

## From PDFs to Procurement: Realizing the Promise of Model- Based Product Definition

Andrew Bank  
Exiger, LLC

## Introduction

In recent years, organizations across the aerospace, defense, and advanced manufacturing sectors have made significant investments in digital transformation. These initiatives are driven by clear imperatives: reduce engineering cycle times, improve data traceability, mitigate compliance risk, and shorten time to market. At the center of these efforts is a concept often cited but rarely achieved in practice: the digital thread – a continuous, integrated flow of data that connects every stage of a product’s lifecycle, from initial requirements to manufacturing, procurement, and sustainment.

Despite its prominence in strategy decks and conference panels, the digital thread remains stubbornly fragmented for many organizations. Data silos persist. Legacy documents resist integration. And crucial product knowledge remains locked in analog formats that cannot be interpreted by machines, easily consumed by humans, or shared between teams.

One of the primary barriers is the lack of a complete, model-based product definition – one that captures not just geometry and tolerances, but **also the complete operational, compliance, and procurement context of a part or assembly**: materials, manufacturing processes, process equipment and operators, procurement lead times, test protocols, inspection methods, quality requirements, regulatory rules, parts availability, supplier risk, and sustainment considerations. Without this full definition, model-based engineering remains siloed and aspirational, rather than integrated and operational.

This paper explores the root causes of this gap, examines why traditional approaches to digitalization fall short, and outlines a practical solution for creating a truly complete, model-based product definition – one that includes not only CAD geometry and PMI but also the standards, specifications, and operational requirements that actually govern how parts are made, tested, procured, and sustained. It is a call to action for engineering, procurement, and compliance leaders – as well as the publishers of external industry information – to rethink their approach to legacy content, and to begin treating it as the strategic asset it truly is.

The [Defense Standardization Program Office’s 2026 Digital Standards Strategy \(DSS\)](#) makes this issue more immediate. It defines digital standards as standards published in machine-readable and machine-interpretable formats for use in digital tools and processes, including models. That is a clear signal that defense programs are moving away from standards as static reference documents and toward standards as usable digital inputs.

This paper focuses on the missing step between that direction and real execution: turning legacy standards, specs, and drawings into structured product definition data that engineering, procurement, quality, supply chain, and model-based tools can actually use.

**Have you read the “prequel” to this article? Read it now:**  
**[“The Hidden Costs of Using Analog Data in Modern Engineering Workflow”](#)**



## What Complete Model-Based Product Definition Really Requires

In many organizations, “Model-Based Definition” (MBD) is implicitly understood as the 3D CAD model (e.g., a CATIA or CREO solid) plus associated Product and Manufacturing Information (PMI). That is a critical starting point, but it is not sufficient as a complete Model-Based Product Definition (MBPD).

A complete MBPD must define exactly *all* requirements for the product, assembly, or part in a model-based way. That includes:

- ◆ Geometry and tolerances
- ◆ Materials and material conditions
- ◆ Heat treatment and surface treatments
- ◆ Manufacturing and process requirements
- ◆ Test protocols and inspection criteria
- ◆ Quality and acceptance criteria
- ◆ Environmental, safety, and regulatory requirements
- ◆ Supply chain and operational constraints (lead times, sourcing limits, export controls, part or material scarcity, etc.)

Today, many of these requirements are captured in purchase orders, process specifications, material standards, and industry standards (SAE, ASTM, MIL, etc.). These artifacts are usually human-readable only (PDF, Word, scanned images). ***Even when PMI references them, they remain external static documents, not model-based data. And even when CAD models include this information, they are usually analog notes within the model.*** This creates a hole in what is often labeled “model-based” – the CAD model is parametric and structured (i.e. machine-readable), while the specifications that drive actual manufacturing and procurement behavior remain opaque to machines.

Closing this gap requires all these requirement sources to become first-class, model-based citizens. A true MBPD must go beyond the 3D model and PMI to incorporate every relevant requirement – including from external industry standards – in a structured, machine-consumable way. Only then can downstream processes – manufacturing planning, procurement, compliance, and sustainment – rely on the product definition as a single, authoritative source of truth.



## Where the Status Quo Falls Short

Here are several tangible areas where the status quo fails to provide sufficient data to power manufacturing, planning, procurement, compliance, and sustainment.

