

Foundations for MEP AI:

Data, Governance, and AI-Enabled Design Processes

Paper 3 of 4: Technology & Implementation

Authors:

- Daniel Månsson - Chief Digital Officer - Bengt Dahlgren, Sweden
- Davor Stjelja - AI Lead - Granlund, Finland
- John Coroner - Chief Operating Officer - Ethos Engineering, Ireland
- Niels Vercaemst - Technical & Innovation Director - Ingenium, Belgium



FIRST Q

Executive Summary


Engineers across Europe use AI daily for specification writing, report generation, and project research. This is valuable work that frees engineers for engineering. But the gap between this productivity-layer usage and AI that can validate an HVAC design or predict cross-discipline clashes remains wide, and it is not an algorithm problem. It is an infrastructure problem: data foundations, governance frameworks, and process structures that most firms have not yet built.

This paper takes two committed positions. First, shadow AI should be channeled, not banned. Firms that block AI tools drive usage underground where it cannot be governed. Second, general LLM chat is necessary but not sufficient for serious MEP AI. Basic AI tools build the baseline literacy every firm needs, but without integration to BIM and calculation engines,

retrieval over the firm's own project data, and built-in audit trails, general chat remains productivity assistance, not engineering AI.

Starting from the messy reality of how MEP firms actually operate (fragmented data, no single source of truth, unmanaged shadow AI, contracts that never contemplated AI), the paper defines what "AI-ready" data and governance look like, maps AI capability from production-ready to emerging, and establishes three governance guardrails that every firm must build. EU regulation is simultaneously building the data highway: five adopted regulations mandate the machine-readable infrastructure that AI needs, creating a dual dividend for firms that invest in integrated data infrastructure rather than fragmented compliance systems.

Paper 4 will translate these foundations into business strategy, investment frameworks, and phased roadmaps.



Why Foundations Matter Now

Zero. That is the number of Nordic construction companies that qualified as “future-built” in BCG’s 2025 maturity assessment [1]. Not one, despite the fact that over 60% of Nordic construction executives have implemented at least one AI use case. A quarter of them report meaningful business impact. The rest have adopted AI without building the infrastructure to make it count.

This pattern is not unique to the Nordics. Among the 27% of AEC firms actively using AI for core tasks, 68% report measurable productivity gains, with nearly half recovering 500 to 1,000 engineering hours per firm annually—on the order of one to three percent of total engineering capacity at a typical mid-sized firm [2]. Global AI spending hit \$301 billion in 2026, up from \$223 billion

the year before [3]. Across the broader construction sector, though, 45% of professionals report zero AI implementation, even as 72% of enterprises overall have at least one AI workload in production [3], [4]. Among engineering employers, 58% report using AI but only 18% use it regularly [5]. The gap between nominal adoption and active integration defines where we are right now.


What matters more is what engineers are using AI for: quality assurance, project tracking, specification review, report generation, material classification. Real tasks that deliver measurable value. What AI is not yet doing is the engineering design itself: routing piping systems, sizing equipment for specific buildings, optimizing cross-discipline coordination. Automating administrative

work is genuinely valuable. But AI’s potential extends into the engineering work itself, and that next step requires foundations that most firms have not built.

Paper 1 in this series established the strategic opportunity: MEP firms hold decades of semi-structured domain knowledge that technology companies cannot easily replicate, and EU regulation is building the data highway whether firms intend to use it or not. Paper 2 described the human dimension, the evolving hybrid roles where engineers become “Tool Builders” who capture domain knowledge into custom automations and “Design Orchestrators” who steer AI-augmented workflows while maintaining engineering rigor. Both papers were right about the direction. But the pace of change has outrun the framing.

Why Foundations Matter Now

The bottleneck is not algorithms or ambition. It is the data, process, and governance infrastructure that separates productivity assistance from credible engineering AI. Without that infrastructure, no MEP AI product can be credible regardless of the underlying model. This paper addresses how to build it, starting from the mess.



Starting From Reality

The Data Mess

If this paper started from a greenfield scenario, the advice would be simple: structure your data, set your governance policies, deploy. But no MEP firm starts from greenfield. The reality is decades of project data scattered across BIM models, PDFs, spreadsheets, emails, scanned documents, CRM systems, and proprietary calculation tools. Naming conventions differ between offices, between teams, and between individual engineers. Ownership is unclear. Manual workarounds are everywhere.

