A SAYARI SUPPLY CHAIN INTELLIGENCE BRIEF . MAY 27, 2026

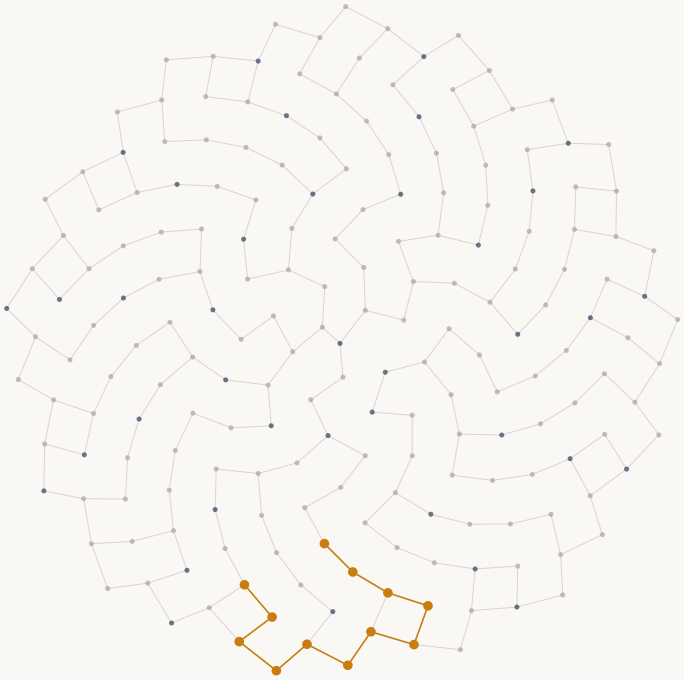

Securing the $10 Trillion AI Supply Chain

Self-reported data isn't evidence. Primary-source records are.

SUPPLY CHAIN INTELLIGENCE

Securing the $10 Trillion AI Supply Chain: The Case for Primary-Data Compliance

Sayari · May 27, 2026

Compliance programs built on supplier questionnaires and probabilistic mapping were designed for a different regulatory era. The enforcement tools now targeting AI supply chains demand something more deterministic: primary-source evidence, traceable to its origin.

The Regulatory Moment

AI infrastructure investment is projected to exceed $10 trillion over the next decade, roughly the combined size of the German and French economies. The hardware behind that build-out — semiconductors, rare earth components, and finished devices — moves through some of the highest-risk supply chain geographies in the world.

Enforcement has followed the investment. The UFLPA has reviewed more than 41,000 shipments, denied more than 23,000, and placed over $4 billion in goods at risk. The EU's CSDDD is expanding forced labor enforcement across the Atlantic. New tariff measures targeting copper, aluminum, and semiconductors range from $195 to $287 billion. The volume of capital moving through these supply chains and the pace of regulatory enforcement are both outrunning the compliance infrastructure most organizations have in place.

The Problem With Existing Programs

Most compliance programs still rely on supplier questionnaires or probabilistic mapping. These tools were built for a different set of expectations and carry three structural gaps that make them difficult to defend under current enforcement standards.

- **No universal entity registrar.** Verifying whether a tier-three supplier is connected to a state-owned enterprise or a sanctioned actor requires reconciling fragmented, inconsistent

data across jurisdictions. The gap is not a matter of effort — it is architectural.

- **HBOMs and risk data are rarely linked.** Hardware bills of materials, where they exist at all, are not connected to origin data or compliance status in any standardized format. A component can appear clean in screening and be sourced from a high-risk smelter that no questionnaire surfaces.
- **Conformity claims are largely unverifiable.** A supplier may assert compliance. A certificate may exist. Neither confirms the current state of the facility or the product. Under the UFLPA's rebuttable presumption of guilt, inference does not hold up at the border. The standard is: here is the verifiable origin of every material in this product, traced to its source.

The Tech Against Trafficking AI Provider Traceability Standard

Published in March 2026 by the Tech Against Trafficking coalition, the AI Provider Traceability Standard defines five mandatory record types that AI providers and their supply chain partners must maintain and be able to transfer. Together, they answer the who, what, when, where, and how of every node in the supply chain.

- **Digital Corporate Records (DCR) — the who.** The legal entity behind each facility: ownership structure, operational control, and cross-jurisdictional identifiers. The foundation for sanctions and export control screening.
- **Digital Facility Records (DFR) — the where.** Physical location, operational certifications, and linkages to state-owned enterprises or restricted entity lists such as the BIS Entity List.
- **Hardware Bills of Materials / Digital Product Passports (HBOM/DPP) — the what.** A component-level inventory linked to verifiable origin and compliance status data. The record most organizations are furthest from being able to produce today.
- **Digital Traceability Events (DTE) — the when.** A time-stamped log of key lifecycle events: procurement, configuration changes, and decommissioning. The chain of custody.
- **Digital Conformity Credentials (DCC) — the how.** A verifiable attestation that a facility meets standard requirements, functioning as a machine-readable compliance certificate.

The standard does not ask organizations to create new data. It asks them to connect data that already exists across ERP, PLM, and trade intelligence systems in a format that can be transferred, validated, and trusted across supply chain tiers. The standard is built on existing

conventions, including EPCIS, to minimize the burden on suppliers providing data in a consistent format.

The Sayari and Source Intelligence Integration

A complete picture of supply chain provenance requires answers to three questions: who you are dealing with, what you are dealing in, and where it comes from. Sayari and Source Intelligence address each leg of that triangle. Together, they close the gap that neither addresses independently.

Sayari covers the *who*: corporate ownership mapping, entity resolution across 250+ jurisdictions, sanctions and export control screening, and trade intelligence that surfaces what is moving through a supply chain in near real time. Source Intelligence covers the *what*: full material disclosures, conflict minerals data, HBOM collection workflows, and supplier compliance surveys. The *where* is where the two datasets converge.

The practical effect is precision. Without full material disclosure, a compliance team screening for risk exposure faces millions of potential entity relationships, many of them false signals. With material disclosure as the backbone, the entity graph narrows to the suppliers, smelters, and facilities actually tied to materials in the product. The question shifts from "could we be exposed?" to "here is the specific node in our supply chain that represents material risk."

The integration maps directly to the TAT standard's eight specific requirements: Sayari handles entity identification (DCR), facility-level risk screening (DFR), and trade flow intelligence. Source Intelligence handles HBOM data collection, full material disclosures, and conformity workflows (DCC). Together, they produce the linked, primary-source record set the standard requires.

Getting Started

Full traceability is not a single step. The organizations furthest along built iteratively, beginning with their highest-risk product lines and the supply chain branches carrying the most enforcement exposure. Three practical entry points:

- **Read the standard.** The TAT AI Provider Traceability Standard is publicly available at the Tech Against Trafficking website. It is deliberately practical: it describes data that already exists inside most organizations.

- **Identify highest-risk exposure.** For AI infrastructure procurement, that typically means rare earth inputs, smelter-level traceability for UFLPA-listed materials, and entity-level screening for state-owned enterprise connections in tier-two and tier-three suppliers.
- **Map existing data sources.** ERP and PLM systems often hold more usable supply chain data than compliance teams have access to. Connecting those internal records to external entity and trade intelligence is frequently the most direct path toward meeting the standard's record-type requirements.

Please visit sayari.com to learn more about how Sayari and Source Intelligence can map their combined data infrastructure to your supply chain, product lines, and enforcement exposure.