never needed to. Google monetized what people did online. The platforms now scaling across allied governments hold what people are, on paper, to the state itself.

The United Kingdom's fight over the NHS Federated Data Platform is not a story about one difficult vendor relationship. It is a preview of what it looks like for a democratic government to try, belatedly, to claw back sovereignty over data it should never have let concentrate in the first place. Permission structures and audit trails cannot undo that concentration once it exists; they can only monitor it. Government organizations still deciding how quickly to scale AI-enabled data-fusion platforms across their most sensitive functions should treat the UK's experience, and the pace of Palantir's own U.S. federal expansion, as the clearest evidence available that sovereignty is won or lost at the moment of adoption, not recovered afterward through better logging.

About the Author

General Manager, Product Group, Permuta Technologies

Martin “Marty” Jennings, Colonel (retired), United States Air Force

For the past 3 years, Marty Jennings has been the General Manager for the Product Group at Permuta Technologies, where he leads a team of Microsoft-certified developers responsible for more than 400 applications supporting unique Federal and Military use cases worldwide. Prior to joining Permuta, he served as Chief Cloud Architect and Chief Engineer for Leidos’ 10-year, $10B GSM-O II contract with the Defense Information Systems Agency (DISA). A 30-year U.S. Air Force veteran, Marty retired in the rank of Colonel in September 2021. His military and civilian career spans software development, cybersecurity, and large-scale enterprise system program management in support of multiple combatant commands and agencies.

LinkedIn: https://www.linkedin.com/in/jentek