CMCS INSIGHTS · AI

AI for PMIS

Sixteen articles on what AI can and cannot do with project data: the three layers of AI on a PMIS, the governance questions to answer first, what arrives with PMWEB 2026, where partners such as Gryps add reasoning and capture, and how to start without a migration.

By Bassam Samman, PMP, PSP, EVP, GPM

CEO & Chairman, CMCS

IN THIS SERIES

16 articles

01	No Data, No AI: Why the Substrate Decides the Answer
02	Three Layers of AI on a PMIS
03	The Approval Gate Is the Product
04	Seven Questions to Answer Before AI Touches Your Project Data
05	Find, Open, Create, Help: What Arrives with PMWEB 2026
06	From Scanned Contract to Structured Record, with the Source Shown
07	Describe the Form, Get the Form
08	What a Project Summary Has to Do Before You Can Act on It
09	Early Warning, and the Gate That Has to Stay Closed
10	Where a PMIS Needs Help: Capture, OCR and the Archive Nobody Has Read
11	Reasoning Across Platforms: When the Answer Lives in Three Systems
12	Twelve Questions to Ask Your PMIS Vendor About AI
13	The Use Cases Your Vendor Will Never Ship
14	Four Organisations, Four Sets of Questions
15	You Do Not Need a Migration to Start
16	AI-Assisted Migration of Legacy Project Data, and What It Will Not Do for You

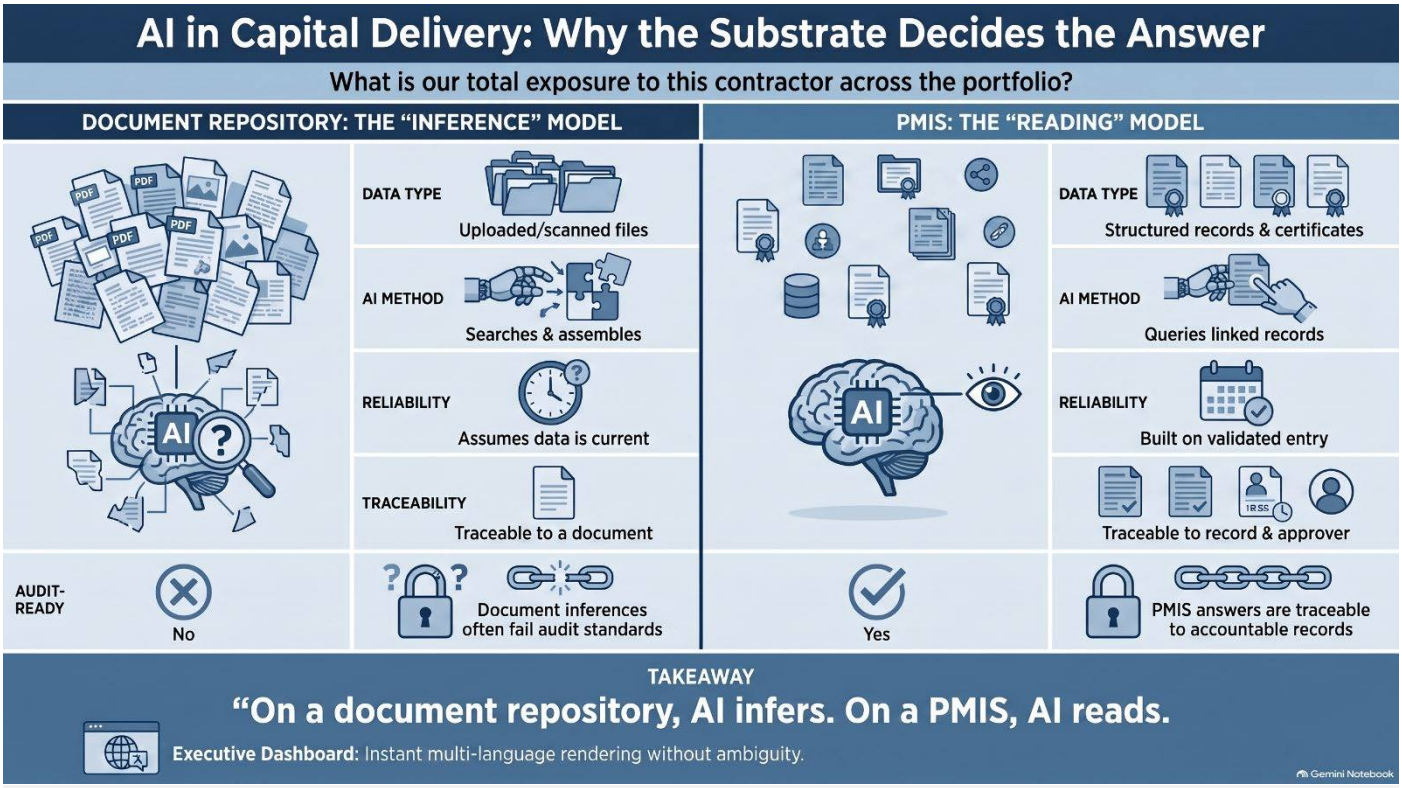
ARTICLE 1 OF 16 — AI FOR PMIS SERIES

No data, no AI: why the substrate decides the answer

Everyone has now seen the demonstration. Someone asks an AI assistant a question about a project, and back comes a fluent, confident paragraph. The room nods. Then somebody asks the question that actually matters, which is where that answer came from, and the honest reply is that it came from a document uploaded eleven months ago by a person who has since left the company.

That is the failure mode of AI in capital delivery, and it is not a failure of the AI. The model did its job. It read what it was given. The problem is what it was given.

Before any organisation evaluates an AI capability for its projects, it needs to answer a prior question that almost nobody asks: what is the AI reading? The substrate underneath an AI system sets a ceiling on the quality of every answer it can produce, and no amount of model sophistication lifts that ceiling.



Before evaluating any AI capability for your projects, answer a prior question that almost nobody asks: what is the AI reading?

The capabilities in circulation are document capabilities

Look at the AI feature lists moving around our industry at the moment. Assisted drafting. OCR intake. Tone matching. Smart search. Policy referencing. Sensitive content detection. Document quality checks. A tracking assistant. Automated follow-ups.

Every one of those is a genuinely useful capability. And every one of them was designed for a world in which documents arrive, get uploaded or scanned, and then need to be understood. They are document management capabilities. Several are excellent document management capabilities. They are being sold, increasingly, with project management labels attached.

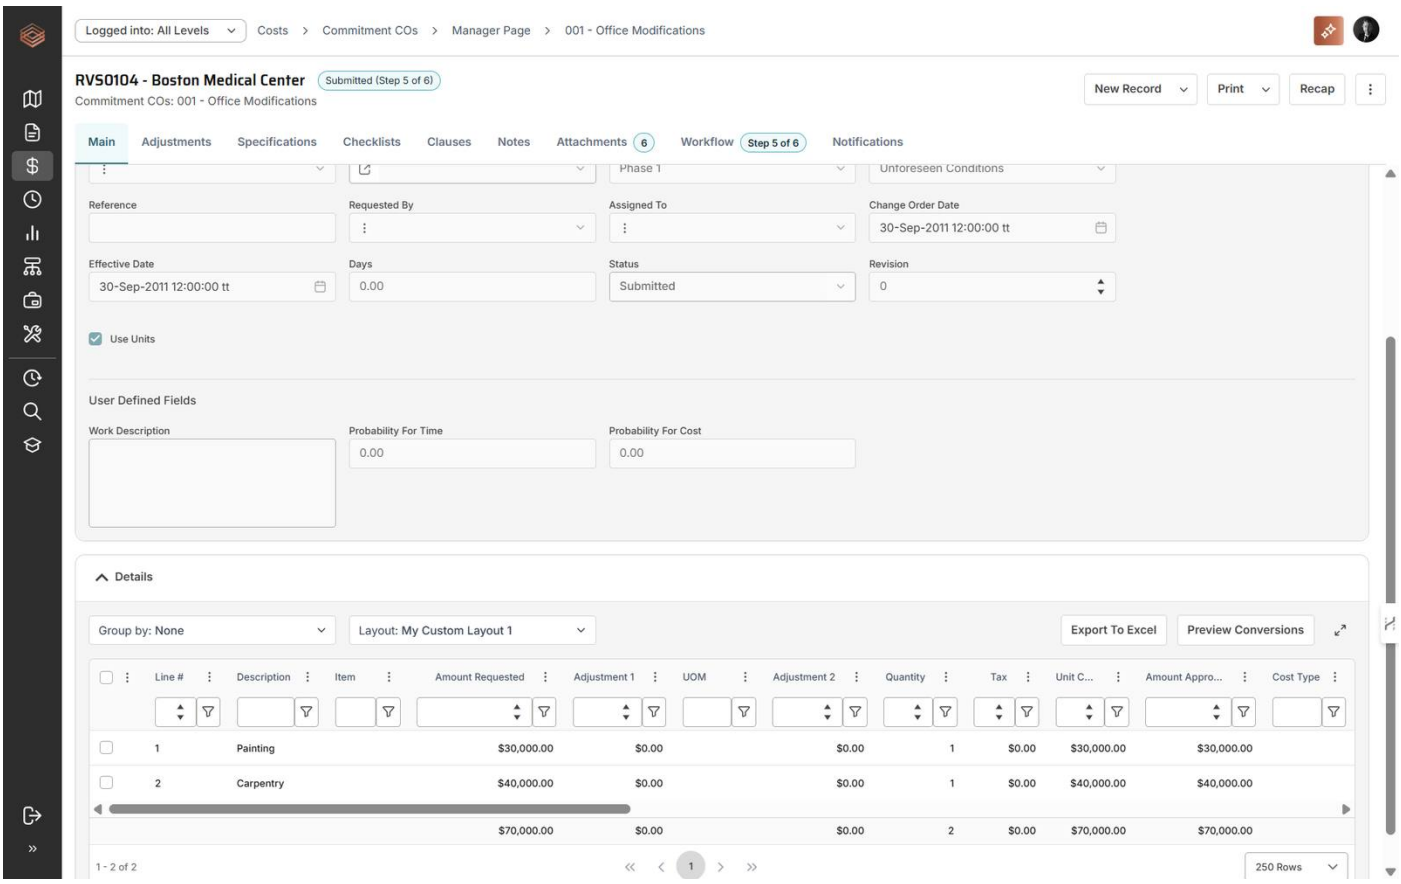
The distinction is not pedantry. It determines what the AI is able to know.

What a PMIS does differently

A project management information system does not wait for a document. It captures information at the point it is created, through a process template designed for that specific transaction.

When a change order is raised in a PMIS, it is not a letter that someone will later scan. It is a record with a defined structure. The cost codes come from a controlled list. The dates are validated. The originating commitment is a relationship, not a sentence referring to one. The approval routing is defined before the record exists, so the question of who authorised it is answered by the system rather than reconstructed from an email chain afterwards.

Multiply that by every transaction on a programme, over several years, and the difference is not incremental. One organisation has a repository of documents about its projects. The other has a structured, validated, permissioned model of what actually happened.



On a document repository, AI infers. On a PMIS, AI reads.

Inference and reading are not the same answer

Ask both systems the same question. What is our total exposure to this contractor across the portfolio?

The document system searches its content, finds the contracts and correspondence that mention the contractor, extracts the numbers it recognises, and assembles a figure. That figure is an inference. It may well be right. But it rests on the assumption that every relevant document was uploaded, that none of them is superseded, and that the numbers in them were never varied by an instruction that took a different form.

The PMIS queries commitments, change orders, invoices and payment certificates that are already linked to that vendor as structured records. The figure it returns is not an inference. It is a reading of the system of record, and every component of it can be opened and inspected.

In most industries the first answer is good enough. In capital delivery it is not, because the answer has to survive contact with a quantity surveyor, an auditor, or eventually a tribunal. An answer that cannot be traced back to its source is an answer that cannot be defended, and an answer that cannot be defended has no operational value no matter how well written it is.

A consequence for executive reporting

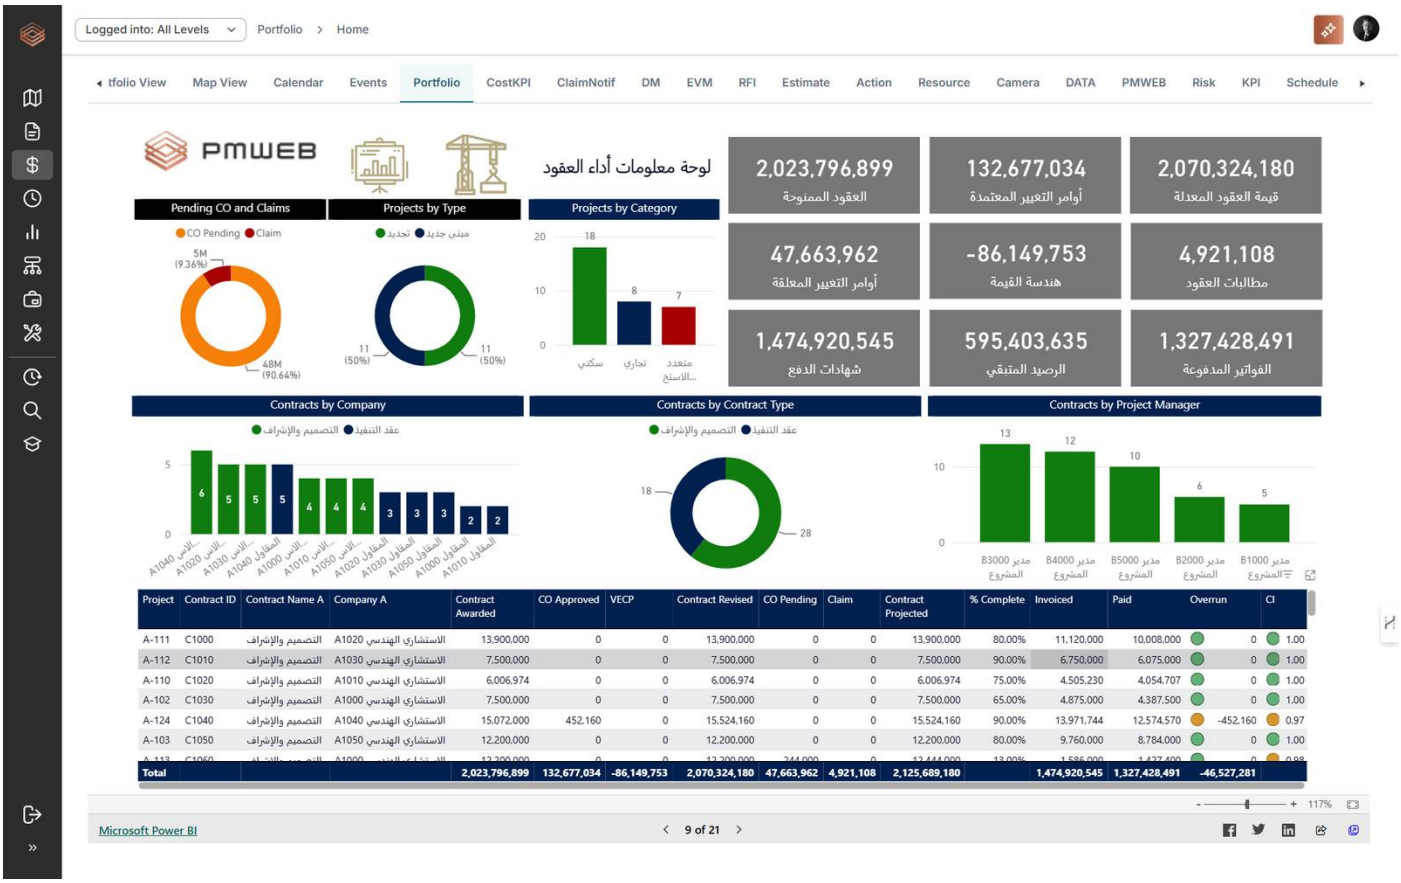
There is a second effect of structured capture that matters particularly in this region, and it is rarely connected to the AI conversation at all.