### Siloed, Analog Data

Engineering and manufacturing teams work in high-stakes environments where precision, traceability, and compliance are mission-critical. Every decision – from selecting a material, to approving a supplier, to updating a specification – depends on accurate and complete information. Yet, for most organizations, this information remains fragmented across a patchwork of formats, systems, and departments. The result is a costly and persistent data handicap.

Much of the critical product knowledge still lives in static PDFs: technical data packages, drawings, customer requirement documents, procurement specifications, and process standards that require manual interpretation and rekeying. Engineers routinely spend hours copying and pasting content between documents and systems, risking transcription errors, version mismatches, missed updates, and interpretation mistakes along the way. When a requirement changes, the implications must be tracked by hand – often without visibility into downstream consequences. When data is transcribed from an analog source to create a derivative – for example, a set of work instructions or a vendor RFP – it proliferates the problem **by creating yet more analog data**.

Critically, this fragmentation applies not only to internal documents but also to the external industry standards that govern many of those requirements. Even when a drawing references “per ASTM B117” or “per AMS 2750,” that connection is a textual pointer to another human-only document, not a data-level link to structured, versioned requirements.

### Outdated Specifications

When outdated specifications make their way into new designs, the impact reverberates across engineering, procurement, and manufacturing. Obsolete requirements often remain hidden until they reach the teams responsible for sourcing or production – who then discover that a material, coating, or process has been discontinued, replaced, or newly regulated. At that point, corrective action must be initiated to resolve the issue.

In defense programs, this often takes the form of a DD Form 339 request – a formal change request sent from logistics or procurement organizations (such as DLA) back to the engineering support activity (ESA) for clarification or authorization. In commercial settings, the same dynamic appears as engineering change notices (ECNs) or supplier deviation requests (SDRs) or worst of all, non-conformance reports (NCRs, which usually surface late in the product development lifecycle, often after first article release). Regardless of the form, the pattern is the same: an outdated spec triggers a costly cycle of rework, administrative back-and-forth, and production delay. The product manufacturing information that was defined in design travels clumsily through the digital thread as static PDF attachments rather than as model-based data with actionable trigger points, machine-

readable flags, and other programmable stage gates to monitor and validate manufacturability, obtainability, compliance, quality, and performance.

Every time a design includes a superseded specification, the organization pays twice – once when engineers unknowingly reuse it, and again when downstream teams must halt production to resolve it. The cost is not just in hours or paperwork; it's in lost time-to-market, supplier confidence, and overall program momentum.

Outdated or superseded industry standards amplify this problem. A drawing may continue to reference an obsolete ASTM or SAE document long after that standard has been revised or cancelled. Without a model-based representation of standards – with explicit versioning, lineage, and replacement mappings – engineers have no systematic way to know that a requirement they are reusing is already out of date. Today, there is some progress toward a model-based representation of standards, but commercial entities have made their needs clear: ***industry publishers must move faster to provide their content as model-based “digital twins” that integrate with modern model-based data platforms and workflow systems.***

## **Lack of Visibility into Regulated Materials and Critical Minerals**

Static documentation also obscures the regulatory and geopolitical realities of material sourcing. A specification may call for a legacy material – such as those containing PFAS, cadmium, or other restricted substances – without any indication of its regulated status or supply chain risk. Similarly, designs may rely on critical minerals that are scarce or restricted from regions like China or Russia. These oversights force procurement teams to scramble for alternatives late in the process, driving up cost and risk. What's missing is not just data, but context – the intelligence needed to make responsible material choices early in design.

## **Long-Lead-Time (LLT) Components**

A similar situation exists for parts and equipment with long administrative and production lead times. When a part is specified on a drawing, the notes rarely call out production lead time, so castings and forgings in particular – often averaging ~200 total days of combined administrative and manufacturing lead time – hide their true schedule impact. Those realities usually surface months later in procurement, when timelines are already committed and schedules start to slip. A single critical component can stall entire programs, especially when it is a prerequisite for downstream assembly steps. Richer product definition related to manufacturing processes can surface these LLT risks earlier and even flag candidates for alternate approaches such as additive manufacturing, where castings and forgings are often leading opportunities for redesign.

## **Supplier (In)Visibility**

Finally, design teams typically have little awareness of supplier health, capacity, or risk exposure. A single specification can quietly lock in dependency on a vulnerable source – one facing financial strain, foreign ownership, compliance violations, or even environmental disruptions. None of this intelligence is captured in static documents, yet all of it can determine whether a part is delivered on time, to spec, and within budget.