Consider scenarios that will sound familiar. A bid emailed out under deadline pressure, revised offline, never saved back to the shared drive. The team works off stale numbers for

days before anyone notices. A pipe-sizing spreadsheet passed person to person with silently changed cells, each engineer trusting the one before them to have checked the formulas. Project emails archived in personal inboxes rather than project mailboxes, invisible to colleagues and invisible to any AI system that might try to learn from them. One firm conducting a data governance audit interviewed 40 senior staff and discovered that no standardized naming convention existed across the organization. Another, deploying an HR system, found 86 different job roles for 220 people.

This is the starting point for AI. Research confirms it is widespread: professionals in construction waste an estimated 20% of work time searching

for or correcting information, and better data management alone could save 10 to 15% of total construction costs [6]. The “corpus problem,” the fact that AI needs thousands of well-structured datasets for training while the industry lacks shared datasets at that scale, is the central barrier to AI-BIM integration [7]. Many BIM models are technically valid but practically useless for AI because they are incomplete, inconsistently modeled, or lack the semantic richness that machine learning requires. The top domain-specific barriers to AI adoption in construction are construction-specific knowledge gaps, workflow integration challenges, and unstructured or heterogeneous data [8]. All infrastructure problems, not algorithmic ones.

Starting From Reality

Shadow AI & The Vibe Coding Question

Against this backdrop, individual engineers have discovered clear value using personal AI tools. Their top use cases are consistent across every study: information research, email drafting, summarization, brainstorming, and translation [2], [3]. These are non-engineering tasks that deliver immediate, visible time savings. But each use of an unmanaged AI tool on client data creates compliance, security, intellectual property, and quality risks that the firm cannot monitor, govern, or learn from [9].

The broader AI market is spending heavily (an estimated \$539 billion in capital

expenditure in 2026), but the gap between investment and realized value is widening [10]. Enterprises now require upfront financial validation and defined ROI timelines before approving AI initiatives, moving from emotion-based to fact-based governance [11]. For MEP firms, the foundations described in this paper are not optional preparation. They are the precondition for any AI investment that boards will approve.

This paper takes a committed position: shadow AI should be channeled, not banned. Firms that block AI tools drive usage underground where it cannot be monitored. The productive governance posture is providing enterprise alternatives with clear boundaries.

A related problem deserves attention. Domain-expert engineers, not software developers by training, are using AI tools to build small applications by describing what they want in natural language and letting AI generate entire features. This is a positive development: we want engineers building tools that solve their domain problems. The goal is domain-expert AI-assisted coding, where engineers use AI tools within approved templates, security patterns, and review processes to create genuinely useful automations.


Starting From Reality

The risk is unstructured “vibe coding,” building entire applications without understanding the generated code, without following the firm’s security or architectural standards, and without assigned maintenance ownership. When that person leaves, the application becomes unmaintainable. This recreates the shared-spreadsheet

fragmentation described above, at higher velocity and greater complexity. The same dynamic is now reshaping software at large: a “SaaS apocalypse” debate has emerged as AI-generated bespoke tools fragment once-consolidated platforms. MEP firms can avoid that trajectory only by establishing structure before fragmentation sets in. The answer is not prohibition but

structure: approved components and templates so that domain experts can build tools productively within a framework that maintains quality. This is a governance question, addressed later in this paper.

But before governance comes a more fundamental question: what does the target state for data actually look like?



What “AI-Ready” Data & Governance Look Like in MEP

Who Owns The Data?

Ask five people at any MEP firm who owns their project data and you will get five different answers. The client typically owns the final deliverables. The consultant retains intellectual property in their methods and design approaches. Suppliers own their product data. None of these boundaries were drawn with AI in mind.

Old contracts never contemplated AI use of project data. New contracts are inconsistent: some clients grant broad data reuse rights, others restrict any use beyond the specific project, and the terms vary between customers and across jurisdictions. The industry has not settled this question, and individual firms must

develop a documented position before they build AI systems on client work.

Internally, the ownership question is equally unresolved. Who ensures that the firm’s own data (naming conventions, classification standards, historical project records) is correct, up to date, and consistently maintained? Role-based data ownership is the practical answer: every data domain needs a named owner, the role accountable for ensuring that the single source of truth for that domain is current, complete, and consistently maintained. Without this, data quality degrades through diffusion of responsibility. Most firms have not formalized it.