Most project communications in the Gulf are captured in English. RFIs, change orders, correspondence and meeting minutes are written in the language the project teams work in. Boards, ministries and senior government stakeholders, however, often expect to review performance in Arabic.

A PMIS bridges that gap without translating a single record. Structured data carries controlled vocabulary: statuses, categories, cost codes, phases, disciplines, value lists. These are not free text. They are defined terms with a fixed set of permitted values, so each one has to be labelled in Arabic only once. From then on, the records the project teams captured in English can be presented in Arabic dashboards and reports for executive review, as the contracts performance dashboard below shows. Every figure is read directly from PMWEB; only the titles, measures and categories change language.

A document repository has no equivalent, because everything inside it is narrative. To present it in Arabic you have to translate it, and translating contractual language introduces exactly the ambiguity the contract was written to remove.

The detail beneath the dashboard stays in the language it was written in. The summary layer, the layer senior stakeholders actually read, follows the structure and can be generated in whichever language the audience needs. That is a property of the substrate, not a feature anyone had to build.



Contracts performance dashboard in PMWEB, generated in Arabic for executive review from records captured in English. Illustrative data.

The question to take into the room

Over the coming months every platform serving this industry will announce AI capability, and the announcements will be difficult to tell apart. They will use the same adjectives. Several will use the same screenshot of a chat panel.

The question that separates them is not what the AI can do. It is what the AI is reading, and whether the answer it gives you can be traced back to a record that someone was accountable for creating.

Ask a vendor to show you the answer. Then ask them to show you where it came from. The second demonstration is the one worth watching.

This is no longer only the buyer's position. In September 2026 PMWEB published its own checklist for owners and construction managers, *Six Questions to Ask Before You Put AI Near Your Project Data*, and two of its six questions are the argument of this article in question form. The first asks whether the AI can show the source record behind every number: whether a cost forecast, a schedule risk score or a contract summary links through to the record it came from, so that a reviewer can open the pay application or the schedule revision behind it and check it. The sixth asks whether it works on your actual records, meaning live commitments, change orders, RFIs, submittals, cost and schedule, rather than a curated set of documents uploaded to a chat window for the demonstration.

Together those two questions are the substrate argument turned into a test. A demonstration built on a curated document set tells you how the AI performs on the documents the vendor chose. A demonstration run against a live programme tells you how it performs on yours.

> **The question that separates platforms is not what the AI can do. It is whether the answer it gives you can be traced back to a record that someone was accountable for creating.**

This is the first of sixteen articles on AI in project management information systems, published weekly. Next week: why AI value on a PMIS arrives in three distinct layers, and why most organisations evaluate only the first.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZCO, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on

Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

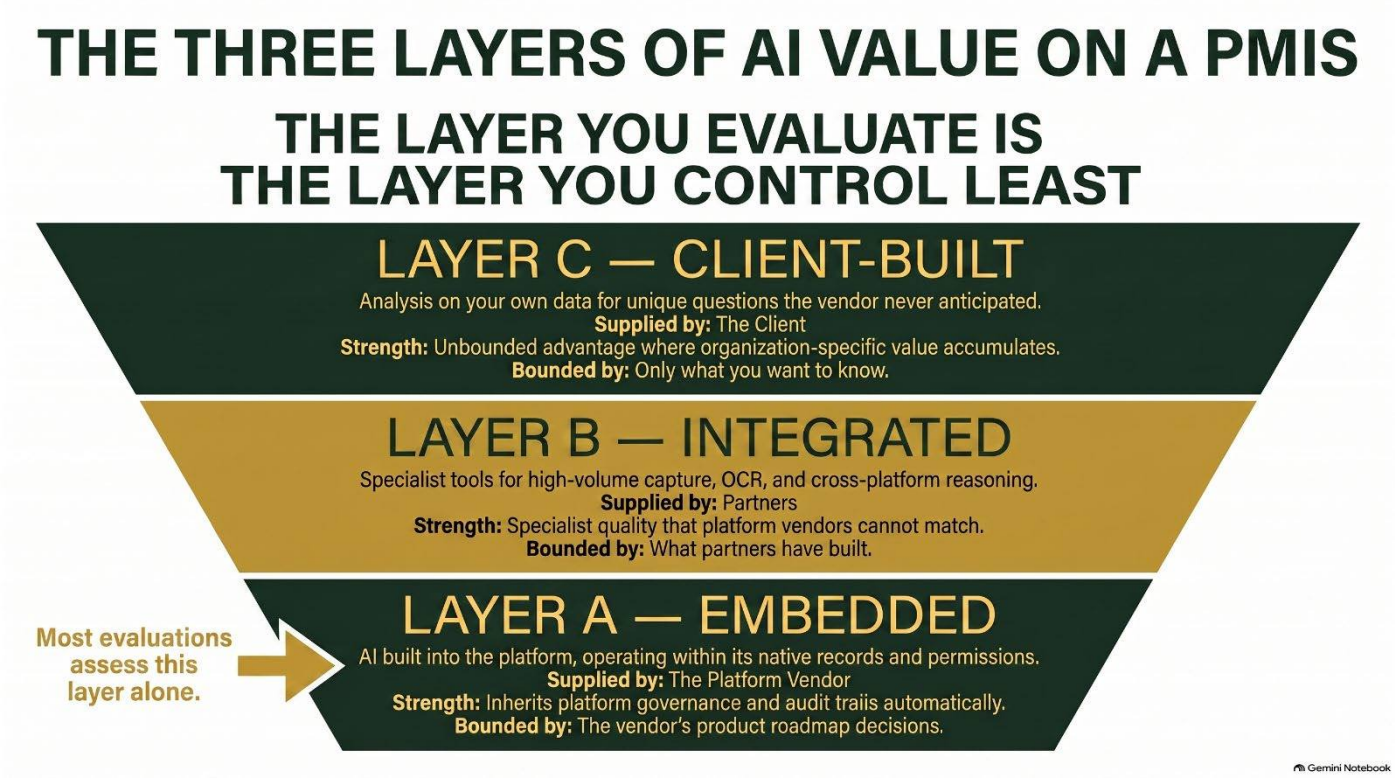
ARTICLE 2 OF 16 — AI FOR PMIS SERIES

Three layers of AI on a PMIS: what the vendor ships, what partners add, what you build yourself

Most organisations evaluating AI for their project systems are looking at one third of the picture, and it is the third they control least.

The evaluation usually runs like this. The vendor presents a roadmap. The team assesses the features against a requirements list. Someone asks about timing. A score is assigned. The exercise is entirely reasonable and it is also incomplete, because the vendor's feature list is only one of three places AI value on a PMIS comes from.

Understanding all three changes the question you are asking. It moves from what AI does your product have to what AI can I run on my project data, and organisations asking the second question end up in a materially better position than those asking the first.



Most PMIS AI evaluation stops at the vendor's feature list, which is the smallest of the three sources of value and the one the client controls least.

Layer A — the AI the platform ships

This is the layer everybody evaluates. Capability built into the PMIS itself, operating on its records, inside its permission model, recorded in its audit trail.

Its great advantage is that it inherits the platform's governance for free. When an assistant built into the system searches records, it searches only what the user was already entitled to see, because it is using the same access control that governs every other part of the application. Nobody had to build that separately and nobody can accidentally bypass it. For a capital programme that inheritance is worth more than most of the features sitting on top of it.

PMWEB states the same principle in its Intelligence Control Framework, the governance architecture for AI in the platform: AI should work under the permission model the owner already administers, not a second one alongside

it. A second permission model is a second thing to configure, a second thing to audit, and a second place for the two to disagree.

Its limitation is that it is defined by the vendor's judgement about what most customers need. Your organisation is not most customers.

Layer B — specialist capability joined to the platform

No PMIS vendor will ever be the best in the world at high-volume document capture. No document capture vendor will ever be the best at commitment management. These are different disciplines with different engineering problems, and the honest position is that a serious AI stack draws on both.

This layer covers two things in particular. At the front, capture and OCR, turning the paper and image estate into something structured. Above the platforms, reasoning that spans systems, because the questions that matter most in capital delivery rarely live inside one application. Which of my delayed activities have unapproved commitment changes behind them is a question that touches the schedule tool and the PMIS at once, and neither can answer it alone.

The trade at this layer is integration effort against specialist quality. Buying a bundled version of a specialist capability from a platform vendor usually means buying the weakest available version of it.

The trade has a governance side as well. A partner capability does not inherit the platform's controls automatically, so the questions a buyer puts to the platform vendor (whose permissions it honours, whether it reads your data or trains on it, what it logs) have to be put to the partner separately. Part III of this series returns to this.

Layer C — what you build yourself

This is the layer almost nobody evaluates, and it is where organisation-specific advantage actually accumulates.

Every PMIS client recognises the following experience. The question you genuinely need answered is not on the report menu. You raise an enhancement request. It is assessed, prioritised, scheduled, and arrives two releases later in a form that is not quite what you asked for. By then the question has changed.

Layer C is what happens when an organisation can point a general AI tool at its own project data and ask that question directly. Not through the vendor's roadmap. Not as an enhancement request. As an analysis, expressed in plain language, produced in an afternoon.

This layer is unbounded in a way the other two are not. Layer A is bounded by the vendor's product decisions. Layer B is bounded by what partners have built. Layer C is bounded only by what your organisation wants to know, and by whether you can reach your own data well enough to ask.

> **Layer A is bounded by the vendor. Layer B is bounded by the partner. Layer C is bounded only by what you want to know.**

The test that separates platforms

That last condition is the whole point, and it is where platforms differ far more than their feature lists suggest.

A platform that exposes your data only through a rate-limited interface has made a decision about who is allowed to analyse your project history, and the answer is the vendor. A platform that holds your data in a dedicated database with a documented schema has made the opposite decision. The feature comparison will not show you this. It is an architectural fact, and it determines whether Layer C exists for you at all.

A later article in this series takes that point apart, because it also governs what happens to fifteen years of history in your current system if you ever decide to move. For now it is enough to note that the question belongs on the evaluation list, and that it is almost never on it.

Why the model matters for evaluation

Assessing a platform on Layer A alone produces a predictable outcome. Every serious vendor is shipping broadly similar capability this year, the scores cluster, and the decision falls back to price or incumbency.

Assessing all three layers produces a different conversation. Layer A tells you what you get on day one. Layer B tells you whether the platform is a good citizen in a stack you will inevitably assemble. Layer C tells you who controls your project history, which is a question about the next ten years rather than the next release.

Layer A tells you what you get on day one. Layer C tells you who controls your project history, which is a question about the next ten years rather than the next release.

This series runs weekly in four parts. Part I establishes how AI on a PMIS should be judged. Part II covers what arrives with the PMWEB 2026 release. Part III covers the partner layer. Part IV covers the use cases organisations build for themselves, and what it takes to start. Next week: why every AI capability worth deploying in capital delivery stops and asks.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

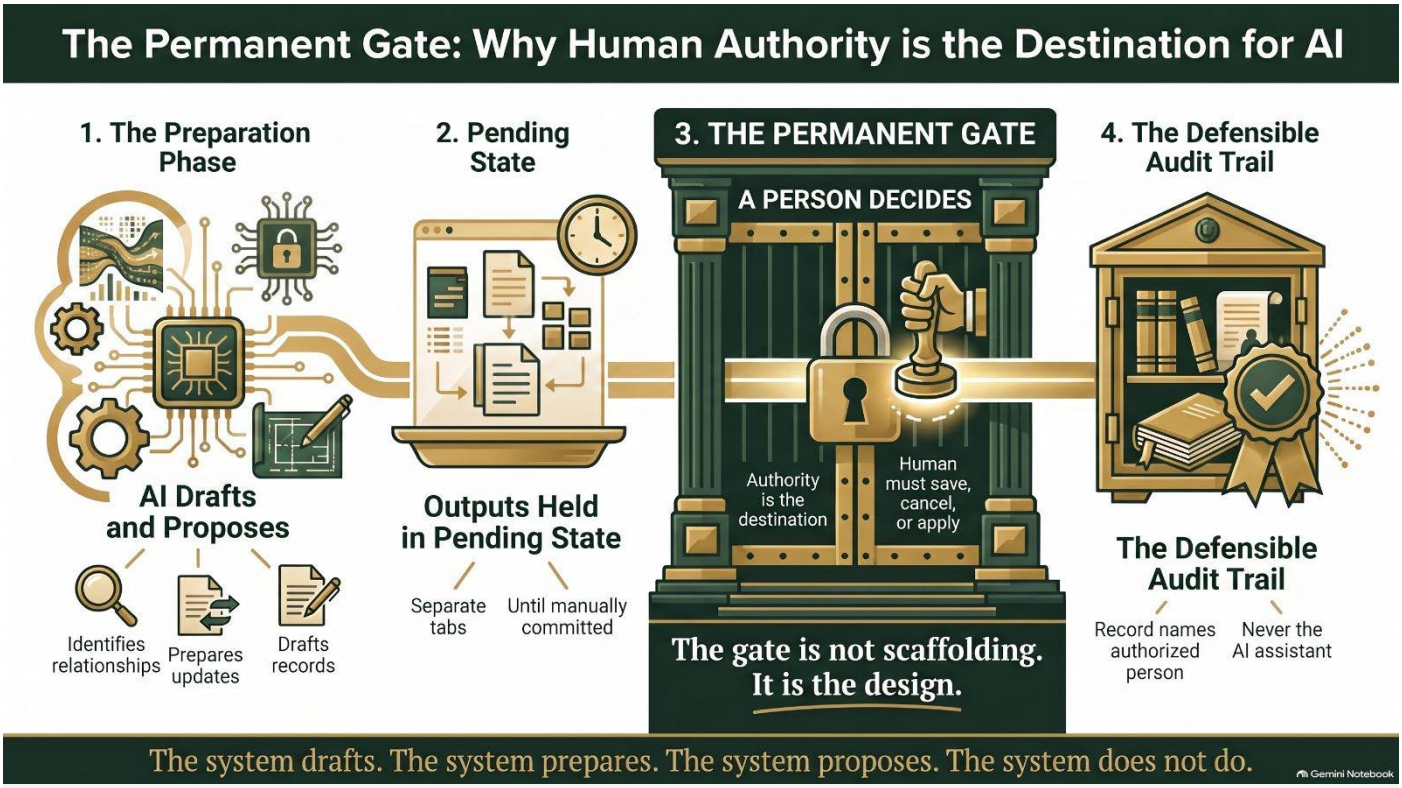
CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

ARTICLE 3 OF 16 — AI FOR PMIS SERIES

The approval gate is the product

There is an assumption running underneath most AI discussion in our industry, and it is worth stating out loud so it can be examined. The assumption is that autonomy is the destination. That today's AI asks permission because it is early, that tomorrow's will not need to, and that the approval step is scaffolding which comes down once the building is safe.