The labor involved in reconciling these disparate sources is not just inefficient – it introduces real operational risk. Missed requirements, overlooked regulated substances, or reliance on outdated specifications can lead to delays, non-compliance, and costly rework. More importantly, it undermines the goal of creating a unified, authoritative product definition that can be leveraged throughout the product lifecycle.

## Data Inconsistency and Format Incompatibility

Another persistent challenge is the inconsistency of data formats across engineering ecosystems. Even within a single organization, product information is fragmented across CAD systems, PLM platforms, Excel spreadsheets, SharePoint folders, Word documents, and scanned PDFs. Each department may have its own tools and conventions, making it extremely difficult to consolidate information into a consistent, model-based structure.

Efforts to establish a digital thread are often thwarted by the need to harmonize this disparate content. While CAD files may support geometric product specifications and some parametric data, they do not contain the rich contextual and compliance-related content found in accompanying documentation. The result is a partial model that lacks the fidelity needed to support downstream automation in manufacturing, procurement, and sustainment.

Furthermore, most of this content is not machine-readable. It cannot be programmatically queried, interpreted, or linked to other systems without human intervention. As a result, any attempt to automate downstream tasks – such as supplier risk analysis, regulatory compliance checks, or digital twin simulation – is severely limited by the quality and accessibility of the underlying data.



## Another Hidden Gap: Industry Standards are Not Model-Based Requirements

Many of the most consequential requirements for a part or assembly do not originate in internal documents at all. They come from external standards bodies: SAE, ASTM, ASME, MIL, IEC, UL, ISO, IPC, and others. These standards define how materials are characterized, how heat treatment is performed, how coatings are applied, how testing is executed, and what constitutes acceptable quality.

In today's workflows, those standards:

- Are referenced textually (e.g., “per ASTM B117”) in drawings, TDPs, or work instructions
- Remain as analog, human-readable PDFs or books
- Are rarely version-controlled or linked at a data level to the parts and assemblies that depend on them

This means that even when the CAD/PMI model is parametric and machine-readable, a large fraction of its governing requirements remains outside of the model, embedded in documents meant for human consumption only. Engineers often assume that “the model is the definition,” but

in practice the model is only a partial definition, pointing outward to a web of unstructured specifications.

The DSS says this shift starts with machine-readable formats such as XML, but it does not stop there. It also calls for machine-interpretable formats, including models, so standards can be used directly in digital tools and processes.

That matters because a note like “per ASTM B117” is not model-based product definition. It is still a pointer to a document that someone has to find, read, and interpret. In a true model-based environment, the governing requirement has to exist as structured, versioned data tied to the part, process, test, and downstream systems that depend on it.

The DSS makes this even more concrete by naming Systems Modeling Language (SysML) as one example of machine-interpretable content and by pointing to the first fully digital DID in ASSIST as a SysML version. If standards remain analog while the rest of the engineering stack becomes digital, MBPD will remain incomplete.

To achieve a true MBPD, these standards must be brought into the model itself – parsed, structured, linked, and versioned so that every reference to “ASTM X” or “AMS Y” corresponds to explicit, machine-consumable requirement objects, not just free text.



## The Limitations of Existing AI and “Digitalization” Tools

To address these challenges, a range of tools including untrained commercial AI models have emerged claiming to “digitize” engineering documentation. However, most fall short because they lack semantic understanding, robust engineering ontology, and authoritative knowledge of external standards, and therefore cannot produce a complete, accurate, ontology-based, or machine-interpretable Model-Based Product Definition (MBPD).

A common example is requirements extraction. Many tools rely on NLP to flag words like “shall,” “should,” or “may,” which may indicate that a requirement exists but does not reliably determine what it refers to – whether it is a tolerance, surface finish, chemical specification, test method, or process constraint – nor does it consistently capture the associated values, ranges, dependencies, or engineering intent required for MBPD.

Without a canonical semantic model to enforce consistency across documents, drawings, and systems, the same engineering concept may be represented in multiple incompatible ways, preventing parts from being grouped, compared, or analyzed reliably. And while general-purpose AI can be powerful, it also introduces the risk of hallucinated interpretations and inferred attributes in complex technical contexts. Finally, most tools do not interpret or structure external standards as reusable, parametric requirement sets – identifying a reference to “ASTM B117” is not the same as transforming that standard into model-based data that can be instantiated, queried, versioned, and validated across programs.