The Practical Target State

What does “AI-ready” MEP data actually look like? Not perfection, but a structured baseline that AI systems can consume. File names that a machine can parse into project, discipline, system type, and version. Building elements classified consistently using IFC types, UniClass, or OmniClass. Standardized formats for calculations, specifications, and reports that encode the firm’s best practice rather than individual engineers’ personal preferences. A systematic approach to consolidating historical project data into formats AI can learn from, not all at once but progressively, starting with the highest-value project types. Logging and traceability for every AI

What “AI-Ready” Data & Governance Look Like in MEP

interaction: who asked what, using which model, based on which sources, and what validation was performed. Project data queryable through APIs, not locked behind proprietary software that must be opened manually.

High-quality, structured data is the prerequisite for trustworthy AI. Inconsistent formats, data silos, and incomplete records undermine AI effectiveness regardless of model sophistication [12]. Governance itself must be framed as enablement: centralized standards with federated domain-level accountability, lifecycle checkpoints from data documentation through deployment to monitoring [13].

Standards Licensing: A Structural Barrier

One barrier deserves specific mention because no individual firm can resolve it. Current standards licensing arrangements (engineering standards published as PDFs, sold under per-user fees, with no AI-use clause) make it effectively impossible to build grounded, code-compliant MEP AI. An AI system checking compliance against ASHRAE 90.1 or EN 16798 needs machine-readable access to those standards with explicit permission for AI processing. The infrastructure for this barely exists, and individual firms are only beginning to negotiate bespoke arrangements. Standards bodies, professional associations, and regulators must address this collectively.



What “AI-Ready” Data & Governance Look Like in MEP

EU Regulation as a Data Accelerant

When regulators mandated XBRL for financial filings, the initial reaction was compliance cost. Within a decade, that structured data became the foundation for AI-driven analytics. XBRL now underpins 216 regulatory frameworks worldwide. Construction is at the same inflection point.

Five adopted EU regulations collectively mandate the same machine-readable data infrastructure that AI requires, creating a dual dividend for firms that recognize it [14], [15]. The Ecodesign for Sustainable Products Regulation (ESPR) mandates Digital Product Passports in JSON-LD format, with construction product

DPPs expected by 2028 to 2030. JSON-LD adds semantic context that resolves every data term to its official EU definition, eliminating ambiguity when comparing products across manufacturers. The revised Construction Products Regulation (CPR) creates a construction-specific data dictionary under Article 78, designed for machine readability and BIM interoperability. The EPBD recast introduces the Digital Building Logbook, a repository for energy performance certificates, renovation passports, and lifecycle carbon data, with member states transposing by May 2026. CSRD requires sustainability reporting in XHTML with Inline XBRL tags. Finland’s Rakentamislaki mandates IFC 4.0.2.1 submission from January

2026, with the national Ryhti system providing API-accessible standardized building data.

The timeline functions as an AI readiness roadmap. In 2025 and 2026, IFC mandates and GWP reporting go live in the Nordics. In 2027 and 2028, DPP standards, ESAP, and Digital Building Logbook frameworks arrive. By 2028 to 2030, construction product DPPs and full lifecycle reporting kick in. Firms that treat each deadline as an isolated compliance task will build fragmented systems. Firms that invest in integrated data infrastructure now will find their AI capabilities compounding with each regulatory phase.

What “AI-Ready” Data & Governance Look Like in MEP

The regulatory picture extends beyond data mandates. The EU AI Act introduces risk-based classification for AI systems, requiring traceability, human oversight, and documentation for higher-risk applications. Most MEP design tools will classify as minimal or limited risk, but operational building management systems controlling HVAC or fire suppression in occupied buildings may qualify as high-risk under Annex III [16]. Firms must track this evolving classification alongside their data infrastructure investments.


With the data target state defined and the regulatory tailwind mapped, the next question is process: how should MEP design workflows change to accommodate AI?



What “AI-Ready” Data & Governance Look Like in MEP

AI-Enabled Design Processes

AI is, to a large extent, programming commoditized for non-programming people. Every AI-augmented MEP workflow, from automated specification review to calculation validation, is ultimately a software pipeline: data in, processing, data out. Many of the conventions that make software development reliable (version control, review gates, traceability, automated testing, operational monitoring[1.1]) are directly relevant to making AI-enabled MEP workflows reliable. This is not about MEP imitating software culture. It is about recognizing that AI-enabled engineering work is software and needs structural discipline.



What “AI-Ready” Data & Governance Look Like in MEP

Some of these processes already happen informally in MEP firms. Stage gates function as review flows, option comparisons happen in meetings, revision logs capture some design history. But none of these operates in a structured, machine-readable way that AI systems could use, and that absence is itself a barrier.