In capital delivery that assumption is wrong, and organisations that build their AI strategy on it will spend the next three years constructing something they cannot use.

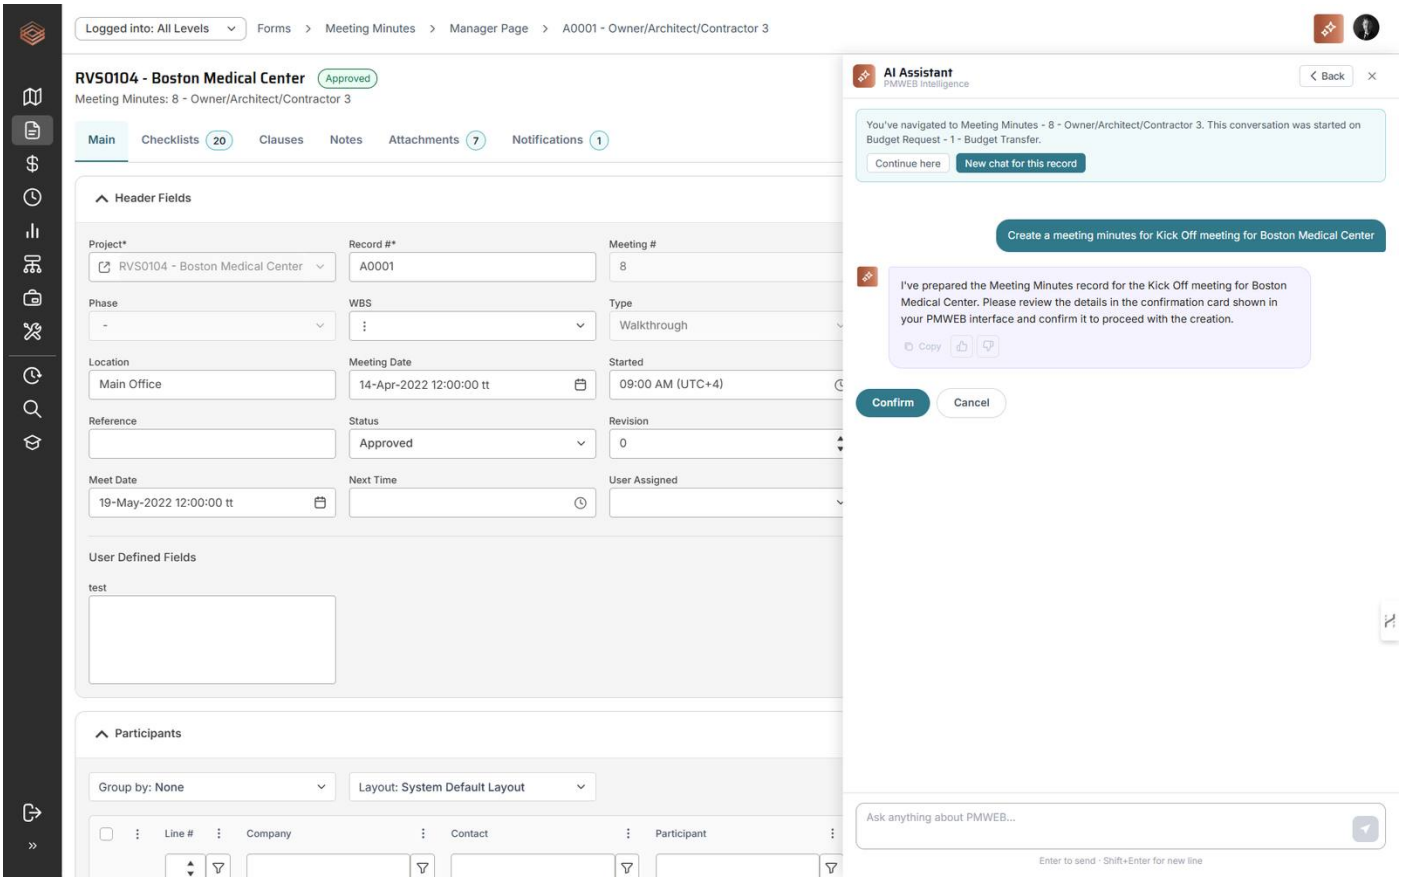


The assumption that autonomy is the destination, and that the approval step is scaffolding to be removed once the building is safe, is wrong in capital delivery.

What the demonstrations actually show

Watch closely what happens at the end of an AI action in a well-designed project system.

The assistant drafts a record. It opens in a separate tab, so whatever the user was working on is untouched and nothing in progress is overwritten. The user reviews the fields. Then the user saves, or the user cancels. The system does not save in the background.



Where the AI has prepared updates to an existing record from a source document, those updates stay pending. They are visible, they are ready, and they are not applied until someone commits them. Where the AI identifies a relationship between two records that were not previously connected, it will surface that relationship for analysis and it will not create the link. Creating the link stays a deliberate act by a person.

Read as a limitation, this is a weakness that later releases will remove. Read correctly, it is the only design that survives contact with how capital programmes are actually governed.

Why a programme cannot absorb an unattended actor

A capital programme runs on a delegation of authority. Every commitment, every variation, every certification, every payment release has a named individual behind it who was authorised to make that decision at that value. The DOA is not administrative decoration. It is the mechanism by which an organisation can say, eighteen months later, who decided this and on what basis.

Introduce an actor into that structure which takes decisions and is not a person, and you have not accelerated the programme. You have put a hole in it.

The consequences do not appear immediately, which is what makes this dangerous. They appear at the first serious dispute, when a party is asked to demonstrate the basis on which a variation was accepted and the audit trail records that it was accepted automatically. They appear at the first audit, when a payment release cannot be attributed to a person authorised to release it. By then the capability has been embedded for two years, and the finding is not about the AI. It is about the organisation's control environment.

The question is not whether the AI is autonomous. It is whether the audit trail still names a person.

The objection, and the answer

The reasonable objection runs like this. If a human still has to approve everything, what exactly was gained? You have replaced typing with reviewing. That is an improvement, but it is not a transformation, and it does not justify the investment.

The objection misidentifies what was scarce.

Approval capacity was never the bottleneck on a capital programme. Senior people can approve a great many things in a day. What they cannot do is work out which of four thousand open records deserves their attention this week. The scarce resource is attention, not authority, and directing attention is precisely what this class of capability does well.

An AI that drafts a change order has saved somebody twenty minutes. An AI that identifies which three of your two hundred open commitments have moved into a position warranting a decision has done something considerably more valuable, and it has done it without taking a single decision itself.

What to look for

When a vendor demonstrates AI capability, the interesting moment is not the output. It is what happens immediately after the output.

Does the action complete, or does it stop? If it stops, what exactly is the user being asked to confirm, and do they have enough in front of them to confirm it responsibly? If the AI has drafted something, can the user see what it was drafted from? If the AI has prepared a change, is it obvious that the change has not yet been applied?

And the question that matters most for anyone who will have to defend the system later: after the user confirms, what does the audit trail say? It should name the person. If it names the assistant, or names nobody, you are looking at a capability that will not survive your first serious dispute.

PMWEB's own buyer checklist for AI, published in September 2026, reduces this to the shortest question on any list of its kind: who signs? AI output should enter the same approval workflow as any other submission, with a named human approver. An AI-drafted pay application review or change order recommendation is actioned only once a named project manager has approved it. A machine-generated forecast, risk flag or summary that takes effect without a person approving it fails the test, however good the output.

The commitment worth understanding

As AI capability in project systems becomes more ambitious, and it will, the property to watch is not how much more it can do. It is whether the gate moves.

Capability that monitors a portfolio continuously, connects an external event to internal exposure and proposes a response is considerably more useful than capability that drafts a form. It is also considerably more consequential when it is wrong. The design commitment that makes the first acceptable is the same one that makes the second acceptable, and it is that the system drafts, prepares and proposes, and a person decides.

That is not a stage on the way to something else. It is the destination.

PMWEB has since written the commitment down. Its Intelligence Control Framework, the governance architecture that sets what AI can access in the platform, how it stays bound by existing permissions and where human judgement stays in the loop, rests on a guiding principle of five words: AI advises, the operator decides. That is the right sentence to hold any vendor to, because it describes where the gate sits rather than how capable the system is.

The system drafts, prepares and proposes, and a person decides. That is not a stage on the way to something else. It is the destination.

Next week: the seven governance questions that stop AI deployments after signature, and what a defensible answer to each of them sounds like.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

ARTICLE 4 OF 16 — AI FOR PMIS SERIES

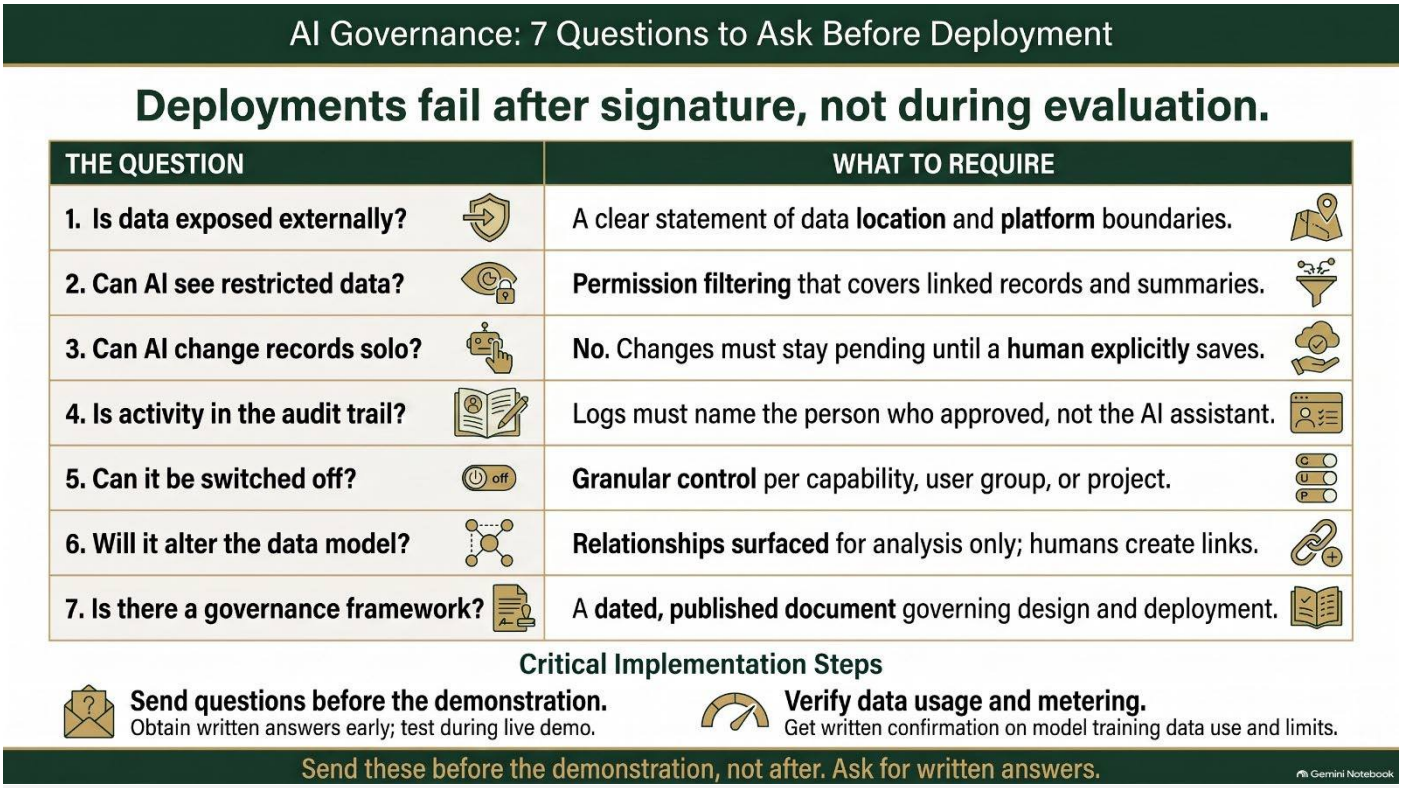
Seven questions to answer before AI touches your project data

AI deployments in capital delivery organisations do not usually fail during evaluation. They fail after signature, in the gap between the project team who wanted the capability and the audit committee, IT governance function or legal team who were asked to approve it late and had no good answers to work with.

The seven questions below are the ones that arrive at that point. They come up in every serious evaluation, they are usually answered in passing during the demonstration, and they are the questions that stop deployments.

They are worth taking into a vendor meeting deliberately, in writing, early enough that the answers can be tested. What follows is each question, what a good answer sounds like as distinct from a reassuring one, and the position PMWEB has stated on each.

PMWEB has since published its own version of this exercise, a six-question checklist for owners and construction managers titled *Six Questions to Ask Before You Put AI Near Your Project Data*. Where its questions sharpen what to require, that is noted below. The two questions it asks that are not on this list are covered at the end.



AI deployments in capital delivery organisations do not usually fail during evaluation. They fail after signature.

1. Is our data exposed to external AI services?

A reassuring answer says the platform takes security seriously. A good answer tells you where the data goes.

This is the first question every IT function asks and it usually determines whether the conversation continues. What you are establishing is whether project data leaves the platform's boundary to be processed, and if so, where it goes and under what terms.

PMWEB's stated position is that customer data remains contained within PMWEB and is not exposed to outside sources.

PMWEB's checklist takes the question one step further: is the AI reading your data, or training on it? Ask where the data goes, whether it leaves your tenant, how long it is retained, and whether it is used to train models shared with other organisations. Then ask whether those answers are in the contract. An answer that appears only in the sales deck is a red flag by the vendor's own definition. What you want is written terms confirming, for example, that change order and pay application data will not be used to train models serving anyone else.

2. Can the AI see data a user cannot?

This is the question most often answered too quickly, and it is harder than it sounds.

A search that spans records is straightforward to permission. A search that spans records and workflow comments, and then assembles a combined summary connecting activity across several linked records, is not. The risk is that the synthesis reveals something none of the individual sources would have revealed to that user. An organisation with a properly segregated programme structure, where commercial teams on different projects are deliberately walled from one another, should test this specifically rather than accept a general assurance.

