Inspectioneering Journal

ASSET INTEGRITY INTELLIGENCE

Are Your Corporate Standards Ready for the AI Age?

Why Requirements Management Systems Are Becoming the Practical Foundation for Guided, AI-Enabled Project Execution

Jeannie Lewis, *Global Division Manager – Engineering Standards at Becht*

VOLUME 32, ISSUE 4

JULY | AUGUST 2026

Are Your Corporate Standards Ready for the AI Age?

Why Requirements Management Systems Are Becoming the Practical Foundation for Guided, AI-Enabled Project Execution

Jeannie Lewis, *Global Division Manager – Engineering Standards at Becht*

Artificial intelligence has made corporate standards exciting again.

That may sound strange to anyone who has spent years managing engineering specifications, inspection requirements, and corporate policies. Standards work is essential, but it is not usually described as exciting. Most organizations still manage standards primarily as controlled documents: PDFs, Word files, SharePoint libraries, companion guidance or overlays, and project-specific exceptions.

AI is forcing a different conversation. Organizations are asking whether AI can help users find requirements faster, summarize complex content, compare specifications, identify missing context, and preserve subject matter expert (SME) knowledge. Those goals are valuable, but they expose a practical challenge: finding the right information is not the same as knowing how and when to apply it.

The energy industry is appropriately cautious about AI in high-consequence environments because AI-enabled decisions require transparency, explainability, safety, reliability, data quality, governance, and human oversight. An "80% answer" may be useful in a general knowledge setting. In technical standards and engineering decisions, the missing context is often the part that matters most.

The quality of the standards data becomes part of the AI control framework. If standards are ambiguous, duplicated, outdated, or trapped in conflicting files, AI will inherit those weaknesses. If requirements are structured, current, contextualized, and traceable, AI can provide more reliable support.

In addition, we are now facing two parallel challenges:

1. **The retirement challenge**: Experienced standards owners know what the standards say, but also why requirements were added, when they apply, and which questions must be asked before giving an answer. As that experience leaves the workforce, organizations risk losing the judgment that connects with the written requirements and real-world decisions
2. **The AI adoption challenge**: Engineers and other technical personnel are already using AI tools to answer standards-related questions. Now organizations must decide whether to provide the governance and capabilities needed for reliable AI use without compromising asset reliability.

When it comes to AI, incorrect answers are not always obvious. They may often sound right and cite real sources, but miss the context needed to be applied correctly, such as process service or local regulatory requirements. That is why experienced SMEs often answer a question with another question. They are trying to understand the assumptions and conditions that determine the correct answer.

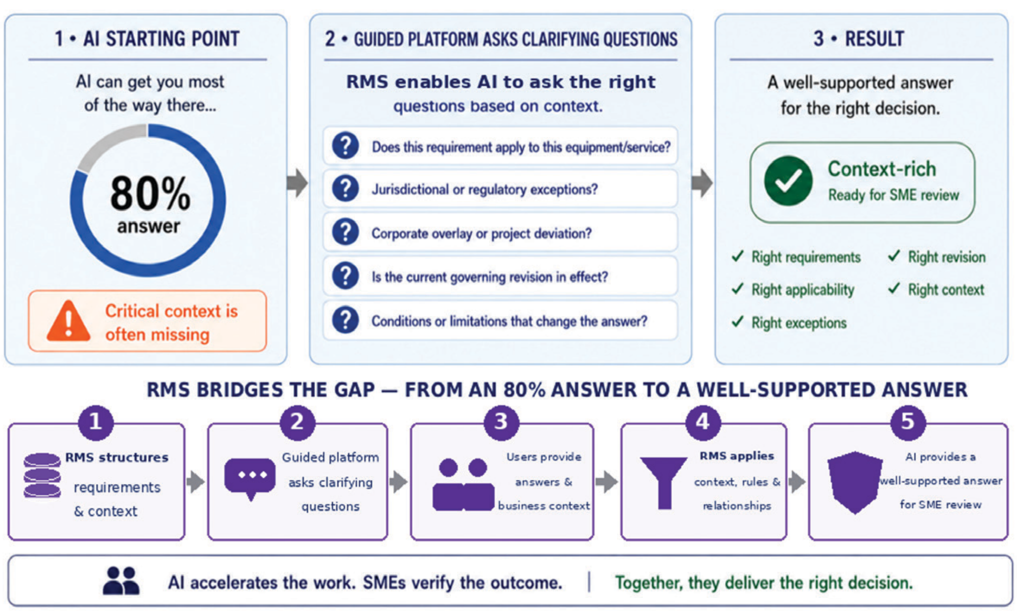

Figure 1. An "80% answer" may be useful for discovery, but technical decisions require context before the answer can be relied on.