In software development, branching and merging allow teams to explore alternatives without losing the main line of work. MEP design lacks this formally, though engineers intuitively do something similar when they save multiple versions [t2.1] of a BIM model or spreadsheet. The aspiration is explicit version control where design options are tracked, compared, and merged with full

traceability, so that an AI system can learn not just from the final result but from the exploration that produced it.

Review flows need the same treatment. Who approves what, at which stage? In most MEP firms, this is understood tacitly but not formalized in a way that AI systems can consume. Making review gates explicit and structured enables AI to understand when a design decision is tentative versus approved, and what validation has been performed.

A particularly interesting gap is the disconnect between design discussions and computational validation. An engineer decides to change a pipe diameter in a meeting. The calculation is updated separately. The BIM model is updated separately

again. Emerging MCP-connected systems link natural language conversations to validated simulation engines, making these connections explicit.

An engineer describes a building modification in natural language; the AI translates it into structured simulation parameters; the physics engine runs the calculation; the result returns with full traceability.

The AI never performs the physics[t3.1]; it orchestrates the exchange between human intent and validated computation [17], [18]. This pattern has been demonstrated at national laboratory scale, with frameworks achieving 50% reduction in computational cost through specialized agents coordinating across building systems.

What “AI-Ready” Data & Governance Look Like in MEP

Then there is the un-logged decision problem. Critical design choices made in meetings or emails, the rationale understood by those present, future engineers unable to trace why a particular system configuration was chosen.

When a project spans years and engineers rotate, lost rationale forces expensive re-evaluation of decisions that were already well-reasoned. Structured logging of why decisions were made, not just what was decided, is both a governance requirement and an AI training asset.

MEP design has fundamentally different coordination patterns than software development: longer timescales, different feedback loops, different regulatory checkpoints. The relevant analogy is to specific

software development tools (version control, review gates, audit trails, automated validation pipelines) that address universal engineering needs for traceability and quality assurance, not to software development methodologies. The emerging discipline of systematic AI lifecycle management provides the closest analogy to what MEP design processes could become [19], [20]. That 87 to 90% of AI initiatives stall before deployment, not due to technical performance but due to lack of operational infrastructure, makes it clear: process foundations matter more than algorithmic sophistication.

Paper 2’s hybrid roles connect here: “Tool Builders” are

the engineers who create and manage these AI-augmented workflows; “Design Orchestrators” work on top of them where AI contributions are documented and traceable.

These process structures are not theoretical. The next section shows what they make possible, starting from demonstrated value.



Serious MEP AI Capability: Value First, Features Second

Where Values Shows Up Today

An AI workflow reads technical documents (structural diagrams, product data sheets, equipment specifications), extracts relevant parameters, and classifies them against project-specific standards. The system automates what was previously manual data preparation. Engineers review and correct the exceptions. The design principle that matters: traceability is built in from day one. Every AI decision is logged with reasoning and source references, but this requires deliberate system architecture, not default AI behavior.

In another case, an AI agent checks supplier technical sheets against project specifications. Stainless steel required? Is it mentioned in the technical

sheet? The cost: approximately five euros in tokens per batch review versus a junior engineer's full day of cross-referencing specifications, checking dozens of technical sheets against project requirements. The quality: comparable to a one-to-two-year junior engineer. Not perfect, but a reliable first pass that catches the majority of issues before human review.

At the enterprise level, a company-wide LLM deployment provides all models to all employees, token-based, at 15 to 45 euros per person per month. [t4.1] Survey-measured time savings: approximately one hour per week per engineer. Top use cases match the shadow AI pattern: information research, emails, summarization, brainstorming, translation. The marginal cost per user is low

enough that even modest time savings justify the investment. The real value is strategic: reduced shadow AI risk and baseline AI literacy across the entire workforce, the prerequisite for everything that follows.

This is the paper's second committed position: general LLM chat is necessary but not sufficient for serious MEP AI. The company-wide rollout is the foundation. It builds baseline AI maturity, gives every engineer practical experience, and eliminates the shadow AI problem. But without integration to BIM and calculation engines, retrieval over the firm's own project data, and built-in audit trails, general chat remains productivity assistance, not engineering AI.

Serious MEP AI Capability: Value First, Features Second

The Honest Counter Thread

Not all evidence shows productivity gains. One rigorous randomized controlled trial found that experienced software developers were 19% slower when using AI coding tools, despite perceiving a 24% productivity speedup [t5.1][21]. A controlled design experiment found that generative AI improved novice performance but reduced general creative self-efficacy across all users, and produced no statistically significant overall performance advantage [22]. Six AI models tested on the same MEP pump station project produced heating load results ranging from 25 to 150 MBH, a sixfold variance that demonstrates why unvalidated AI output is not engineering [23].