PMWEB's stated position is that AI access follows existing platform permissions, and that results are permission-filtered including where they connect activity across linked records.

The checklist frames the requirement precisely: AI should work under the permission model the owner already administers, not a second one. Restricted claim positions, confidential vendor pricing and contract values visible only to some staff must not surface in an AI summary prepared for someone who cannot open the underlying records.

3. Can the AI change records without a human?

The answer you want is no, and you want to see it rather than hear it.

Ask for a demonstration that ends in the AI preparing something, and watch what happens next. Prepared updates should remain visibly pending. A drafted record should require an explicit save. The distinction between the system having done something and the system having proposed something should be obvious on screen to a user who is not paying full attention, because that is the condition under which most work is actually done.

PMWEB's stated position is that the assistant will not finalise an action such as saving a record without explicit user approval, and that prepared updates remain pending until saved or cancelled.

The checklist names the person at the end of the chain. AI output should go through the same approval workflow as any other submission, with a named human approver before anything is actioned. Its version of this question is the shortest on either list: who signs?

4. Is AI activity in the audit trail?

A capability that acts on your records and is not recorded in the audit trail has created a category of event your control environment cannot see.

What you are looking for is that AI-assisted actions are audited to the same standard as any other action, and specifically that the trail names the person who approved rather than the assistant that prepared. This matters most in the situation nobody plans for, which is a dispute two years later where the basis of a decision has to be reconstructed.

PMWEB's stated position is that permission controls and a full audit trail are retained as automation increases.

The checklist is specific about what the log should hold: the user, the timestamp, the query, the response and the resulting action. It also gives the test case. If someone asks later why contingency was drawn down in March, the record should show what was asked of the AI, what it returned, who reviewed it and what they decided. Ask

whether the query and the response are logged, not only the resulting change to the record, because the change alone does not show what the approver was told. "The system suggested it" is not an answer that survives a dispute.

> **A capability that acts on your records and is not recorded in the audit trail has created a category of event that your control environment cannot see.**

5. Can it be switched off?

This is the underrated question on the list.

An all-or-nothing capability forces an organisation into a single decision it is not ready to make. Capability that can be enabled individually allows a phased rollout, which is the only kind that survives contact with a large programme. It also allows a restricted-scope deployment, where a capability is available to a commercial team and not to a site team, or enabled on one programme while another is mid-audit.

Ask what the unit of control is. Per capability, per user group, per project, per environment. The answer determines how much of your rollout you actually control.

PMWEB's stated position is that each AI capability can be enabled or disabled individually.

6. Will it alter the data model?

An AI that creates relationships between records to improve its own analysis is altering the structure of your project data, and it is doing so on the basis of inference.

The distinction to establish is between surfacing a relationship for analysis and creating one in the system. The first is analysis. The second is a change to the record, and it should require a person.

PMWEB's stated position is that relationships can be surfaced for analysis, but creating a link requires a deliberate user action.

7. Is there a published governance framework?

The difference between a vendor with a governance framework and one without is not the presence of a document. It is whether the answers to the previous six questions are design commitments or current behaviour.

Current behaviour changes without notice. A design commitment is a position the vendor has published and can be held to. Ask for the framework, ask when it was published, and ask what in it constrains the product roadmap.

PMWEB's published AI governance architecture is the Intelligence Control Framework, which defines how AI capabilities are designed, deployed and governed within the platform.

In its September 2026 checklist PMWEB describes the framework as covering three things: what AI can access, how AI stays bound by existing permissions, and where human judgement stays in the loop. It says the framework was built to answer governance questions of this kind structurally rather than case by case, and states its guiding principle in five words: AI advises, the operator decides. That is the shape of a design commitment. The test is still the one above: ask for the document, its date, and what in it constrains the roadmap.

Vijai Viswanathan, PMWEB's Chief Product and Technology Officer, put the balance plainly when the checklist was published: the potential of AI in capital project management is enormous, but so is the responsibility.

Two positions worth having in writing

Two things are worth obtaining from any vendor in written form rather than as demonstration commentary, because both will be relied upon later.

The first is what happens to your data. Not only whether it leaves the platform, but whether it is used to train or improve models, and how long AI-related data is retained. This is the question the audit committee asks, and a verbal assurance in a demonstration is not something an audit committee can accept.

The second is whether AI usage is metered. Several platforms are introducing token, query or per-seat allowances for AI capability. An allowance changes rollout fundamentally, because a capability people are told to ration does not get adopted at all. PMWEB has stated that AI usage carries no limit.

Two more questions, from the vendor's own checklist

The seven questions above are the ones that stop deployments after signature. Two of PMWEB's six are not on this list, and they are worth adding because they expose a weak capability before it gets that far.

Can it show the source record behind every number?

Every figure the AI produces, whether a cost forecast, a schedule risk score or a contract summary, should link to the record behind it, ideally in one click. A forecast that cites the pay application or schedule revision it rests on can be checked by the reviewer. A figure with no traceable source can only be believed or disbelieved, and nobody in capital delivery should have to make that choice.

Does it work on your actual records?

Ask for the demonstration to run on live programme data (commitments, change orders, RFIs, submittals, cost and schedule) rather than on a curated document set or files uploaded to a chat window. A capability that performs well on a prepared sample and has never been run against an active programme has not yet been tested on what you are buying it for.

Using the list

Send the seven questions, and the two above, before the demonstration rather than after it. Ask for written answers. Where a vendor tells you that an answer is still being worked through, that is a more useful response than a general assurance, and it should be recorded as such rather than treated as a failure.

Any vendor whose answers to these seven are solid has told you something important about how they think, and it will be visible in the rest of the product.

> **Send the seven questions before the demonstration rather than after it, and ask for written answers. A vendor giving you a complete set of positive answers is not describing a product.**

That concludes Part I of this series. Part II examines what the PMWEB 2026 release brings, beginning with the on-demand assistant.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

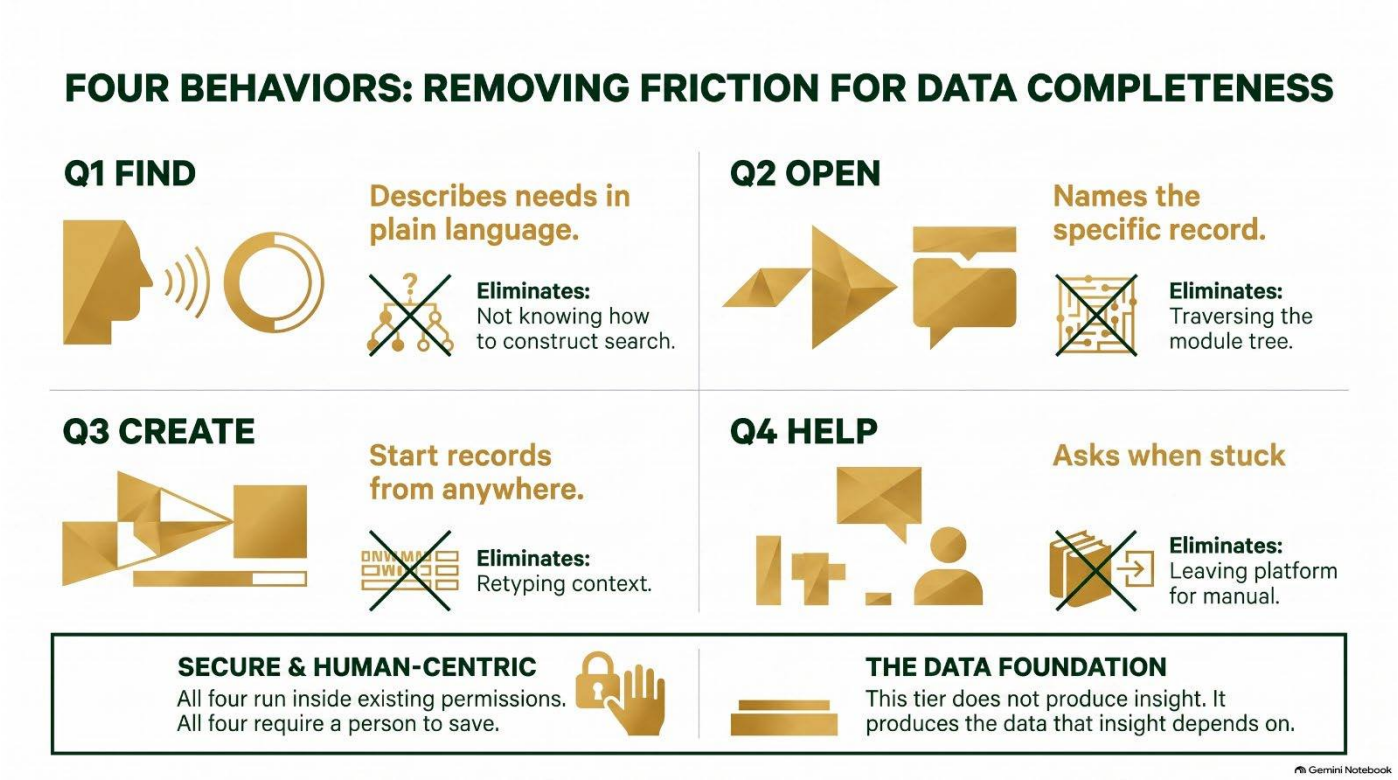
CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

ARTICLE 5 OF 16 — AI FOR PMIS SERIES

Find, open, create, help: what arrives with PMWEB 2026

Ask a project director why the data in their PMIS is thinner than it should be and you will rarely be told that people do not understand the system. You will be told that the people who have the information do not enter it, and the people who enter it do not have the information.

That gap is usually described as a training problem. It is not. It is a friction problem, and it is worth understanding precisely, because the first AI capability arriving in PMWEB 2026 addresses exactly that and nothing more ambitious.



Four behaviours, available wherever the user happens to be working: find, open, create and help.

> **Adoption is a data completeness problem before it is a training problem. Every workaround a user adopts removes data from the system of record.**

The cost of knowing where things are

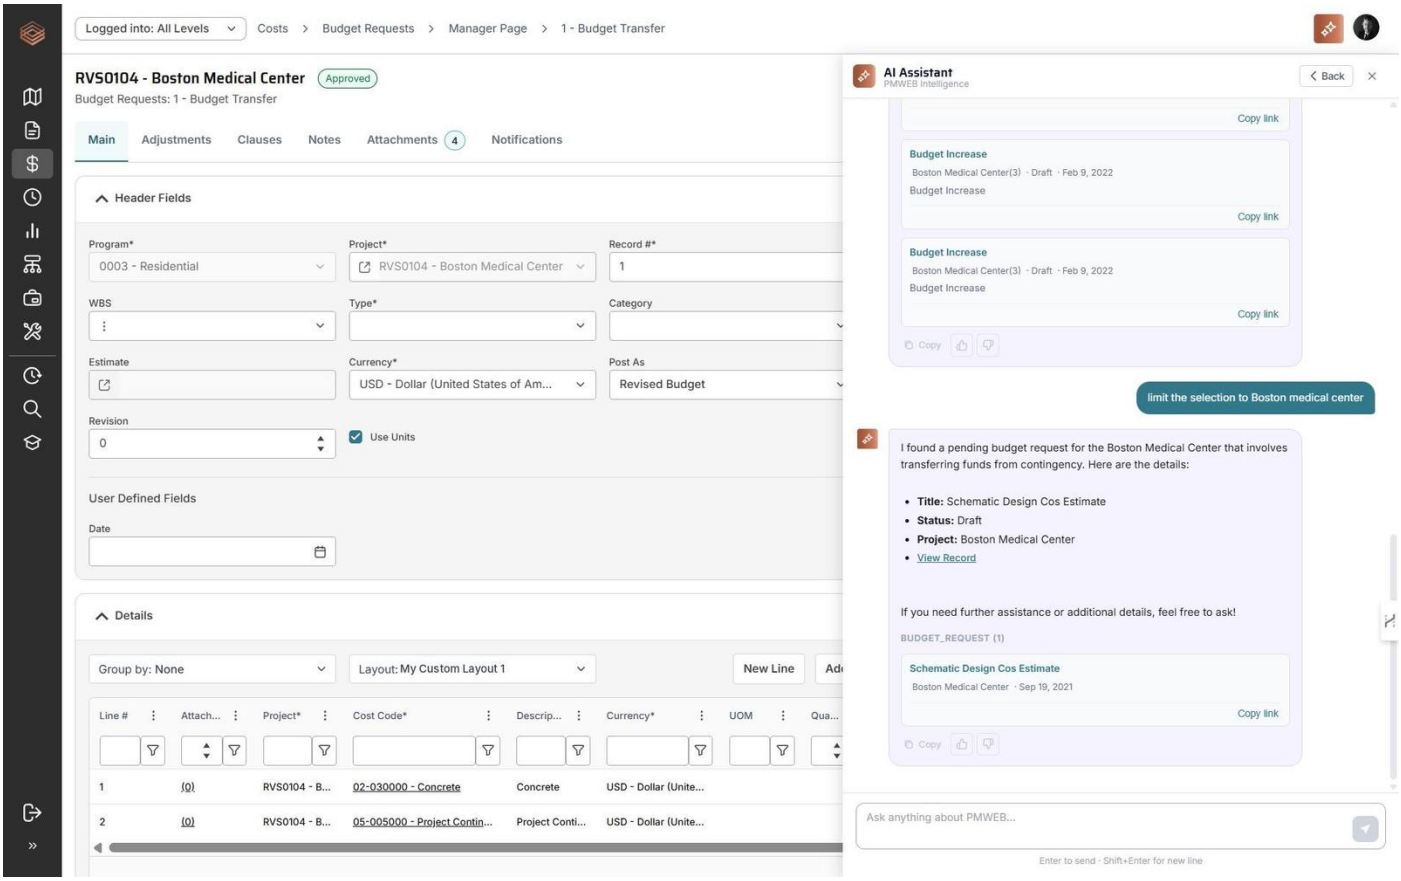
A configured PMIS is a large application. Modules, sub-modules, a record tree, filters, saved views. An experienced user navigates it without thinking, which is precisely why the cost is invisible to the people who evaluate the system.

The occasional user is the one who matters here. The site engineer who raises four records a month. The commercial manager who needs one figure before a meeting. The consultant who was given access for a specific purpose and uses it fortnightly. For each of them the navigation is not automatic, and the cost of working out where something lives is high enough, often enough, that the rational response is to do it another way. They send an email instead. They ask someone. They keep it in a spreadsheet and enter it later, or not at all.

Every one of those decisions removes data from the system of record, and the aggregate of them is the reason the PMIS does not reflect the programme as fully as it should.

Find

The user describes what they need in plain language and the assistant searches records across PMWEB, returning linked results.



The important qualifier is that it searches what the user is authorised to see. This is not a permission layer bolted on afterwards; the assistant uses the same access control that governs every other part of the platform, which means it cannot be configured into an inconsistent state and cannot be accidentally bypassed. For an organisation running segregated commercial teams across programmes, that inheritance matters more than the search itself.

