

How to detect contract leakage across ERPs with graphs

Procurement leaders detect contract leakage across ERPs by replacing text searches with graph-based attribute matching. When procurement teams connect material master records and supplier agreements within an industrial ontology, off-contract spend and missed tier rebates become visible. The quality of the data foundation decides what agents and analytics can do.

In multi-ERP operating environments, enterprise resource planning tables isolate item records and purchase histories within local databases. Standard relational queries rely on shared primary keys, common part numbers, or identical item descriptions. Because distributed plants log identical items under distinct naming conventions, relational queries fail to reconcile purchase records with negotiated commercial terms.

Organizations overcome these visibility gaps by moving away from static database extraction. To identify unmanaged variance across decentralized facilities, procurement teams stop querying records and start navigating relationships.

Why relational ERP queries miss contract leakage

Relational ERP architectures store procurement data in normalized tables separated by plant codes, purchasing organizations, and legal entities. Finding pricing discrepancies across these tables requires SQL joins across purchase order lines, invoice records, material masters, and outline agreements.

These table joins require identical foreign keys or standardized item descriptions. In global operating environments, distributed business units create duplicate material records with localized descriptions, missing manufacturer codes, and differing units of measure.

Relational database joins fail in multi-ERP landscapes because distributed tables lack shared semantic meaning. When local systems record the same physical item under unique material numbers, relational queries treat these purchases as distinct, unrelated goods.

Traditional spend analytics pipelines extract transactional tables into central data lakes or data warehouses to run text-matching scripts.

Machine learning models trained on pure text patterns match words rather than engineering specifications. These models produce false positives on distinct technical components while missing true duplicates described in different languages.

When queries cannot bridge these description discrepancies, off-contract purchases pass undetected. Procurement category teams miss volume aggregation targets, pay unnegotiated spot prices, and lose negotiated tier discounts.

How industrial ontologies resolve cross-ERP entity discrepancies

Industrial ontologies define categories, technical attributes, and physical relationships across materials, services, and suppliers. Rather than treating item descriptions as arbitrary strings of text, an ontology interprets procurement records through explicit technical rules.

The ontology organizes knowledge around 30,000 industrial categories in any language. It parses free-text descriptions into structured technical properties, such as dimensions, material grades, pressure ratings, and international standard references.

This semantic entity resolution solves the structural limitation of multi-ERP environments. When local plants purchase the same item under local conventions, the ontology resolves the records to a common semantic definition.

Consider an identical hexagon head screw purchased across three separate facilities:

- **Lyon plant:** *Vis tête hexagonale ISO 4017 M16x60 classe 8.8 zingage électrolytique*
- **Ludwigshafen plant:** *Sechskantschraube DIN 933 M16x60 8.8 galvanisch verzinkt*
- **Teesside plant:** *HEX SET SCREW M16 X 60 GR 8.8 ZINC PLATED*

These three records share no part numbers, share very few words, and reference two distinct technical standards. Relational queries and text-matching algorithms treat them as three unrelated items.

The ontology confirms ISO 4017 and DIN 933 describe identical fully threaded fasteners at M16 (both specify 24 mm across flats, though widths differ at M10, M12, M14, and M22). Matching attributes include: fully threaded hex head; M16 diameter; 60 mm length; property class 8.8; and electroplated zinc finish.

By isolating verified physical attributes from localized phrasing, procurement systems establish entity equivalence across distributed enterprise systems. For complex operations, mapping multilingual spend data with AI techniques provides the prerequisite normalization that makes global cross-plant comparisons possible.

Building semantic digital twins of procurement data

Once an ontology defines domain attributes, the enterprise generates a Semantic Digital Twin for every material, service, supplier, asset, and contract. A Semantic Digital Twin serves as an AI-enriched digital representation defined by its verified technical attributes and relational connections.

Technical Attribute Management (TAM) builds and governs the twins, extracting attributes, classifying items, and detecting duplicates alongside existing platforms like SAP S/4HANA and SAP MDG. TAM extracts properties from legacy records, assigns international standard taxonomies, and links related entities into a coherent network.

Instead of maintaining a flat material master record, the twin contains detailed properties:

- **Mechanical attributes:** dimensions, thread pitch, load capacity, and material composition.
- **Operational attributes:** operating temperatures, pressure tolerances, and environmental limits.
- **Commercial relationships:** qualified manufacturer references, parent supplier corporate structures, and valid pricing agreements.

When applied to commercial contracts, twins represent negotiated terms as structured data nodes. Price brackets, minimum order quantities, and annual volume rebate

thresholds connect directly to the underlying material definitions.

Connecting contracts directly to material twins bridges the operational divide between sourcing negotiations and plant-level requisitions. Procurement leaders inspect active spend against negotiated contracts across every connected ERP instance.