AI is not a universal accelerant. Its value depends on the task, the integration quality, and whether the foundations are in place. The evidence suggests AI works best where the task is well-defined, the validation criteria are clear, and the data is structured, which is precisely the target state this paper defines.

What Makes MEP AI “Serious”

Beyond general productivity, serious MEP AI requires four vendor-neutral capability features. First, robust integration with BIM/MEP tools and calculation engines: not a standalone chat window, but AI that reads and writes to the same models and databases engineers use, connected to validated physics engines for

any calculation that matters. Second, retrieval over the firm’s own data, grounded in the firm’s project history, standards, templates, and accumulated expertise rather than generic internet training. A question about how the firm has historically sized chillers for hospital projects should draw on the firm’s actual project database, not a statistical average. Third, built-in governance: role-based access, usage logging, audit trails, and approval workflows embedded in the tool, not bolted on as a policy afterthought. Fourth, traceability: every AI-assisted output links back to its sources and reasoning. This is not optional. It determines whether liability caps hold and insurance coverage applies.

Serious MEP AI Capability: Value First, Features Second

The Capability Picture

The gap between what AI can do in research and what is commercially available defines the infrastructure challenge. Clash detection and filtering (AUC 0.972), HVAC fault detection at scale (15,000+ buildings, up to 25% energy cost reduction), compliance checking (97% precision on building codes), and estimating and quantity takeoff (95% scale detection accuracy) are all commercially shipping with proven ROI [2], [24]. One tier down, physics-informed neural networks achieve sub-1% deviation from reference solutions in thermal modeling, reinforcement learning delivers 22% average energy savings across 126 building studies,

and energy modeling surrogates replicate simulations within 0.07% error at 60 times the speed, all peer-reviewed and demonstrated in real buildings but not yet embedded in commercial MEP tools [24], [25]. At the emerging edge, MCP-connected simulation workflows, LLM-to-simulation pipelines achieving 100% accuracy on validation cases, and multi-agent design systems coordinating across building disciplines show architectural promise, with early implementations from national laboratories and universities [17], [18].

The gap between “validated” and “commercially available” in this picture is exactly the infrastructure gap this paper

addresses. Closing it requires data foundations, process structures, and governance frameworks, plus the integration architecture that makes validated capabilities deployable in practice [6], [26].



Serious MEP AI Capability: Value First, Features Second

FIGURE 3 • THE MEP AI CAPABILITY MAP

What MEP AI can actually do — today and next

From production-ready to emerging. The distance between “validated” and “commercially available” is precisely the infrastructure gap this paper addresses.

PRODUCTION-READY

[2] [24]

Commercially shipping
with proven ROI today

Clash detection
& filtering

0.972
AUC

HVAC fault
detection at scale

15-25%
energy-cost reduction

Code-compliance
checking

97%
precision on building codes

Estimating &
quantity takeoff

Shipping
automated, in commercial tools

VALIDATED

[24] [25]

Peer-reviewed in real builds
— not yet in commercial tools

Physics-informed
neural networks

<1%
deviation vs. reference (thermal)

Reinforcement-
learning control

22%
avg. energy savings · 126 buildings

Energy-model
surrogates

0.07%
error · 60× faster than simulation

EMERGING EDGE

[17] [18]

Architectural promise; early pilots
from national labs & universities —

MCP-connected
simulation flows

human intent → validated physics
NOT yet commercial

LLM-to-simulation
pipelines

100% on validation cases (pilots)

Multi-agent
design systems

cross-discipline coordination

Serious MEP AI Capability: Value First, Features Second

Eight Diagnostic Questions

How ready is your firm? These questions cut across data, capability, and governance:

1. Can you find all district heating projects your firm has completed in the last decade? If the answer requires asking individual engineers, you lack basic data discoverability.
2. Are your BIM deliverables in validated IFC format with consistent classification? IFC compliance is both a regulatory requirement and an AI prerequisite.
3. Where are your engineers using AI today, and do you know? If you cannot answer, you have shadow AI.
4. Who reviews AI-generated content before it reaches a client? Informal review is not a protocol.
5. Can an AI system query your project database through an API? If accessing historical data requires opening proprietary software, it is not AI-accessible.
6. Does your professional liability insurance mention AI? Check the endorsement schedules, not just the base policy.
7. Have you measured time savings or error reduction from AI tools your engineers use? Perceived value is not measured value.
8. If a regulatory auditor asked “was AI involved in this calculation?”, could your project documentation answer that? The audit trail determines whether liability caps hold.