It is also the second question on PMWEB's own AI checklist for buyers: does it respect the permissions already in place? The Intelligence Control Framework answers it in the terms that matter, which is that AI works under the permission model the owner already administers rather than a second one.

Open

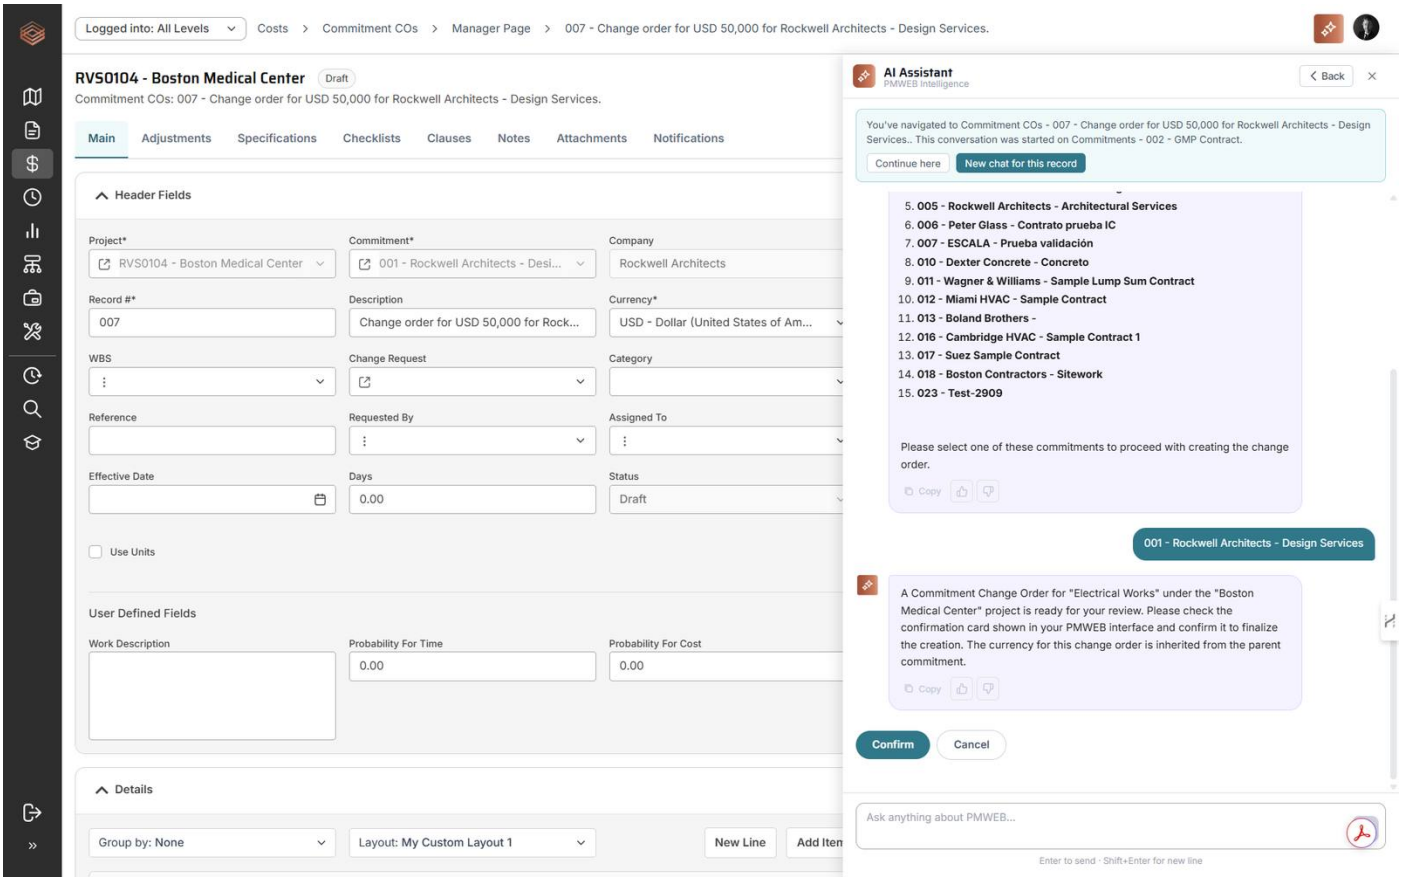
The assistant takes the user directly to the best-matching record, removing the traversal through the module tree entirely.

This sounds like the smallest of the four and it may be the one that changes behaviour most, because it removes the specific friction that stops occasional users. The user no longer has to know where a record type lives in order to reach an instance of it.

Create

An authorised record or adaptive form can be started from anywhere in the platform, without navigating to the originating module first.

The part worth examining is what happens next. The assistant reads the page the user is on and populates a draft from it. In the demonstration of a commitment change order, eight of twenty-one fields were pre-filled from context before the user typed anything.



Eight of twenty-one is worth sitting with, because it is a more honest number than a percentage would suggest. It is not most of the record. It is the portion that could be derived from where the user already was, and it is the portion that is tedious to retype and easy to get wrong. The remaining thirteen fields are the ones that require judgement, which is the correct division of labour.

The draft opens in a separate tab, so nothing in progress is disturbed, and it is not saved until the user reviews the fields and saves it. As covered in an earlier article in this series, that gate is a design commitment rather than a limitation of an early release.

Help

Contextual guidance, suggested prompts, follow-up questions and access to previous assistant sessions, delivered inside the working context.

Product help has historically lived somewhere else — a manual, a portal, a training session six months ago. Help that arrives where the user is stuck is a different proposition, and for a platform whose adoption depends on occasional users it may be worth more than the three capabilities above it.

Why this tier matters more than it appears to

None of this is insight. There is no analysis here, no early warning, no recommendation. It is a set of capabilities aimed squarely at the friction that stops people entering data and finding it.

That is exactly why it is the right thing to ship first.

> **Every capability further up depends on the data being there. Everything else in AI for project systems is downstream of that.**

An organisation that deploys sophisticated analysis on a thinly populated PMIS gets sophisticated analysis of an incomplete picture, delivered with more confidence than the underlying data deserves. Removing the friction that keeps the record incomplete is not a modest first step ahead of the interesting work. It is the precondition for the interesting work.

One commercial point worth noting

Several platforms are introducing usage allowances for AI capability — token budgets, query limits, per-seat entitlements. PMWEB's stated position is that AI usage carries no limit.

For a capability of this kind that is not a minor commercial detail. Capability aimed at occasional users only works if occasional users are allowed to use it freely. A team told to ration something does not adopt it; it reverts to the email and the spreadsheet, and the friction problem returns intact.

What to do about it

Two practical points for organisations already running PMWEB. This capability requires PMWEB 2026, so establish your current version and your upgrade path now rather than at release. And AI capability can be enabled individually, which means a phased rollout is available to you — worth planning deliberately rather than switching everything on at once.

For organisations not on PMWEB, the thing to take from this is the sequencing. When a vendor shows you AI analysis before they have shown you AI that makes the system easier to populate, ask what the analysis is going to be reading. And ask to see it run on your live records rather than a prepared sample. PMWEB's own checklist makes that the last of six questions to put to any vendor, including itself.

> **Everything else in AI for project systems is downstream of the data being there. Removing the friction that keeps the record incomplete is not a modest first step. It is the precondition for the interesting work.**

Next week: turning a signed contract into a structured record, and why showing the source matters more than the extraction.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation

and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

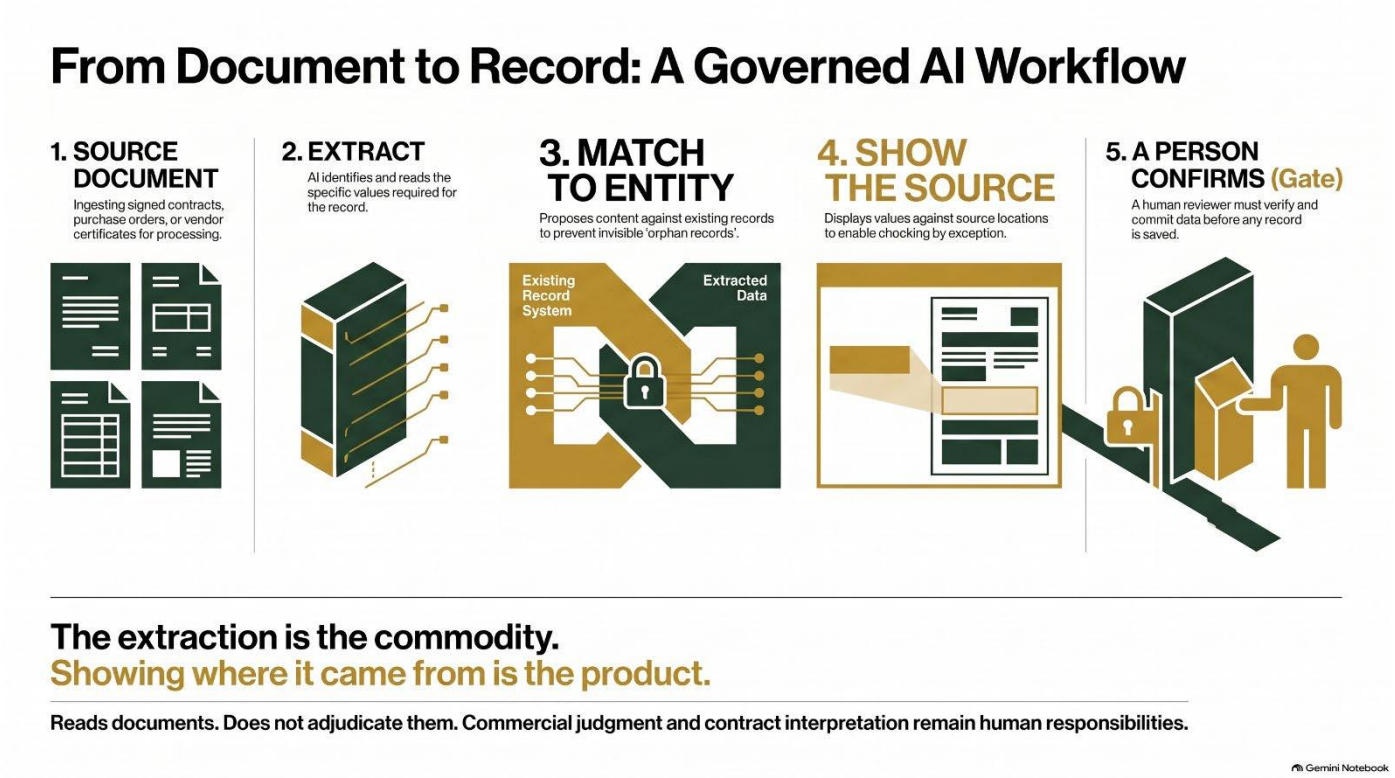
ARTICLE 6 OF 16 — AI FOR PMIS SERIES

From scanned contract to structured record — with the source shown

A contract is signed. A purchase order is issued. A vendor submits a bond certificate. Each of these arrives as a document, and none of them is governed until it becomes a record.

That conversion is one of the least discussed and most expensive activities in capital delivery. Somebody reads the document, finds the values that matter, decides which existing record they belong to, and types them in. On a large programme this happens thousands of times, it is done by people whose time is worth considerably more than data entry, and it is where a meaningful share of data quality problems originate.

Extraction is the obvious application of AI to this, and every vendor in the market is building it. What separates a useful implementation from a dangerous one is not the extraction.



A contract is not governed until it becomes a record. That conversion is one of the least discussed and most expensive activities in capital delivery.

What the capability does

In PMWEB 2026, the assistant reads a source document — a scanned contract, a purchase order — and extracts the information needed to create the corresponding record.

So far this describes every extraction product ever built. Two further behaviours are what make it usable in a governed environment.

It proposes where the content belongs

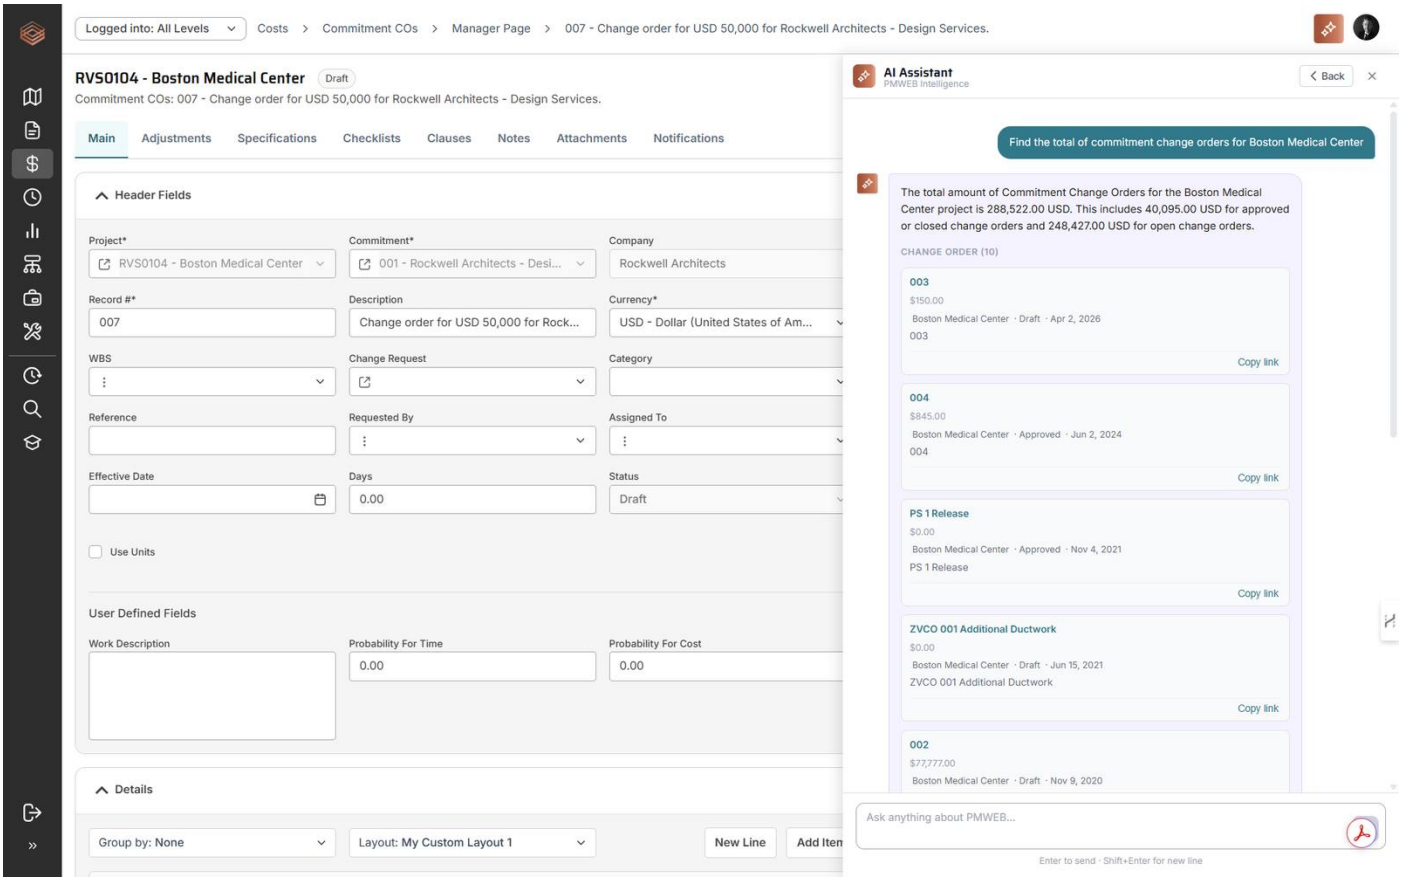
Extracted content is matched to the relevant PMWEB entity rather than being used to create a standalone record.