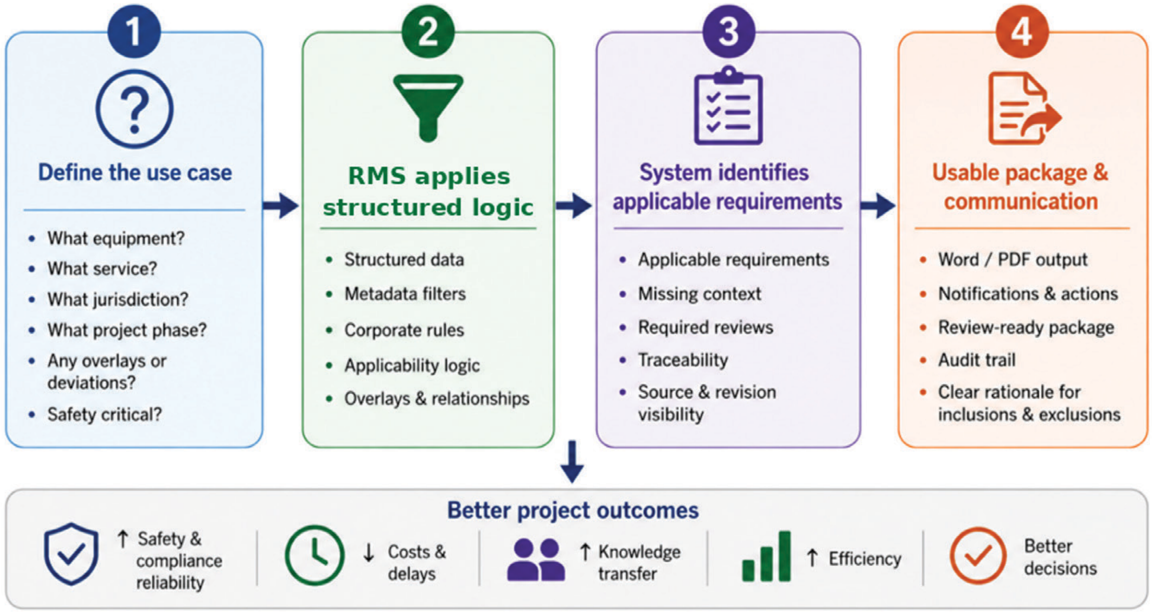

Figure 2. A context-aware RMS helps users define the use case, apply structured requirements logic, identify applicable requirements, and produce usable outputs while preserving traceability.

Many organizations will start by pointing large language models (LLMs), the AI technology used to interpret and respond to natural-language questions, at PDFs, Word documents, SharePoint sites, and historical project files. That can help with discovery, but it does not solve the harder problem of applicability. Rather than treating standards simply as content for AI to retrieve, the larger opportunity is to make standards easier to apply through guided workflows. These workflows lead users through the questions and inputs needed to determine which requirements apply to a specific situation, rather than relying on the user to know what context to provide upfront. At the same time, they create the structured foundation AI needs to support technical work responsibly.

That is where a requirements management system, or RMS, becomes important. An RMS is a system for managing individual requirements with associated structured, governed data that still allows users to view, review, publish, and export standards in familiar document formats.

In an RMS, each requirement can carry its own applicability logic, intent, interpretation, revision history, and similar associated metadata. Doing so allows the system to:

- Guide users through the required inputs
- Determine what applies
- Explain why it applies
- Preserve traceability
- Provide both users and AI tools with better context than prose alone

Figure 2 illustrates the shift from document retrieval to guided execution. The system walks the user through the use case, applies structured logic, identifies what applies, and produces an optimized, usable output.

Managing Standards Knowledge

Corporate standards contain invaluable knowledge built through decades of use and sweat equity. They capture operating experience, engineering judgment, reliability considerations, local interpretations, safety expectations, preferred practices, and experience-based decisions. But standards contain two kinds of knowledge: the words that are written down and the judgment required to apply them correctly. The latter is where the aforementioned retirement challenge matters.