As a result, organizations relying on general AI or “digitalization” tools alone tend to end up with a partial and fragile product definition requiring extensive manual validation and constant maintenance. Achieving a complete MBPD requires more than extraction or pattern recognition: it requires semantic consistency, authoritative reference data, and an integrated approach that combines multiple complementary technologies.



## A More Capable Approach: Full-Stack Document Intelligence and Structured Requirements Graph

A more effective strategy begins with a different premise: the goal is not simply to extract data, but to understand it in context and convert it into structured, interoperable knowledge that can power the digital thread. Comprehensive MBPD requires an ensemble of technologies working together. GenAI provides pattern recognition and language understanding. Ontology-based semantic models supply meaning and structure. Technical databases deliver authority. Together, they transform unstructured engineering content – drawings, specifications, standards, and test methods – into a coherent and complete model-based definition.

This requires an engineering AI system that combines multiple layers of intelligence, including:

- Semantic engineering knowledge models that represent how parts, materials, processes, coatings, and corresponding standards relate to each other
- Industry-aligned ontologies that provide the structure for semantic models and enforce semantic consistency
- Authoritative, machine-readable metadata (and/or full-text) for industry standards (SAE, ASTM, ASME, MIL, etc.)
- Curated technical databases containing references to regulated substances, specification hierarchies, and historical revision data
- Manufacturability and production logistics that factor in elements like available machinery and capabilities, machine and operator scheduling, and throughputs and capacities
- Supply chain mapping and risk mitigation that analyzes procurement risks like regulatory compliance, materials and parts availability (or lack thereof), and supplier reliability

This is where SWISS fits. The DSS calls for a move from human-readable documents to machine-readable and then machine-interpretable content. SWISS provides that bridge by extracting requirements from drawings, specifications, and standards and turning them into structured, traceable data.

XSB has already demonstrated the basics of this approach in prototype work. Requirements captured in SWISS were mapped into SysML in Cameo, with traceability back to the source and links to external specifications and standards. That shows the playbook is real: legacy documents can become structured requirement data first, and model-ready outputs second.

SWISS is not a publishing system or a replacement for modeling tools. It does the heavy lift most organizations still avoid: turning legacy technical content into reusable digital requirement data.

Crucially, this approach must not stop at internal documents. It must also ingest and interpret external standards, normalizing their contents into a structured representation that can be consistently reused across programs. ***The system must be able to treat an industry standard as a data source, not just a citation.***

Conceptually, the solution follows a repeatable pipeline:

1. Ingest internal documents (drawings, work instructions, purchase specifications) and external standards (SAE, ASTM, ASME, MIL, etc.).
2. Parse and normalize the content, identifying key text elements, references, tables, and figures.
3. Map the extracted content to an ontology of parts, materials, processes, and other requirement types.
4. Track revisions and hierarchies, such as “this AMS spec replaces that QQ spec” or “this issue of a standard introduces new restrictions.”
5. Generate structured, machine-readable requirement objects that can be bound to specific parts, features, or assemblies.

When these components are combined, the system can make inferences that are beyond the reach of traditional OCR, NLP, or general-purpose AI tools. For example:

- A reference to QQ-A-601 is correctly understood as indicating the use of aluminum sand casting, which implies both a material and a manufacturing process.
- A reference to Aluminum Alloy 6061 in one drawing, composition 6061 or Aluminum Alloy UNS A96061 in another should consistently identify the same material concept.
- A mention of QQ-P-416 is identified not simply as a document number, but as a specification for cadmium plating, enabling the system to flag potential compliance concerns related to hazardous substances.
- References to standards can be validated in real-time and flagged for correction if the standard has been updated or replaced.
- Long lead time processes can be highlighted for designers and sustainers so they can choose alternates or factor into their timelines.
- Requirements which can be misinterpreted are accurately understood; for example, a “cast” reference can be clearly interpreted as a “casting” process.
- Part numbers and manufacturers can be examined in advance of getting into the supply chain, enabling the avoidance of bigger and more expensive problems downstream in procurement cost and delays.
- Manufacturing resources such as required machinery, heat treatment ovens, and human operators can be anticipated and scheduled far in advance for greater efficiency.
- References to regulated substances and scarce critical minerals (e.g. PFAS, antimony, yttrium, gallium) are mapped to known supply chain risks, including export restrictions or obsolescence threats.