This addresses a failure mode that is easy to underestimate. An extraction engine that reads a change order and creates a change order has done its job by its own measure. But if that record is not connected to the commitment it varies, it is an orphan. It will not appear in the commitment's revised value. It will not show in exposure reporting against the vendor. It exists, it is accurate, and it is invisible to every question anyone will actually ask.

A database accumulating orphan records degrades quietly. Nothing breaks. The totals are simply wrong, and they stay wrong until someone reconciles manually and discovers why.

It shows you where every value came from

Each extracted value is displayed against its location in the source document. The reviewer verifies rather than trusts, and confirms before the record is saved.



This is the design decision that determines whether the capability saves any time at all, and the arithmetic is worth doing explicitly.

Extraction you cannot verify has to be checked in full. Someone opens the source document, reads it, and compares every extracted value against it — which costs approximately what typing the record would have cost, plus the risk that checking is more boring than typing and therefore done less carefully.

Extraction with the source shown is checked by exception. The reviewer's eye goes to the values, sees where each one came from, and stops at the one that looks wrong. The work is no longer reading a document. It is confirming a mapping.

This is the requirement PMWEB places first on its own buyer checklist for AI: can it show the source record behind every number? Extraction is the point at which that requirement is either met or lost. A value that arrives in the

record with its source location attached stays traceable for the life of the record. A value that arrives without one becomes, from that moment, a number nobody can account for.

> **Unverifiable extraction costs what typing costs. Extraction with the source shown is checked by exception. That difference is the entire business case.**

Where the gate sits

Nothing is saved without a person. The reviewer confirms before the record is created, and where the AI has prepared updates to an existing record, those updates remain pending until someone applies them.

For an organisation whose contract records feed payment certification, this is not a nicety. A record created from a document without human confirmation is a commercial position taken by a machine, and it will eventually be a commercial position nobody can account for.

What it does not do

One boundary is worth stating clearly, because it is the point where expectations run ahead of the capability.

This reads documents. It does not adjudicate them. A clause that is ambiguous in the contract is ambiguous in the extracted record. Where a document contains terms that conflict with an earlier instrument, extraction will faithfully surface what the document says, and deciding which instrument governs remains a commercial judgement made by a person who understands the contract chain.

That is the correct division, and it should be presented as such rather than glossed. A capability that claimed to resolve contractual ambiguity would be a capability nobody responsible should deploy.

What to establish before you deploy it

Four questions worth answering with your own documents rather than a vendor sample.

First, what is the document estate actually like? Extraction performs differently on a born-digital PDF and on a scanned photocopy of a fax from 2011, and most programmes have both. Test on the worst of what you hold, not the best.

Second, who reviews? The value of the capability is realised at the confirmation step, so it matters whether that step lands with someone who understands what the values mean. Extraction reviewed by someone without commercial context is extraction that gets confirmed without being verified.

Third, what is the exception path? When the reviewer sees a value that is wrong, what happens — do they correct it in place, reject the extraction, or escalate? A capability without a defined exception path produces a queue of documents nobody wants to own.

Fourth, where do the documents go? A scanned contract carries prices, terms and positions a counterparty expects to stay confidential. Establish whether documents processed by the AI leave your tenant, how long they are retained, and whether they are used to train models serving other organisations. As PMWEB's own checklist advises, get the answers in the contract rather than the sales deck.

> **This reads documents. It does not adjudicate them. That is the correct division of labour, and it should be presented as such rather than glossed.**

Next week: generating an adaptive form from a plain-language description, and what that does to the economics of process change.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

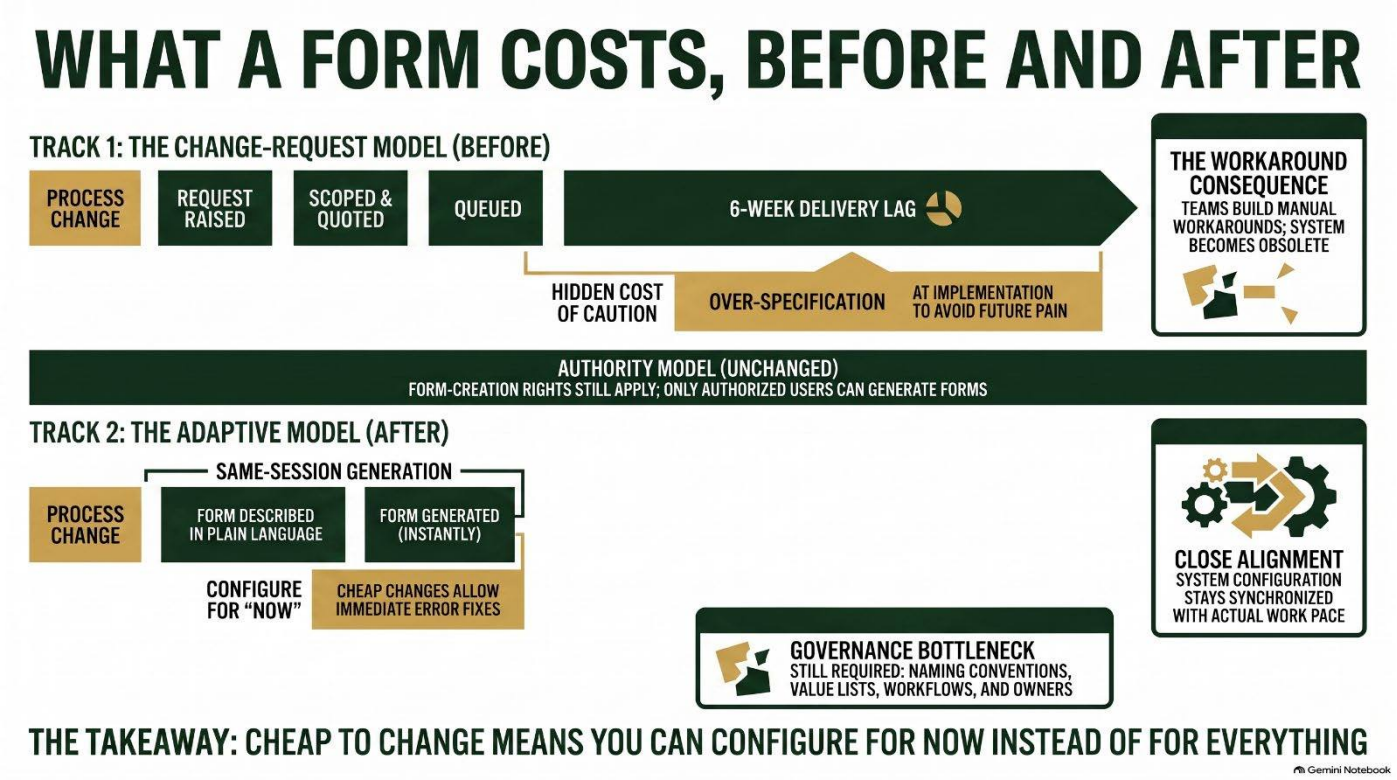
ARTICLE 7 OF 16 — AI FOR PMIS SERIES

Describe the form, get the form

Here is a pattern familiar to anyone who has run a configured PMIS for more than two years.

A process changes. The form in the system no longer matches how the work is actually done. Someone raises a change request. It is scoped, quoted, queued, and scheduled behind three other things. Six weeks later it arrives, by which point the team has developed a workaround, and the workaround is now the process. The form is updated anyway. Nobody uses the new field.

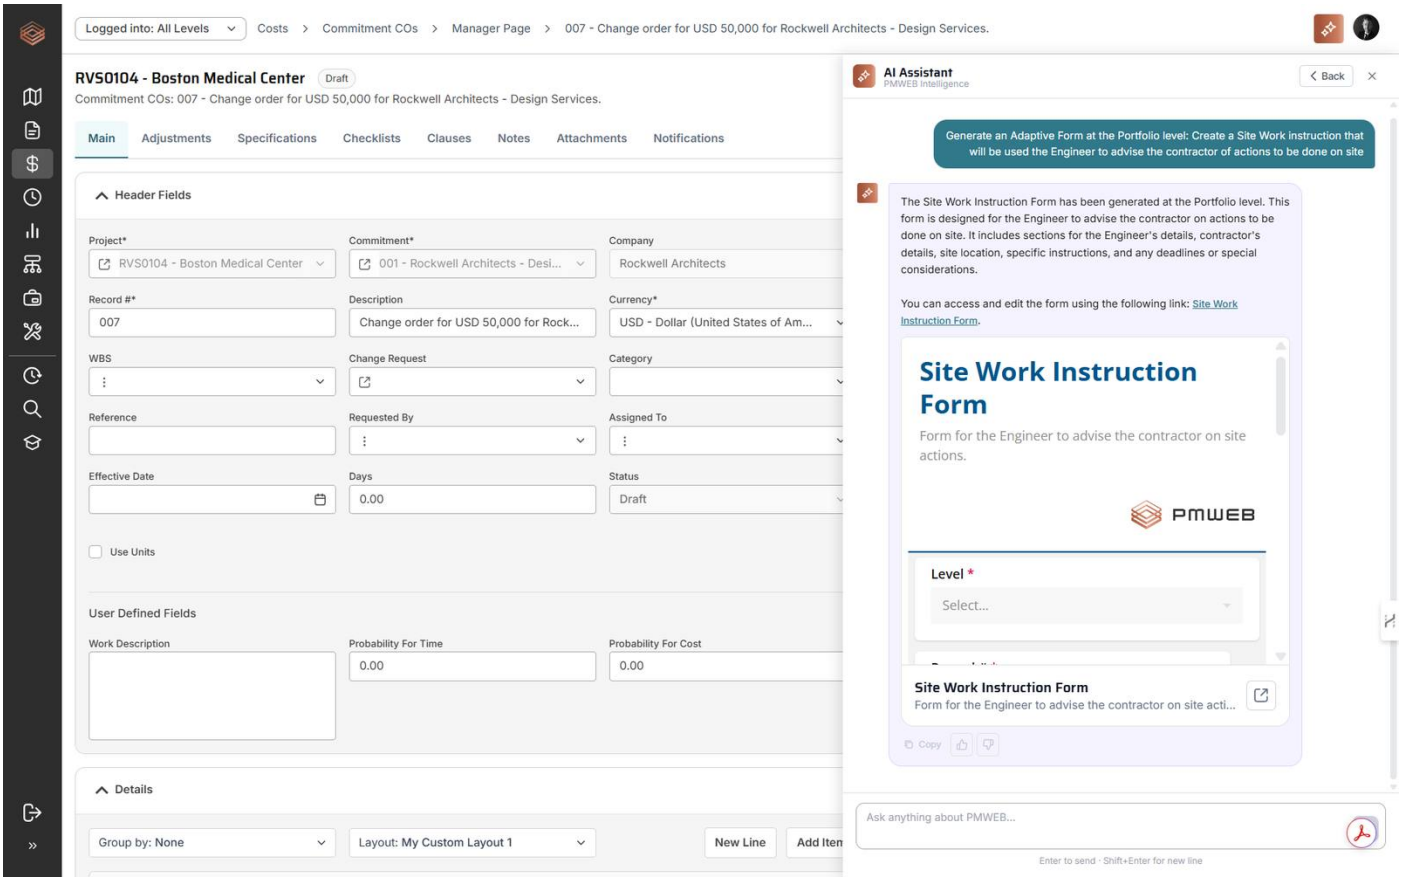
Repeat this a few times and an organisation stops asking. The configuration freezes at the point the implementation ended, the gap between the system and the work widens quietly, and the shortfall is eventually attributed to the platform.



Repeat the change request cycle a few times and an organisation stops asking. The configuration freezes, the gap between the system and the work widens quietly, and the shortfall is eventually attributed to the platform.

What changes

In PMWEB 2026, an adaptive form — its sections and its fields — can be generated from a plain-language description. In the demonstration, a safety inspection form was produced from a description of what it needed to capture.



The mechanics matter less than the change in cycle time. A form that would have been a change request becomes a task completed in the meeting where the need was identified.

The control that keeps this safe

The first reaction from anyone responsible for a large configuration is alarm, and it is the correct instinct. A system in which any user can generate forms is a system that will contain four hundred forms within a year, of which perhaps sixty are used and none reconcile for reporting.

Form-creation rights still apply. This is not every user redesigning the platform. It is the people who already hold that authority doing in minutes what previously took a cycle.

The distinction is worth being precise about when this is presented internally, because the objection will be raised and the answer is straightforward: the authority model does not change. What changes is how long it takes the people who already have that authority to act on it.

It is the same principle PMWEB's Intelligence Control Framework applies across the platform, stated as AI advises, the operator decides. Here the AI proposes the structure of a form. A person who holds form-creation authority decides whether it becomes part of the configuration.

What it means for a PMO

The interesting consequence is not speed. It is that process governance becomes iterative.

Under a change request model, every configuration decision carries the weight of being difficult to reverse. That makes people cautious in the wrong direction. They over-specify at implementation, because a field they might need later is cheaper to add now than in eighteen months, and the result is a form with fifty fields of which twenty are populated. The cost of that caution is paid every day by every user.

When a form is cheap to change, the correct strategy inverts. Configure for what the process needs now. Fix it when it is wrong. A form that is not working in month two gets corrected in month two rather than surviving until the next major review, and the system stays close to the work.

> **The cost of a change request is not the change request. It is the over-specification that organisations adopt to avoid needing one.**

What it does not solve

Generating a form is not governing one, and the distinction should be stated plainly rather than discovered later.

A generated form still needs naming conventions consistent with the rest of the system, so that reporting can find it. It still needs its fields mapped to the controlled value lists the organisation uses, or it produces free text where the rest of the platform produces structured data — and free text is precisely the substrate problem this series began with. It still needs a workflow, an owner, and a decision about who sees it.

None of that is generated. The bottleneck moves from building the form to governing it, which is a considerably better place for the bottleneck to sit, but it does not disappear and no organisation should plan as though it has.

The practical position

For a PMO, the sensible response is to treat this as a change in the economics of configuration rather than as a new feature.

It means the configuration standard becomes more important than it was, because the volume of configuration decisions is about to increase. Write down the naming convention, the value list policy and the approval route for a new form before the capability is enabled, not after. Organisations that do this get an iterative system. Organisations that do not get four hundred forms.