It is difficult to capture, maintain, and apply this knowledge using traditional document-based systems. Some knowledge is written directly into specifications and procedures, which adds length and bloat. Some knowledge is buried in overlays, exceptions, review comments, and past project decisions—managed as separate specifications or draft copies in network folders. Some knowledge lives only in the brains of SMEs. And as engineers and inspectors retire, that institutional knowledge leaves with them.

Historically, organizations have tried to solve this within document-based systems in three ways, none of which is fully satisfactory:

- **Embed explanation in the specification**: Preserves context but adds length, complexity, and maintenance cost
- **Separate rationale into companion documents**: Keeps the base specification cleaner but creates a risk that the documents become out of sync as requirements change over time
- **Rely on experienced SMEs**: The most common approach, and the most fragile, because that knowledge goes away when they retire

Even when documented, many questions remain hard to answer in practice. Which requirements apply to this asset? What changed since the last revision? Which requirements are mandatory for regulatory, PSM, or other compliance reasons, and which may be

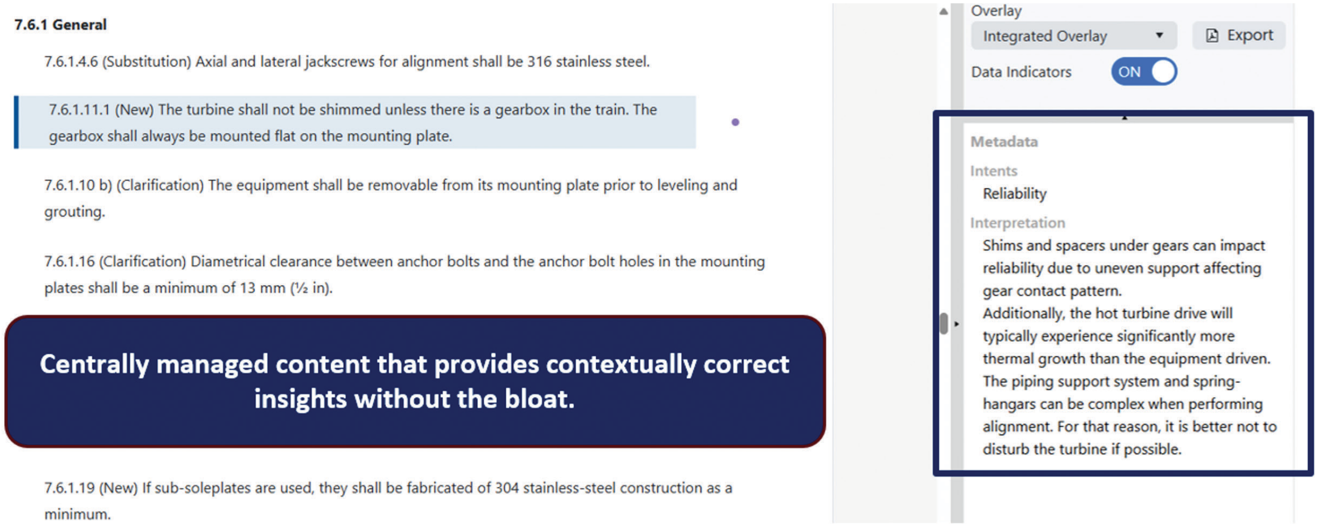

Figure 3. Metadata can keep specifications concise while allowing users to access the associated rationale or SME guidance when needed.

modified through an approved project deviation? What is modified by an overlay? What was removed from the project-specific package, and why?

A PDF or Word document may contain the requirement, but it usually lacks the structured context needed for AI or a less-experienced engineer to provide a response or apply it reliably. Applicability logic, hierarchy, rationale, overlays, exceptions, and relationships between requirements are often embedded in prose, scattered across related documents, or assumed by experienced SMEs.

Without that structure, the user has to know what context to provide before asking the question. That may be possible for an experienced SME, but it is not a reliable model for broader use with smaller teams lacking experienced SMEs, or for AI tools expected to support technical decisions.

An RMS changes how knowledge of standards is managed. Instead of managing only the document, the organization manages the individual requirements within the document, along with associated structured data - information captured in consistent, defined fields rather than only in free-form text. Think of a controlled picklist in Excel, where "Carbon Steel" is always recorded the same way instead of appearing variously as "Carbon Steel," "CS," or other variations. Each requirement can be supported with metadata such as applicability rules, asset type, service, jurisdiction, criticality, rationale, revision history, and relationships to other requirements. Because the associated data is structured, requirements can be more easily sorted, filtered, and related based on project context, while also providing AI tools with more consistent information to work from.