The result is a unified requirements graph – such as the SWISS knowledge graph – in which each reference to a standard in a drawing or model is no longer just a text string; it becomes a link to a specific, versioned requirement set with parameters and constraints available for reasoning, validation, and automation. In effect, the system builds a structured “requirements graph” around

each part or assembly: a network of linked entities representing geometric features, material choices, process steps, test methods, standards clauses, and supply-chain constraints. ***This graph is the model-based product definition***, and it enables downstream workflows in engineering, procurement, compliance, and sustainment to operate from a single, authoritative, machine-interpretable source of truth.



## Creating Complete Model-Based Product Definition

The result of this approach is the ability to generate a comprehensive, digital-thread-ready product model from legacy documentation. Rather than being limited to geometry or basic requirements, this model includes:

- Engineering requirements
- Materials and alloys
- Manufacturing processes
- Surface treatments and coatings
- Test protocols and inspection criteria
- Quality assurance procedures
- Regulatory and compliance indicators
- Supply chain risk markers
- Lifecycle sustainment considerations
- External industry standards (SAE, ASTM, ASME, MIL, etc.), represented as structured, reusable requirement sets

Practically, this means the MBPD is not just an annotated 3D model – it’s a combination of:

- A geometric model (features, tolerances, GD&T)
- A requirements model that captures every requirement, including those originating in external standards and process specifications

Together, these form a complete definition of what the product is, how it must be manufactured, tested, and under what constraints it can be procured and sustained. This model is not an abstract representation. It is a concrete, structured data object that can be integrated into enterprise platforms such as Windchill, TeamCenter, or 3DX, and can feed automated processes across the entire product lifecycle.

In a defense setting, complete product definition may need to support more than one output form. It should populate enterprise systems, support accessible repositories of interface specifications and supporting documentation, and, where needed, feed model-based representations such as SysML. The DSS is clear that the goal is not one mandated tool. The goal is usable formats that work across tools and organizations.

SWISS is built for that. It turns legacy documents into structured requirement data with clear meaning, traceability, and revision context, then delivers that data into the systems and models that need it.



## MBPD as the Front Door to Enterprise Execution

A complete MBPD is not the finish line. It is the point where engineering intent becomes usable across the enterprise. When requirements are structured and traceable – linked to standards, revisions, materials, processes, and test methods – procurement, quality, compliance, and sustainment teams can work from the same underlying definition the engineers are using, instead of reconstructing it later from analog documents.

The DoD’s AI Strategy reinforces this from another angle. It calls for federated data catalogs, exposed interfaces, and enough documentation for third-party integration without prime contractor support. Those goals are hard to meet when the governing requirements still live in scattered PDFs.

A complete MBPD gives those downstream efforts something they can actually use. When requirements and interfaces are structured and traceable, they can be cataloged, shared, checked, and reused across tools and organizations.



## Enabling a “Bill of Operational Features” (BoOF)

Once all aspects of the MBPD are represented in this structured way, organizations can go a step further: deriving a Bill of Operational Features (BoOF, coined by Jan de Nijs at Lockheed Martin).

A BoOF is essentially an operational view of the MBPD – a structured list of the functional and operational features that must be realized, controlled, or verified throughout manufacturing, test, and sustainment. It is enabled by, and dependent on, having all requirements model-based:

- If a heat treat requirement is defined only by a PDF, it cannot reliably appear in a BoOF.
- If a test protocol is only referenced textually (“per ASTM B117”), its operational implications can’t be systematically propagated.
- If a material requirement is not tied to supply-chain risk data, operational readiness cannot be accurately assessed.

By converting these requirements into structured, machine-readable objects that can be bound to parts and features, a BoOF can be automatically or semi-automatically generated for each part or assembly, providing a bridge between engineering intent and operational execution.



## Illustrative Use Case: Bringing Standards into the Model

Consider a legacy drawing for a critical aerospace component that calls out:

- Base material per AMS 4130
- Heat treat per AMS 2750
- Cadmium plating per QQ-P-416
- Salt spray testing per ASTM B117

In a document-only world, these are four textual references that require human interpretation. An engineer or quality specialist must chase down each document, interpret the applicable clauses, and manually ensure that processes and purchase orders align with the latest versions.

In a model-based world using full-stack document intelligence:

- AMS 4130 is parsed into a structured material requirement with explicit composition, mechanical property ranges, and acceptable process conditions.
- AMS 2750 is represented as a set of heat treat process parameters and instrumentation requirements, linked to furnace capability data.
- QQ-P-416 is interpreted as a cadmium plating requirement, automatically linked to hazardous substance rules and potential substitutes.
- ASTM B117 becomes a test protocol object with clear exposure duration, conditions, and acceptance criteria.