> **Write down the naming convention, the value list policy and the approval route for a new form before the capability is enabled, not after. Organisations that do this get an iterative system. Organisations that do not get four hundred forms.**

Next week: what a project summary has to do before anyone can act on it.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in

Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

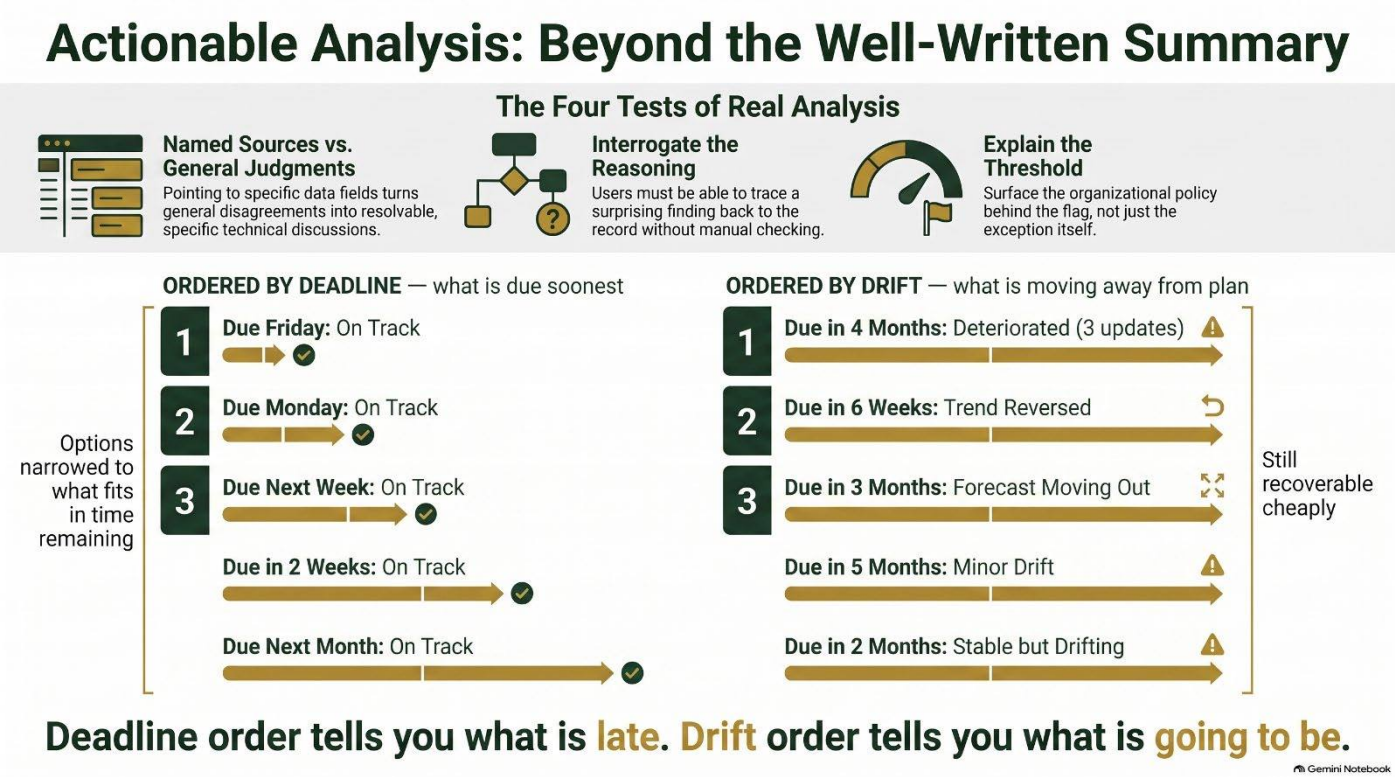
ARTICLE 8 OF 16 — AI FOR PMIS SERIES

What a project summary has to do before you can act on it

This article is about what capital delivery needs from AI-assisted analysis, not about what any platform ships today. The tests below are worth applying to whatever your current system produces.

Summarisation is the capability every vendor is demonstrating this year, and it demonstrates beautifully. A record opens, a paragraph appears, the paragraph is well written and broadly correct, and the room is impressed.

Then the capability meets a contested programme, and a well-written paragraph turns out to be worth very little.



A summary that cannot withstand the question "on what basis?" is not analysis. It is an opinion delivered in a confident register.

Why fluency is not the test

In capital delivery, a reading of the situation is rarely accepted quietly. It is challenged — by the contractor, by the client's commercial team, by the cost consultant, eventually by whoever is appointed if the matter escalates. A summary that cannot withstand the question "on what basis?" is not analysis. It is an opinion delivered in a confident register.

Four tests separate analysis that survives that question from analysis that does not.

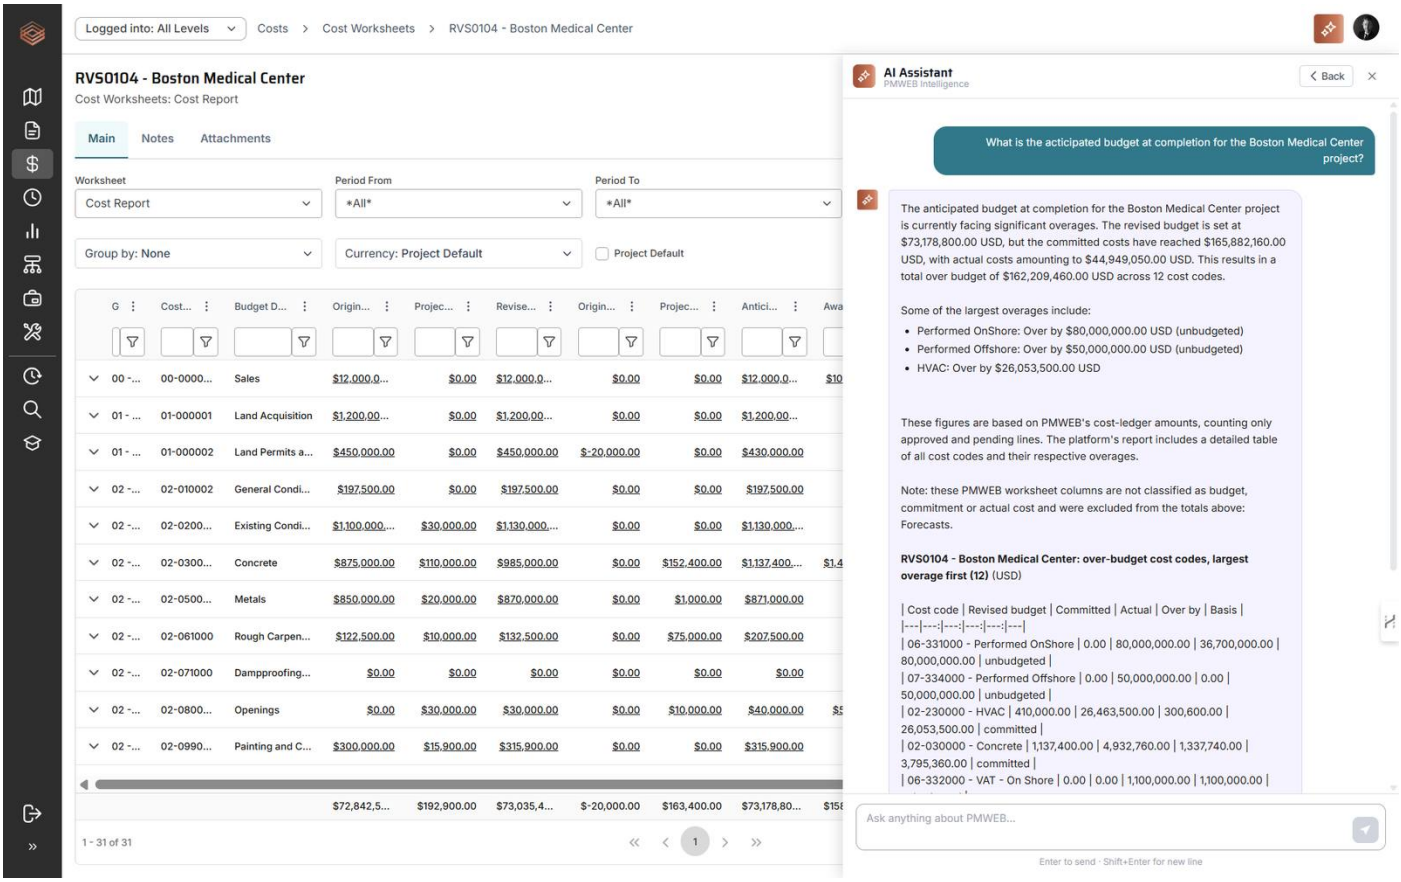
Test one: does it name its sources?

A summary that states a conclusion and a summary that names the fields the conclusion was drawn from look similar on screen and are entirely different instruments.

The first can only be accepted or rejected. The second can be examined. When somebody disagrees with it, the disagreement becomes specific — they are disputing a value in a named field rather than disputing a judgement — and specific disagreements get resolved, while general ones become positions.

Require that any reading points to the data behind it. If it cannot, treat it as a prompt to go and look rather than as a finding.

PMWEB puts this test first on its own buyer checklist for AI, published in September 2026, as a question: can it show the source record behind every number? Its example is the right one. An AI forecast that cites the pay application or schedule revision behind it, which the reviewer can open and check, is analysis. A figure whose only provenance is that the system suggested it is not.



Analysis that names its sources: the PMWEB 2026 AI Assistant states the basis of its answer and lists every cost code behind the figure. Illustrative data.

Test two: can you interrogate it?

A fixed output is a report. You take it or you leave it. What is needed is the ability to ask the follow-up question and have the reasoning traced back to the record.

The distinction shows up in the moment that actually matters, which is when the summary says something surprising. If the response to a surprising reading is to open the module and start checking manually, the capability has produced a suspicion rather than an answer, and it has cost more than it saved.

Test three: is the threshold explained, not just applied?

Take a change order sitting at sixty-two per cent of the original commitment value. A system that flags it has told you a number crossed a line. A system that explains the threshold, shows the commitment history behind the figure, and lets you see the change orders that accumulated to it has told you something you can act on.

Thresholds are organisational decisions, and they are frequently inherited from a policy nobody has revisited. Analysis that surfaces the threshold alongside the exception lets you question whether the threshold is right, which is often the more valuable finding.

Test four: does it rank by drift, or by date?

This is the test most systems fail, and it is the one with the largest practical consequence.

Almost every task list, dashboard and action report in our industry is ordered by deadline. What is due soonest appears first. This is intuitive and it directs attention to the wrong things, because by the time a deadline is close the options have narrowed to whatever can be done in the remaining time.

Ranking by drift asks a different question: what is moving away from where it should be? An item due in four months whose trend has deteriorated across three consecutive updates is more deserving of attention this week than an item due on Friday that is on track. The first is still recoverable cheaply. The second requires nothing but for someone to finish it.

> **Deadline-ordered lists tell you what is late. Drift-ordered lists tell you what is going to be late, which is the only version you can still do something about.**

The two orderings produce genuinely different top-five lists from the same portfolio. That is not a refinement of the same output. It is a different instrument, and organisations running deadline-ordered attention are structurally late.

The same tests apply to reporting

The four tests extend directly to self-service reporting and to analysis of the working grid, which is where a great deal of day-to-day project work actually happens.

Producing a report from a plain-language description is a genuine advance for organisations where every new report is a request to somebody else. But a generated report is subject to exactly the same requirement as a generated summary: it must be traceable to the data underneath it. A report that cannot be reconciled to the records it came from will be reconciled by hand the first time somebody senior disagrees with it, and after that it will not be trusted.

Analysis you cannot trace is decoration. It may be accurate decoration, and it will still be treated as decoration the moment it matters.

Apply the tests where you are

None of this requires waiting for a new release. The four tests apply to the reports, dashboards and exception lists your organisation produces today, most of which were not generated by AI at all.

Take your current portfolio dashboard. Can you trace any number on it back to the records behind it without leaving the screen? Can you ask it why? Does it show you the threshold or only the breach? And is it ordered by what is due, or by what is moving?

Most organisations find that their existing reporting fails at least two of the four. That is worth knowing before adding a capability that will inherit the same weaknesses and express them more fluently.

Most organisations find that their existing reporting fails at least two of the four tests. That is worth knowing before adding a capability that will inherit the same weaknesses and express them more fluently.

Next week: what it takes for a system to raise its hand rather than wait to be asked, and why the value collapses if the decision point is removed to make it faster.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZCO, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

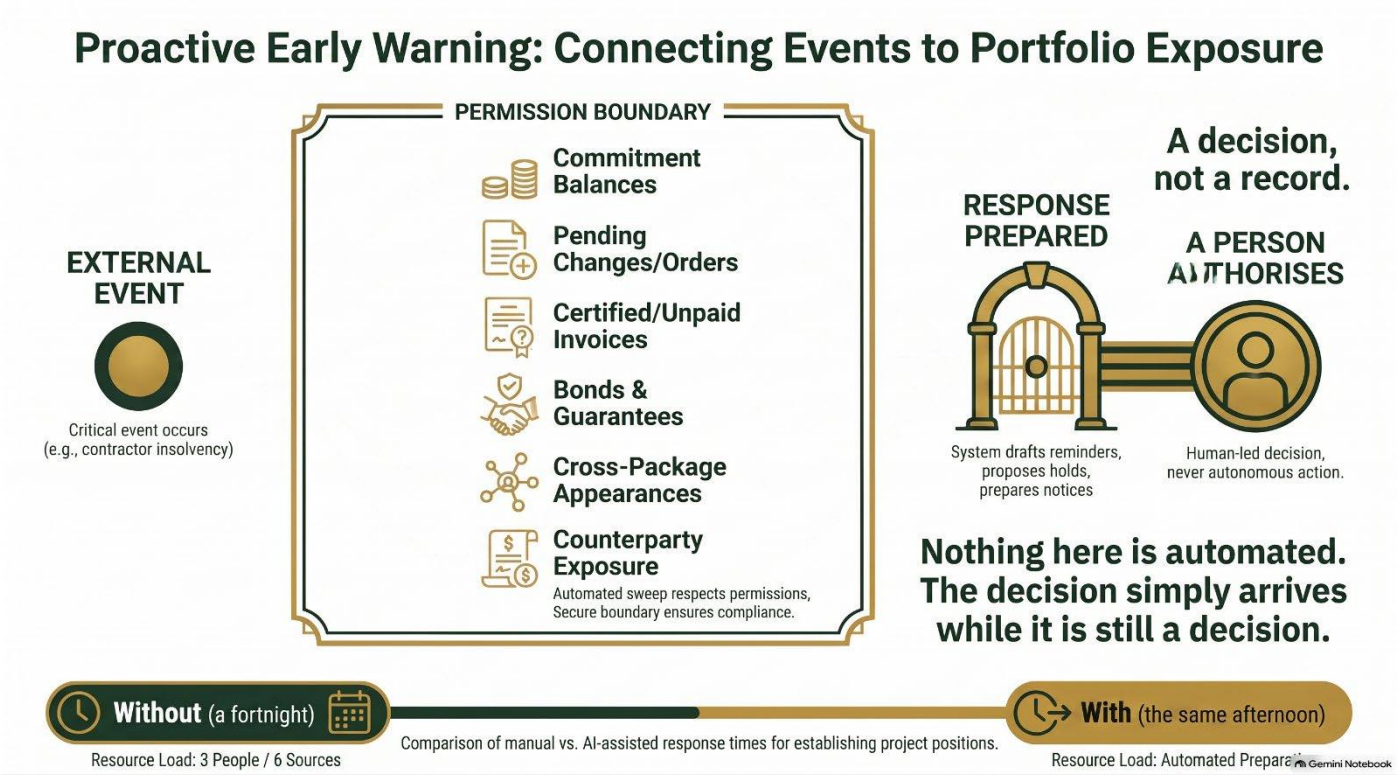
ARTICLE 9 OF 16 — AI FOR PMIS SERIES

Early warning, and the gate that has to stay closed

This article is about what capital delivery needs from AI-assisted monitoring, not about what any platform ships today. The scenario below is illustrative.