That structure helps the system guide engineers through the right inputs, generate optimized requirement sets, and make the basis for inclusion or exclusion easier to review. It also gives AI tools the context needed to ask better clarifying questions and provide more reliable support for technical use cases. In doing so, standards evolve from static documents into governed, contextualized requirements that improve quality and support reliable engineering decisions.

Capturing the "Why" with Knowledge Transfer

Technical requirements exist to transfer knowledge based on experience, such as prior failures, near misses, reliability issues, regulatory interpretations, site service conditions, or lifecycle cost considerations. Unfortunately, that reasoning is often the first thing lost when standards are reduced to short requirements or passed from one generation of engineers to the next.

When the rationale is missing, requirements can also become difficult to remove. A requirement may have been added because of a specific failure, local condition, or historical preference that no longer applies. But without the original reasoning, standards teams may struggle to determine whether the requirement still belongs. Traditional document-based approaches tend to force users to add more explanation into specifications and make them harder to use, or separate the rationale into companion material that can drift from the live requirement.

An RMS allows requirements to remain concise, while the explanations behind them are stored as metadata. The user sees a clean specification by default, and can access the associated rationale or SME guidance when needed. Engineers and standards owners can now understand not only what a requirement says but also why it exists and how that rationale should guide future decisions.

Easy access to metadata also supports different audiences. A third-party supplier may only need the requirement. An owner's engineer may need the requirement plus intent, interpretation, and compliance status. A standards committee may need change history and rationale.

> "Technical requirements exist to transfer knowledge based on experience, such as prior failures, near misses, reliability issues, regulatory interpretations, site service conditions, or lifecycle cost considerations. Unfortunately, that reasoning is often the first thing lost when standards are reduced to short requirements or passed from one generation of engineers to the next."

For AI, this is critical. An LLM grounded only in terse requirements may miss context or intent, while one exposed to uncontrolled commentary may confuse explanation with obligation. Structured metadata allows the system to distinguish requirement, rationale, guidance, interpretation, commentary, and verification. In short, on-demand metadata keeps specifications lean while preserving the insights behind them.

Guided Requirements Application and Optimization

Many organizations start by asking whether AI can allow people to find standards and requirements faster. Better search is useful, but it is not enough to help users determine which requirements apply. That is where guided requirements application becomes valuable, with or without AI.

Consider a simple question: "What is the industry standard for turbines?" A general-purpose chatbot may provide a broad answer. A machinery SME would recognize that the question is incomplete. Steam turbine or gas turbine? General or special purpose? Procurement, design, inspection, operation, or maintenance?

For inspection and mechanical integrity work, reliable standards application is both a risk management and efficiency issue. The wrong answer can affect inspection scope, repair planning, material selection, supplier requirements, jurisdictional compliance, or verification activities. These decisions shape project execution while influencing safety and operational performance. Similar to an experienced SME, AI needs the right context from approved, reliable sources to support technical decisions and determine when and how requirements apply.

A system using the right technical context, such as equipment type, key equipment attributes, or process services, should recognize that a user's prompt or question is incomplete and help to ask a better question rather than simply providing a response. This is where structured requirements change the workflow. In an RMS-enabled environment, the system can guide the user through one practical input at a time. Each answer determines which questions appear next and which requirements remain applicable.

For example:

1. A project user begins by selecting "pressure vessel" as the asset type
2. That selection invokes vessel-specific inputs such as design code, material of construction, design temperature, insulation status, orientation, and process service
3. Those inputs then trigger further information needed for the project, such as additional insulation details, internal component design, corrosion-under-insulation considerations, or other applicable bid-package requirements

This guided process helps users work through the approved rules and conditions within the standard that determine what applies more efficiently. Less-experienced engineers can see which inputs matter and why, while experienced users can move faster without relying on memory or manually checking multiple documents. In both cases, the system makes important context harder to miss.

Once the workflow has captured structured inputs, the RMS or AI can use those inputs to apply predefined rules and determine what should happen next. For example, can help identify incomplete information, flag unusual selections, challenge out-of-norm assumptions, summarize changes, compare supplier exceptions, and explain why a requirement was included. In that model, the RMS or AI supports a governed workflow using structured data rather than guessing from uncontrolled document fragments.