All of these are bound to the specific part and relevant features in the MBPD. If QQ-P-416 is superseded or cadmium usage is restricted in a target market, the system can automatically flag affected parts and propose alternatives. If furnace capability changes, the system can identify parts whose heat treat requirements can no longer be satisfied.

The result is a living, model-based definition that keeps pace with changes in standards, regulations, and capabilities, rather than an analog drawing that silently drifts out of date.



## Why It Matters: Outcomes and ROI

The implications of this capability are significant and measurable. Organizations that adopt this approach can realize tangible benefits across multiple dimensions:

### Reduced Cycle Times

By automating the extraction, classification, and structuring of product data, engineering teams can dramatically reduce the time required to review documentation, validate compliance, complete TDP reviews and supplier packages, and prepare content for downstream use.

## **Lower Labor Costs**

Manual rekeying, cross-referencing, and change management activities represent significant hidden costs. Automating these tasks can deliver labor savings of up to 65%, freeing engineers to focus on higher-value design and decision-making activities.

## **Improved Data Integrity and Fewer Errors**

By eliminating the human hand-off between documents and systems, organizations can reduce the likelihood of transcription errors, missed requirements, or version control issues that lead to rework or non-compliance.

## **Faster Time to Market**

Engineering delays often ripple through to procurement, manufacturing, and delivery schedules. By accelerating the front-end data preparation process, companies can bring products to market more quickly and with greater confidence.

## **Fewer Procurement Fire Drills and Faster Alternate Sourcing**

When requirements are structured and traceable – not buried in static text – procurement can evaluate alternates faster and with fewer iterations back to engineering. That reduces late-stage re-sourcing, supplier re-qualification, and schedule churn caused by restricted materials, obsolete parts, long lead-time processes, or supplier constraints discovered after release.

## **Greater Compliance and Risk Visibility**

By linking specifications to regulatory databases and supply chain intelligence, organizations can proactively identify and mitigate risks related to hazardous materials, obsolete standards, restricted export content, and supplier risk exposure (availability, capacity constraints, ownership, disruption) so teams can act before issues surface as late-stage rework or procurement delays.

## **Enhanced Readiness and Reusability**

Structured product definitions can be reused across programs and platforms, enabling better sustainment planning, fleet readiness, and long-term cost control. The same structure also improves sustainment by making it easier to identify substitute parts, materials, and qualified sources when designs are reused or supply conditions change.



## **Conclusion**

Achieving a true model-based engineering environment requires more than good intentions and a modern PLM system. It demands a commitment to transforming the unstructured legacy content that underpins every product and program. It also requires making engineering intent usable downstream, so procurement, suppliers, quality, compliance, and sustainment teams don't have to reinterpret analog documents to do their jobs.

The technology to do this exists today. With advanced document intelligence – powered by engineering AI, semantic understanding, and technical domain knowledge – organizations can finally bring their legacy documentation into the digital age.

Most importantly, it is now possible to include all requirements in that transformation, not just those already present in CAD and PMI. By converting external standards, material and process specifications, and operational constraints into structured, machine-readable data, organizations can create a complete model-based product definition that truly reflects how the product must be designed, built, tested, and sustained. That includes the sourcing and risk context that determines whether the product can be built and sustained on schedule – materials restrictions, long lead-time processes, obsolete parts, and supplier constraints.

By investing in a comprehensive model-based product definition, companies don't just create better models. They create better outcomes: faster decisions, lower costs, higher quality, greater readiness – and fewer late-stage procurement and supply chain surprises.

Human-readable documents will still matter, but in a model-based environment they cannot be the only place where the real requirements live.

The data you need is already in your possession. Now is the time to make it work for you.



## For More Information

Exiger's SWISS platform uses domain-specific, semantic AI trained with proprietary ontologies and data on parts, materials, and manufacturing processes. SWISS turns analog drawings, standards, and technical documents into structured, machine-readable product definition. Organized in the SWISS Knowledge Graph and available by API, that data can flow by reference into PLM/PDM, QMS, MES, and supply-chain workflows – helping teams catch outdated specs, regulated materials, critical mineral vulnerabilities, long lead-time processes, and supplier constraints before they become delays, rework, procurement emergencies, or compliance liabilities.

**Learn more at: <https://www.exiger.com/products/digital-thread-and-parts-intelligence/>**