Connecting twins within the Enterprise Data Space

The enterprise connects these individual entities into the Enterprise Data Space (EDS), the enterprise's semantic context layer. The EDS is the semantic environment that emerges once an organization links its Semantic Digital Twins across enterprise domains.

Within the EDS, data points do not live in isolated rows or columns. Every material twin links to its approved suppliers, historical purchase orders, physical plant installations, and active master service agreements.

Procurement teams leverage this connected architecture through domain-specific AI agents. AI agents are software components that read the EDS through open APIs and the Model Context Protocol (MCP), then act on what they find.

These agents inspect complex procurement networks to unmask contract leakage that relational queries overlook. By traversing the graph, an agent traces purchases made under spot purchase orders back to active corporate outline agreements for the same technical items.

The functional mechanics of querying relational tables differ from navigating the EDS on every count:

- **Data structure.** Relational: tabular relational databases and data lakes. EDS: a connected graph of Semantic Digital Twins.
- **Item matching logic.** Relational: text strings, material numbers, and vendor codes. EDS: harmonized technical attributes and ontology rules.
- **Multi-language capability.** Relational: relies on literal translation or exact keyword matches. EDS: resolves technical equivalence in any language.
- **Multi-ERP consolidation.** Relational: manual table joins across custom extraction schemas. EDS: semantic digital twin resolution across systems.
- **Leakage identification.** Relational: identifies variance only when part numbers match. EDS: detects off-contract spend across divergent descriptions.
- **Maintenance workflow.** Relational: periodic manual database cleansing batches. EDS: continuous governance integrated into daily creation.

Organizations seeking comprehensive cross-plant governance find that enterprise spend analytics for fragmented data requires this shift from static string matching to dynamic semantic modeling.

Four steps to detect contract leakage with relationship graphs

Procurement organizations implement graph-based contract leakage detection through an established engineering workflow. Each phase expands visibility while anchoring transactional activity to governed data models.

1. **Ingest and extract multi-ERP records.** Extract material masters, service sheets, purchase order lines, and active agreement clauses across all enterprise ERP systems. These sources include central SAP instances, legacy plant-level databases, and specialized contract repositories. The ingestion process preserves source references while staging raw text for attribute extraction. No operational data is destroyed or overwritten in source systems during this stage.

2. **Standardize properties and generate entity twins.** Process the ingested records through domain ontologies to extract discrete technical attributes. Creactives TAM extracts dimensions, material grades, performance parameters, and industrial standards from unstructured text. The system then instantiates Semantic Digital Twins for every physical material, vendor entity, and contract clause. Local catalog numbers resolve into unified twin profiles that carry cross-system references.
3. **Establish relational links across commercial networks.** Connect the generated twins within the EDS. In this phase, relationships connect materials to their parent assemblies, authorized alternative components, approved supplier sites, and negotiated price agreements. Contract terms, including tiered rebates and site-specific discounts, attach as active properties to the relevant material categories. The resulting semantic graph models how goods and commercial agreements interact across global operations.
4. **Deploy agents to identify spend variance and uncollected rebates.** Deploy AI agents to traverse the EDS and evaluate enterprise-wide transactions against active contract models. Agents trace purchases that bypass corporate framework agreements, quantify pricing discrepancies, and surface missed volume rebate thresholds. The agents deliver structured audit trails showing where local buyers selected non-contract vendors for items already covered under corporate pricing.

Recovering value across global procurement networks

Undetected contract leakage drains enterprise profitability when negotiated purchasing agreements remain disconnected from day-to-day requisitions. Traditional data management strategies depend on uniform system compliance, an assumption that collapses across global supply operations.

Relational databases and manual extract routines cannot bridge the linguistic and structural gaps inherent in enterprise ERP environments. Industrial ontologies and Semantic Digital Twins provide the connective framework necessary to unify fragmented procurement landscapes.

Enterprises such as TotalEnergies, Enel, and Subsea7 work with Creactives to build continuous semantic integrity across their global supply networks.

Establishing deep semantic relationships between materials and agreements eliminates the blind spots that allow contract non-compliance to persist. As procurement leaders evaluate AI-driven operational models, the integrity of the underlying data network determines overall analytical success.

Creactives customers achieve 15–20% savings on addressable spend. By transitioning from disconnected data queries to navigated semantic relationships, global procurement teams secure compliance across every operating site.