The value comes from making the rules and conditions within the standards explicit and structured. Once that logic is structured and governed, it can be applied consistently across repeatable workflows such as inspection planning, procurement packages, repair scopes, commissioning checklists, and project execution plans.

Applicability Leads to Better Project Requirements

Applicability is the difference between "this requirement exists" and "this requirement applies here." It turns a broad corporate standard into a usable project requirement set. In static documents, users determine applicability manually by reading long standards, interpreting notes and exceptions, checking overlays, and deciding what applies. That process consumes time, adds cost, reduces quality, and introduces variation.

When determining applicability, organizations generally take one of two approaches:

- A build-up approach starts with an empty package and adds requirements that appear to apply
- A build-down approach starts with the full baseline and removes what does not apply

For high-consequence work, a build-down approach can be more defensible because SMEs can review not only what remains in the package, but also what was excluded and why.

In a structured requirements environment, applicability can be tied

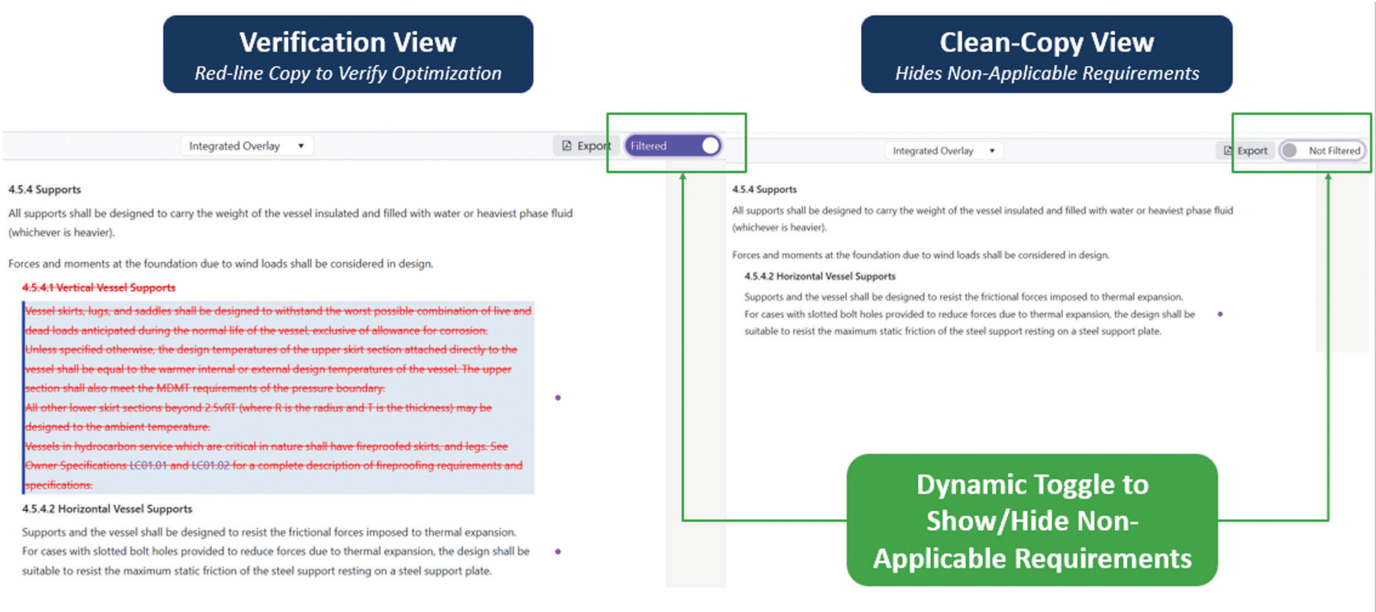

Figure 4. A build-down approach can create optimized project or bid packages while preserving visibility into what was included, what was excluded, and why.

to factors such as equipment classes, process service, jurisdictions, and other relevant project conditions., This allows teams to develop requirement sets that better reflect the actual scope of work. One practical example is project or bid package generation. Instead of sending suppliers every requirement, teams can issue optimized packages that remove non-applicable requirements while preserving the basis for what was included and excluded.

Cost savings from this type of optimization can be both direct and indirect. Reduced review time, fewer clarification cycles, and fewer change orders lower project costs, while better compliance with applicable requirements helps reduce missed obligations and improve asset reliability over time.