The contrast with “just use Claude” is stark. A general LLM chat handles the admin tasks that appear in every survey: research, emails, summaries. It cannot do the engineering tasks in the capability tiers above without the integration, retrieval, and governance features that make outputs trustworthy. That is the capability bar.

Guardrails, Decision Rights, & Accountability

Governance makes AI deployment professionally responsible. Without it, firms are flying blind while the liability rules shift against them. Three guardrails are load-bearing.

Data-Training IP Rights & Audit Trail

Before training or grounding AI on client project data, firms need explicit contractual clarity on reuse rights. Old contracts never contemplated this. New contracts are inconsistent between customers and countries. The industry has not settled it. Firms need a documented position before they build any proprietary AI system on client work.

The level of documentation should be proportional to AI's role. When AI is a minor

workflow tool, similar to how engineers already use simulation software without documenting every keystroke, a lightweight record suffices: which tool, what task, who validated. When AI generates substantial content that directly appears in deliverables, more detailed documentation is warranted: inputs, outputs, validation performed, and who approved the result for professional use. The principle: an auditor should be able to understand AI's contribution to any deliverable, but documentation overhead should not exceed what we already accept for other engineering tools. "Human Oversight Warranties" are emerging as conditions of AI-related coverage; the audit trail is the evidence that those warranties are met [27], [28].

AI Application & Supplier Approval

A named role or committee must approve which AI applications and vendors are permitted, with documented rationale. Modern AI products are not just models. They are integrated systems combining language models, tools, data pipelines, and interfaces. Approval should evaluate the complete application, not just the underlying model. This includes explicit criteria for when new applications are evaluated and how existing approvals are reviewed.

Guardrails, Decision Rights, & Accountability

Equally important: an explicit stance against vendor lock-in. Today's best model may not be next month's. OpenAI's enterprise LLM market share dropped from 50% in 2023 to 27% in 2026 [29]. A 2026 survey of 542 C-level executives found that 74% would face day-to-day disruption if their primary AI vendor became unavailable, and of those who attempted vendor migration, 58% faced failures or significantly greater effort than expected. One documented migration cost \$315,000 for 40 AI workflows after a vendor failure [30], [31]. Infrastructure locked to a single provider is a liability.

The practical stance: model-agnostic architecture, portable data, and open formats. Build on open integration standards

rather than proprietary vendor protocols.

Shadow AI Channeling & Vibe Coding Governance

Enterprise AI alternatives must be provided for any workflow involving client data, with clear boundaries on what enters external AI systems and reported routes for discovering new tools. The strategic response is the framework Paper 2 introduced: Govern, Map, Measure, Manage [32].


Separate governance is needed for domain-expert AI-assisted coding (engineers building tools within approved templates, review processes, and deployment pipelines) versus unstructured vibe coding (building applications

without understanding the generated code or following any development standards). Domain-expert AI-assisted coding should be enabled and encouraged. It is how engineering knowledge becomes reusable tooling. But it requires an approved component and template library, a security and brand review gate before deployment, maintenance ownership assigned before any tool goes live, and a rule that no AI-generated application touches client deliverables without review—the structural counterweight to the “SaaS apocalypse” pattern flagged earlier.

Guardrails, Decision Rights, & Accountability

Who Decides What

Every firm answers these questions differently, but every firm must answer them. Data ownership must be role-based, with clear accountability for naming conventions, classification standards, and quality. AI application and supplier choices require a cross-functional decision involving IT, engineering leadership, legal, and data governance. The engineer of record remains accountable for AI-generated content in deliverables. AI does not reduce professional liability [33]. The governance framework must specify handoffs and escalation paths, not create a committee that approves everything.



Guardrails, Decision Rights, & Accountability

The Liability & Insurance Reality

The urgency behind these guardrails is concrete. Insurance companies are actively withdrawing AI coverage. ISO/Verisk introduced standardized AI exclusion endorsements (CG 40 47 and CG 40 48) effective January 2026, broadly excluding bodily injury, property damage, and personal injury arising out of generative AI [34]. The phrase “arising out of” requires only a causal connection, not proximate cause, significantly lowering the exclusion threshold.

Standard Nordic engineering contracts (KSE 2013 in Finland, ABK 09 in Sweden, NS 8401 in Norway, ABR 18 in Denmark) contain zero provisions addressing AI use. Liability

caps in these contracts protect consultants under ordinary negligence, but if courts determine that deploying AI without adequate verification constitutes gross negligence, those caps evaporate entirely. The engineer of record retains full liability regardless of AI involvement.