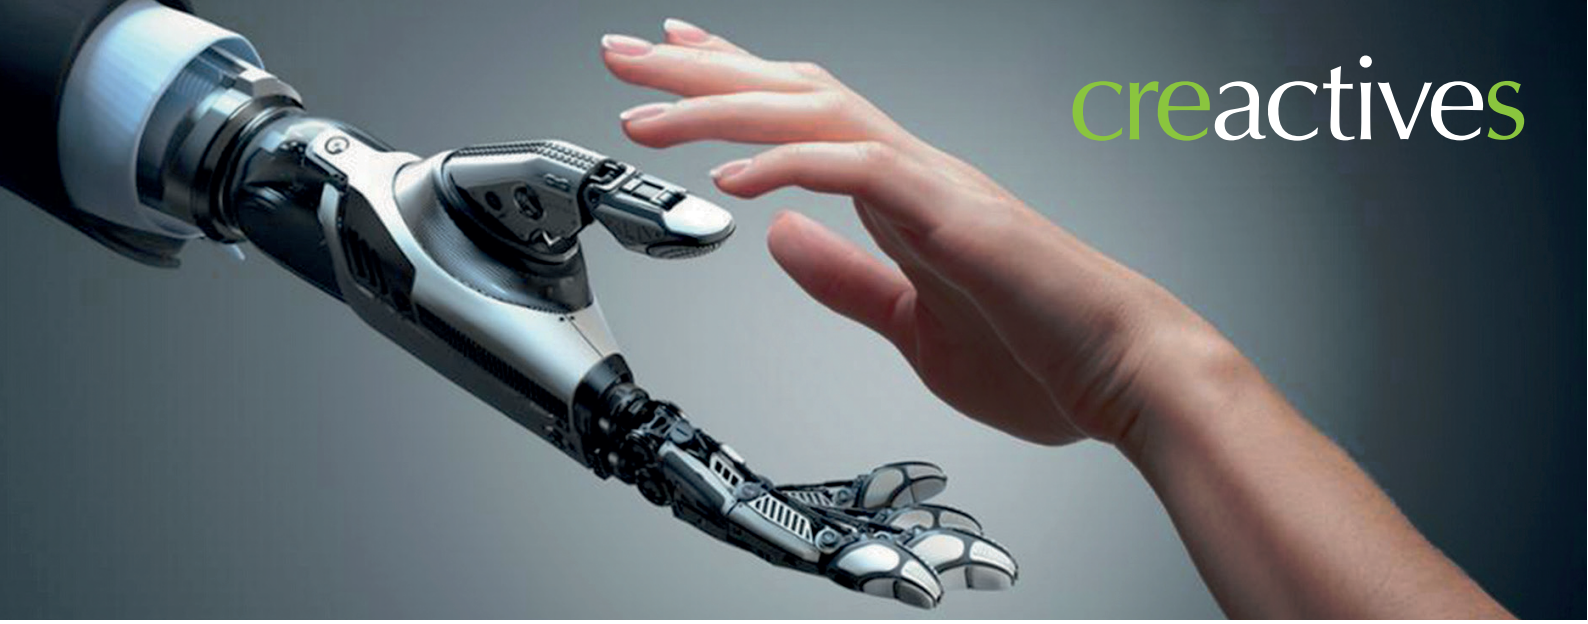
Frequently asked questions

Why do relational ERP queries fail to find contract leakage across multiple systems?
Relational ERP queries depend on identical part numbers, primary keys, or exact text descriptions across normalized database tables. When separate operating facilities log identical items under localized descriptions and distinct naming conventions, relational joins fail to recognize equivalence across systems.

How do industrial ontologies resolve discrepancies between duplicate records?
Industrial ontologies parse unstructured descriptions into discrete technical properties, such as dimensions, material grades, and industrial standards. By matching engineering attributes rather than text strings, the ontology confirms physical equivalence across records written in different languages or formats.

What is the role of semantic digital twins in tracking contract compliance? A semantic digital twin links the technical attributes of an item directly to negotiated commercial terms, including tiered pricing, minimum order quantities, and rebate thresholds. This link enables procurement teams to evaluate daily purchasing activity directly against valid corporate agreements.

How do AI agents detect contract leakage within a connected data space? AI agents traverse the semantic context layer through open application interfaces to audit multi-ERP transactions. They trace purchase orders made under spot pricing back to active corporate outline agreements for matching technical items, highlighting unmanaged spend variance.



creactives

Creactives builds a trusted Data Foundation that creates a Digital Twin of your Master and Spend Data. Our AI-powered solutions eliminate errors, optimize legacy systems, and unlock a wide range of use cases across Procurement & Supply Chain.

SPEND ANALYTICS

MATERIALS & SERVICES
MASTER DATA GOVERNANCE

PROCUREMENT GUIDANCE
AUTOMATION

CREACTIVES: SOLUTIONS SELECTED OVER YEARS BY OPINION LEADERS





"Creactives is a best-of-breed tool that offers standard spend analytics features, such as data cleansing and analysis. In addition, it specializes in materials MDM and is built for automated parts and material classification.".

Abigail Ommen, Research Analyst & Analyst Production Manager

Creactives Group SPA
Via Fermi, 4 – 37135 Verona – Italy

creactives

creactives.com
linkedin.com/company/creactives-spa