For inspection and mechanical integrity, applicability is especially important. Inspection requirements may depend on equipment history, damage mechanisms, process service, risk ranking, jurisdictional rules, prior repairs, and safety criticality. AI can help users navigate that complexity, but only if the underlying requirements have been structured with the right context.

RMS Capabilities to Consider

Organizations evaluating an RMS should look beyond document storage and ask whether the approach can support the way engineering requirements are actually managed and applied.

Some important considerations include:

- Managing requirements individually: Can requirements be identified, traced to their source, revised, and supported with information such as rationale and applicability?
- Determining applicability: Can project context, such as equipment, service, jurisdiction, material, risk, and site requirements, be used to help determine which requirements apply?
- Maintaining traceability: Can SMEs understand where a requirement came from, what changed, what was excluded or modified, and why?
- Supporting familiar engineering processes: Can standards owners and engineers continue to review, compare, approve, and publish requirements in a way that fit existing management of change processes?
- Guiding users through application: Can the approach help users identify information needed to determine applicable requirements and produce usable outputs for activities such as procurement, inspection, or project execution?
- Preparing information for AI use: Is information organized consistently, traceable to its source, appropriately controlled, and clearly separated between requirements, guidance, rationale, interpretation, and commentary?

The goal is not to force every activity into a single platform, but to help requirements move through the project and asset lifecycle without losing context needed to understand and apply them correctly.

Standards development organizations (SDOs) are also moving in this direction—increasingly recognizing that structured, machine-readable requirements are necessary for standards to function reliably in automated and AI-assisted environments.

Implementation

The most effective starting point for RMS implementation is a bounded set of standards or a recurring engineering task where the pain points and potential value are clear. Many corporate standards collections include hundreds of interconnected documents with cross-references, overlays, and discipline dependencies, making an enterprise-wide rollout impractical. At the other end of the spectrum, converting a single specification in isolation rarely reflects how standards are actually used. For example, an organization might start with procurement specifications that are repeatedly used to create bid packages, or with an inspection and test plan workflow that depends on a defined set of standards and applicability rules.

A practical first phase should be large enough to support a real business process, but focused enough to validate the approach.

1. Structure the content: Identify requirement statements, separate requirements from guidance or commentary, extract metadata, identify internal references, and preserve links back to the source document
2. Validate the logic: SMEs review applicability rules, overlays, exclusions, rationale, and required inputs so the system reflects approved standards logic
3. Prove the output: Generate a usable package, such as a bid package, inspection plan, or review-ready requirement set, and let users see what was included, what was excluded, and why

The digital conversion effort is a legitimate concern, but it should not be viewed as purely a manual exercise. Some platforms can help accelerate the first pass by identifying candidate requirement statements, suggesting tags, and preserving links back to the source document. SMEs still need to validate the structure, applicability, and logic, but they should not have to start from a blank screen.

This phased approach matters because AI is already being used with corporate standards, whether organizations are ready or not. The outcome depends on whether AI will rely on uncontrolled PDFs, outdated project files, and scattered SME notes – or on structured, governed requirements with the context needed to apply them correctly. Critically, the same structured foundation that supports AI-enabled project execution also positions organizations for integration with inspection databases, maintenance systems, and emerging industry data standards—but that value starts with getting the requirements right.

Conclusion

Better standards data and guided workflows help project teams today, regardless of whether or not AI is involved. They reduce rework, improve communication with suppliers, support jurisdictional compliance, and preserve the “why” behind requirements before that knowledge walks out the door.

More importantly, they help ensure requirements are easier to understand, apply, trace, and maintain, while preserving the context needed for reliable technical decisions.

Organizations that prepare their standards for both people and AI will realize the greatest value. If you are already evaluating AI tools for technical work, now is the right time to assess standards readiness. ■

For more information on this subject or the author, please email us at inquiries@inspectioneering.com.

CONTRIBUTING AUTHOR

Jeannie Lewis
Ms. Lewis is the Global Division Manager - Engineering Standards at Becht, where she leads BechtPractices®, an advanced Requirements Management System that enables organizations to digitally transform and improve how they access, manage, and use corporate engineering standards and specifications. With 20 years of industry experience, Ms. Lewis specializes in product strategy, customer success, enterprise implementations, and engineering knowledge management. She partners with organizations worldwide to develop, implement, maintain, and continuously improve digital engineering standards programs, helping modernize how they govern and apply critical engineering knowledge across global operations.