The result is a three-way mismatch: the AI vendor bears effectively zero liability through clickwrap license agreements that cap liability at fees paid; the engineer bears full professional liability; and the insurer is withdrawing the safety net. This is not a future concern. It is the current state as of 2026 [35]. The governance guardrails described above are the practical response: documentation and audit trails

that demonstrate responsible AI use, vendor diversification that limits dependency risk, and contractual clarity that addresses data rights before they become disputes.



Looking Ahead: From Foundations to Shared Platforms

When data is structured, governance is in place, and design processes are traceable, a different class of MEP AI becomes viable. Not standalone chat tools that help with emails, but integrated systems that retrieve the firm's actual project history to inform new designs, grounded in the firm's accumulated expertise made searchable and reusable. Systems that connect to validated calculation engines through structured integration protocols, ensuring that AI-assisted design decisions are grounded in physics, not language patterns.

Systems that generate auditable design documentation satisfying both regulatory requirements and professional liability standards, making AI-

augmented deliverables more defensible than traditional ones. And systems that learn from every project, so that the data asset compounds over time, each project making the next one better.

No single firm needs to build all of this alone. The foundations described in this paper (data standards, governance frameworks, integration protocols) are largely pre-competitive. Shared data standards and classification systems reduce the cost of making data AI-ready for every firm individually.

Joint governance templates codify best practice for AI audit trails, vendor evaluation, and data training policies, so each firm does not reinvent the

wheel. Ecosystem platforms provide common integration infrastructure while preserving each firm's proprietary domain knowledge and competitive advantage.

What clients and investors should reasonably demand from any "next-generation" MEP AI solution: integration with existing engineering tools, retrieval over the firm's own data, built-in governance and traceability, vendor independence, and demonstrated value on real engineering tasks, not marketing claims built on general-purpose AI benchmarks.

Looking Ahead: From Foundations to Shared Platforms

Paper 4 will translate the infrastructure defined here into business strategy: investment frameworks for sizing and

sequencing infrastructure spend, phased roadmaps aligned with the EU regulatory timeline, and positioning

strategies for the three strategic paths defined in Paper 1. The foundations come first. The strategy builds on them.



References

1. Boston Consulting Group, "Digital and AI in Nordic Construction," BCG, Aug. 2025.
2. Bluebeam, "AEC Technology Outlook 2026," 2026.
3. S. Reddy G, "AI Adoption Statistics 2026: 67 Enterprise and SMB Data Points." 2026. Available: <https://medhacloud.com/blog/ai-adoption-statistics-2026>
4. RICS, "AI in Construction Report 2025," Royal Institution of Chartered Surveyors, 2025.
5. Institution of Engineering and Technology, "IET skills for a digital future 2023 survey," Institution of Engineering; Technology (IET), Stevenage, UK, 2023.
6. "CIB." 2025.
7. M. M. Ibrahim Mahmoud, "Toward a Robust AI for BIM: Dataset Challenges and Implications for the AEC Industry," in ASCAAD 2025 conference proceedings, 2025. Available: https://papers.cumincad.org/data/works/att/ascaad2025_046.pdf
8. "Adopting Large Language Models in the Construction Industry: Drivers, Barriers, and Strategic Implications from China," MDPI, 2023, doi: 10.3390/buildings15234296.
9. Hartman Executive Advisors, "Navigating the Security, Risks, and Governance of Generative AI in Construction Businesses." 2023. Available: <https://hartmanadvisors.com/navigating-the-security-risks-and-governance-of-generative-ai-in-construction-businesses/>
10. S. Rogelberg, "One AI Bubble Has Already Burst. The Next One—a 'Rare' Kind—is Still Growing," Fortune, 2026, Available: <https://fortune.com/2026/03/29/ai-bubble-already-burst-rare-one-still-growing/>
11. D. Brabham, "Enterprise AI Trends 2026: How Leaders Measure ROI and Risk." 2026. Available: <https://research.etr.ai/etr-data-drop/enterprise-ai-trends-2026-how-leaders-measure-roi-and-risk>
12. Striim, "The Challenge of Data Quality and Availability—And Why It's Holding Back AI and Analytics." 2025. Available: <https://www.striim.com/blog/data-quality-availability-ai-analytics/>
13. Databricks, "AI Governance Best Practices: How to Build Responsible and Effective AI Programs." 2025. Available: <https://www.databricks.com/blog/ai-governance-best-practices-how-build-responsible-and-effective-ai-programs>
14. "Ecodesign for Sustainable Products Regulation (ESPR, Regulation (EU) 2024/1781)." 2024.
15. "Rakentamislaki (Building Act 751/2023)." 2023.
16. Credo AI, "Latest AI Regulations Update: What Enterprises Need to Know in 2026." 2025. Available: <https://www.credo.ai/blog/latest-ai-regulations-update-what-enterprises-need-to-know>
17. C. Avila, D. Ilbay, and D. Rivera, "Human-AI Teaming in Structural Analysis: A Model Context Protocol Approach for Explainable and Accurate Generative AI," Buildings, vol. 15, p. 3190, Sept. 2025, doi: 10.3390/buildings15173190.
18. "Communication," MDPI, 2025, doi: 10.3390/buildings16051018.
19. Amazon Web Services, "What is MLOps? - Machine Learning Operations Explained." 2025. Available: <https://aws.amazon.com/what-is/mlops/>
20. AppRecode, "MLOps Workflow: From Definition to Production-Ready Pipelines." 2025. Available: <https://dev.to/apprecode/mlops-workflow-from-definition-to-production-ready-pipelines-1lbt>
21. METR, "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity." <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>, July 2025.
22. "The Impact of Generative AI on Architectural Conceptual Design: Performance, Creative Self-Efficacy and Cognitive Load," 2025.
23. J. Fruehan and M. Doyle, "See How AI Can Transform MEP Design Decisions," CSE Magazine, 2026, Available: <https://www.csemag.com/see-how-ai-can-transform-mep-design-decisions/>
24. P. et al., "Artificial Intelligence for the Built World*," 2025, doi: 10.1016/j.aei.2025.104225.
25. "Hybrid Machine Learning for Smart Building Fault Detection: Integrating Logistic Regression, Random Forest, SVM, and LSTM Networks in Management, Occupancy, and Space Utilization," 2023.
26. E. Pestana, A. Paice, and S. West, "Optimizing MEP Design in Early AEC Projects Through Generative Design," Automation in Construction, vol. 165, 2024, doi: 10.1016/j.autcon.2024.105566.
27. M. Wiggins, "Generative AI Governance Framework: Building Enterprise Oversight." 2026. Available: <https://agility-at-scale.com/ai/generative/genai-governance-and-oversight/>
28. N. Gaudio, "AI Governance Auditing: Best Practices Every Enterprise Needs in 2026." 2026. Available: <https://www.sweep.io/blog/ai-governance-auditing-best-practices-every-enterprise-needs-in-2026>
29. K. Waehner, "Enterprise Agentic AI Landscape 2026: Trust, Flexibility, and Vendor Lock-in." 2026. Available: <https://www.kai-waehner.de/blog/2026/04/06/enterprise-agentic-ai-landscape-2026-trust-flexibility-and-vendor-lock-in/>
30. R. Flanigan, "Nearly 3 in 4 Enterprises Say Losing AI Vendors Would Disrupt Core Business Operations." 2026. Available: <https://zapier.com/blog/ai-vendor-lock-in-survey/>
31. Swfte AI, "Breaking Free: How Enterprises Are Escaping AI Vendor Lock-in in 2026." 2026. Available: <https://www.swfte.com/blog/avoid-ai-vendor-lock-in-enterprise-guide>
32. "From Fear to Wisdom: A Maturity Model for Humane, Values-Driven AI," White paper.
33. A. Buttell, "Who is Liable When AI Goes Wrong?" Communications of the ACM, 2023.
34. ISO/Verisk, "CG 40 47 and CG 40 48: Generative AI Exclusion Endorsements." 2026.

First Q Network

- First Q Network is an alliance of leading Mechanical & Electrical (M&E) engineering consultancies, with members from 15 countries across Europe. United by a shared commitment to excellence, innovation, and sustainable design, the network enables its members to operate seamlessly across borders while maintaining strong local expertise.
- Through structured collaboration, working groups, and knowledge sharing, First Q members combine expertise and resources to address complex built environment challenges, ensure local compliance in international projects, and deliver high performance engineering solutions across a wide range of sectors.
- With a combined team of over 4,400 professionals and a strong focus on emerging areas such as artificial intelligence, smart buildings, low carbon design, and digital engineering, First Q Network transforms complexity into smarter, more sustainable outcomes, offering clients a trusted platform for delivering future ready projects across Europe.



Members