Every system discussed so far in this series waits to be asked. You open a record, you run a search, you request a summary. The system is a resource and you are the one who has to know to consult it.

The next thing worth wanting is a system that raises its hand.



A monitoring capability that watches only the workflow will report a healthy programme right up until the reconciliation.

What has to be watched

Continuous monitoring is usually described in terms of workflow events — an approval overdue, a submission missed, a milestone passed. That is the easy half and it is not where the exposure hides.

The harder and more valuable half is inconsistency in the data itself. A revised document whose terms no longer match the record that governs them. A commitment whose certified value has drifted past its revised value. A vendor appearing under two entities. None of these generates an event. Nothing is overdue. The workflow is behaving perfectly and the underlying position is wrong, and it will stay wrong until somebody happens to look.

A monitoring capability that watches only the workflow will report a healthy programme right up until the reconciliation.

A scenario worth working through

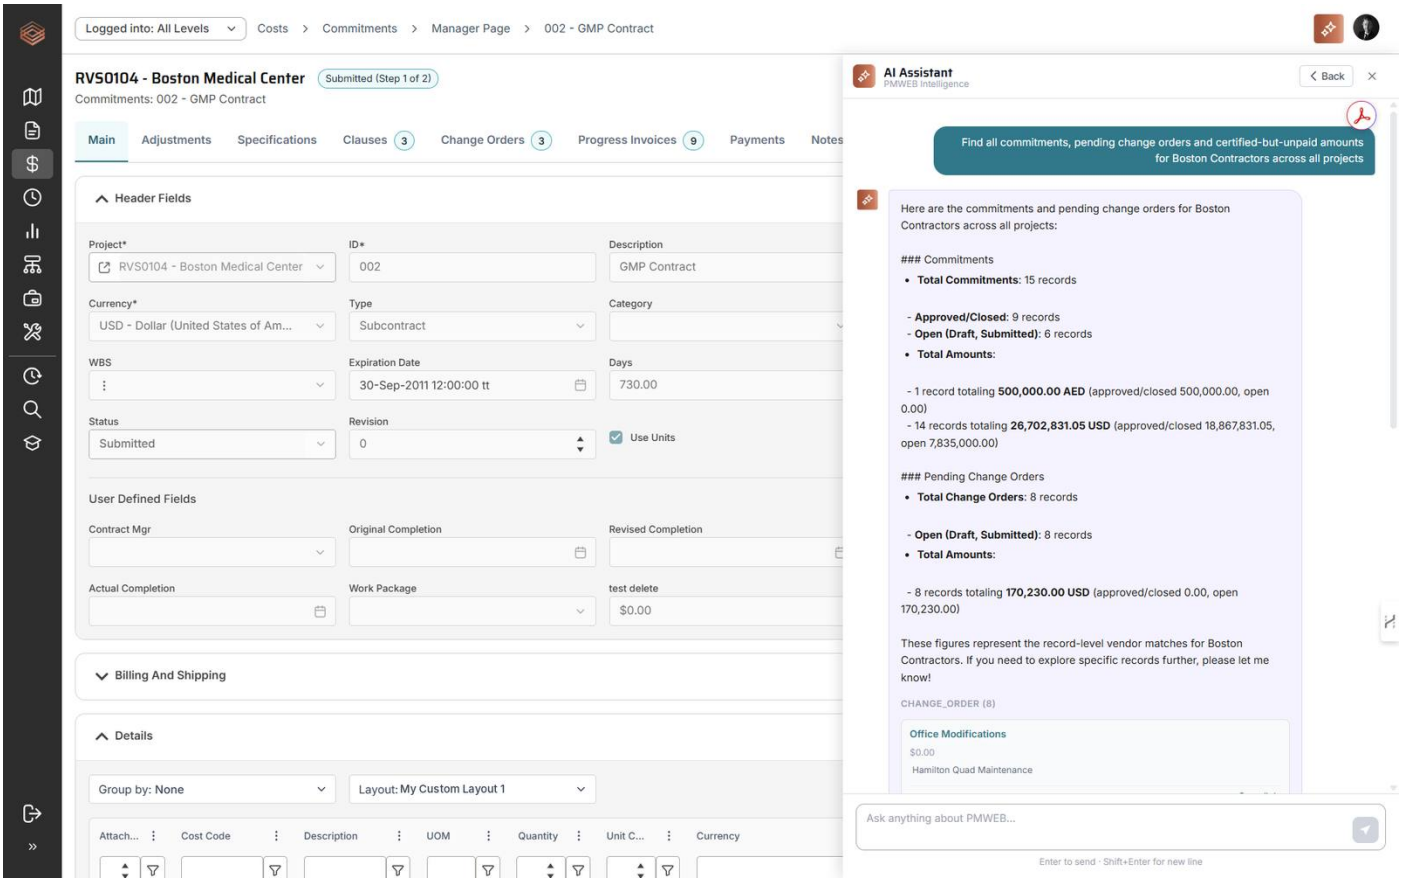
Suppose a contractor on one of your packages enters insolvency proceedings. This is not a rare event in this industry and the sequence that follows is familiar.

The question is not whether you hear about it. You will hear about it, probably from the news before anyone tells you formally. The question is how long it takes to establish your position.

What is the remaining balance on their commitments. What change orders are pending against those commitments and at what value. What has been certified but not paid. What bonds and guarantees are in place and are any close to expiry. And the question that is hardest and matters most: where else in the portfolio does this counterparty appear, including as a subcontractor on packages held by somebody else.

In most organisations that exercise takes a fortnight and is assembled from six sources by three people, at least one of whom is on leave. The answer arrives after the commercially useful window has closed.

A monitoring capability worth having connects the external event to the internal exposure and then sweeps the portfolio for the same counterparty — returning, say, six further commitments nobody had connected. Permission-filtered throughout, so nobody sees anything they were not already entitled to see.



Establishing the position in minutes, not a fortnight: the PMWEB 2026 AI Assistant gathers one counterparty's commitments and pending change orders across every project. Illustrative data.

You will hear about the event. The capability worth having is the one that tells you what it means to you, this afternoon rather than next fortnight.

The exception layer

The same principle applies to exposure that arises internally rather than externally.

Overbilling against a revised commitment value, with the specific invoices that produced the figure named rather than a percentage asserted. A vendor's insurance or bond approaching expiry, detected before it lapses rather than during the claim that follows. A revised document compared against the record that governs it, with the

mismatches surfaced — payment terms moved from thirty days to cash on delivery, retention reduced from ten per cent to five.

Each of these is discoverable today by someone who goes looking. The difference is that nobody has time to go looking at everything, so the looking is sampled, and the exposure that matters is the one that fell outside the sample.

And then the response, which is where the discipline is required

Detection is only useful if something follows, and this is the point at which the design decisions matter most.

A capability that identifies an expiring bond can draft the reminder. A capability that identifies overbilling can propose restricting invoice approval while the position is resolved. A capability that identifies counterparty exposure can propose a payment hold that blocks the workflow and signals the issue across that vendor's other commitments. It can draft the notice to the other project leads affected.

Drafted. Proposed. Prepared. Never done.

The reminder is drafted for a person to review, amend and send. The payment hold is proposed for a person to authorise. As established earlier in this series, the audit trail must name the person, and the more consequential the action the more that matters. A payment hold placed automatically is a commercial position taken by a machine against a counterparty who may be in dispute with you, and it is precisely the kind of action that will be examined line by line later.

PMWEB states the principle behind its Intelligence Control Framework in five words: AI advises, the operator decides. Monitoring is where that principle is tested hardest, because it is where the advice arrives with the most urgency attached.

The objection worth answering

If a person has to authorise everything, what was gained? The system found the problem faster, but a human still does the work.

What was gained is the interval. In the insolvency scenario, the difference between establishing your position on the afternoon of the filing and establishing it a fortnight later is the difference between acting on it and documenting what happened. Nothing about the decision was automated. The decision simply arrived while it was still a decision rather than a record.

That is the honest value proposition for this class of capability, and it is a strong one. It does not require anybody to accept an autonomous actor inside their control environment, and organisations should be sceptical of any version that asks them to.

What to ask for

When this capability is offered to you, three questions establish whether it is the version worth having.

Does it monitor the data or only the workflow? Does the sweep respect existing permissions, including where a result connects activity across records the user reaches separately? And at the end of the chain, when a response has been prepared, what does the system do — and what does the audit trail record?

The answer to the third should be that it waits, and that the trail names a person.

PMWEB's September 2026 checklist adds the detail that makes the third question testable. Every AI interaction should be logged with the user, the timestamp, the query, the response and the resulting action. That matters more for monitoring than for any other class of capability, because the responses it prepares (a payment hold, a restricted

approval, a notice to a counterparty) are the ones most likely to be examined later. When someone asks why a vendor's payments were held in March, the record should show what the system flagged, what it proposed, who reviewed it and what they decided.

> **Nothing about the decision was automated. The decision simply arrived while it was still a decision rather than a record.**

That concludes Part II. Part III turns to the layer no PMIS vendor supplies: capture at archive scale, and reasoning that spans the systems a programme actually runs on.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZCO, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

ARTICLE 10 OF 16 — AI FOR PMIS SERIES

Where a PMIS needs help: capture, OCR and the archive nobody has read

A well-run PMIS holds an accurate, structured, permissioned account of everything that was entered into it. That is a considerable asset and this series has spent eight articles explaining why it matters.

It is also a fraction of what a capital programme produces.

The rest sits in letters, meeting minutes, transmittals, RFI correspondence, technical submittals, site diaries, photographs, drawing revisions and an email estate that nobody has ever measured. On a programme of any age this material outweighs the structured record by an order of magnitude, and it carries most of the contractual history — including, reliably, the parts that become contentious.



No PMIS closes that gap on its own, and a vendor who claims otherwise is describing a document management product.

> **No PMIS closes the gap between the structured record and the archive on its own, and a vendor who claims otherwise is describing a document management product.**

Two capabilities that get confused

The gap is usually discussed as one problem and it is two, with different engineering behind them.

The first is capture. Turning paper, scans, images and unstructured files into text and structure at volume. This is a mature discipline with its own accuracy engineering, its own failure modes and its own specialists, and it is measured in documents per hour and error rates per field.

The second is reasoning. Connecting what was found in that material to what the records say — which contract governs this letter, which commitment this instruction varies, whether the position described in correspondence matches the position in the system.

Products are frequently sold as though these are one capability. They are not, and an organisation that buys a reasoning layer without solving capture has bought something with very little to read.

Why your PMIS vendor is the wrong supplier for the first

This is worth saying plainly, including by those of us who implement a PMIS for a living.

High-volume capture at archive scale is not adjacent to project controls. It is a different engineering problem, and the organisations that are good at it have spent years on the specific failure cases — the photocopy of a fax, the drawing with handwritten markup, the scanned page bound at an angle, the table that spans three pages with a repeating header.

A platform vendor building capture as one feature among forty will not match that, and buying it bundled generally means buying the weakest available version of a capability you could have bought properly. The bundle is convenient. Convenience is not the criterion when the output feeds a contractual record.

> **Buying a specialist capability as a bundled feature usually means buying the weakest available version of it.**

What good looks like at this layer

Four properties separate capture worth deploying from capture that generates a review queue nobody will work through.

Confidence scoring on every extracted value, so that the system distinguishes what it is sure of from what it guessed. A system that presents all output with equal certainty forces uniform review, and uniform review does not scale.

Exception routing driven by those scores, so that human attention goes to the values that need it. This is the same drift argument from earlier in this series applied to document processing: attention is the scarce resource, and directing it is most of the value.

Human verification designed as a fast confirmation step rather than a re-reading. The reviewer should be confirming a mapping, not reading a document.

And traceability back to the source page, for the same reason it mattered in the PMIS: an extracted value that cannot be traced is a value that gets re-checked manually the first time it is disputed.

The decision that determines whether the project finishes

Here is the pattern that sinks archive projects, and it is a scoping error rather than a technology one.

An organisation decides to bring its legacy correspondence under control. The scope is set as extracting governed records from the archive. The archive turns out to contain forty thousand documents. The pilot goes well on two hundred. The full run is estimated, the number is uncomfortable, the project is paused for review and it does not restart.

The alternative is to separate two activities that get merged.

Classification is cheap. Working out what each of forty thousand documents is — its type, its date, its parties, which project and which contract it belongs to — is fast, largely automatable, and immediately useful on its own. It converts an undifferentiated pile into a searchable index. Most of the operational pain of a legacy archive is not that the content is unreadable; it is that nobody can find the four documents that matter.

Extraction is expensive. Pulling governed values out of a document and turning them into records carries a verification cost per document that does not fall much with volume.

Classify everything. Extract from the fraction that warrants it — the live obligations, the open commitments, the disputed history. Organisations that scope it this way finish. Organisations that scope it as one exercise generally do not.

What to ask a capture partner

Test on your worst documents, not a clean sample. Every product demonstrates well on a born-digital PDF and your archive is not made of those.

Ask what happens to a value the system is unsure about, and require the answer to be that it is flagged rather than filled.

Ask how output reaches the PMIS. A capture capability that produces a spreadsheet somebody re-keys has moved the work rather than removed it.

And ask about entity matching specifically, because it is the same requirement as inside the platform: content extracted from a document needs to attach to the record it belongs to, or you have built a second archive alongside the first.

And ask what happens to the documents. An archive holds contract prices, claim positions and correspondence a counterparty expects to stay confidential, and a capture partner processes it at volume. The question PMWEB puts on its own buyer checklist applies here with more force: is it reading your data, or training on it? Establish where documents are processed, whether they leave your tenant, how long they are retained and whether they train models serving anyone else, and get the answers in the contract rather than the sales deck.

Classify everything. Extract from the fraction that warrants it — the live obligations, the open commitments, the disputed history. Organisations that scope it this way finish.

Next week: the questions that cross three systems at once, and the layer that answers them.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington

University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

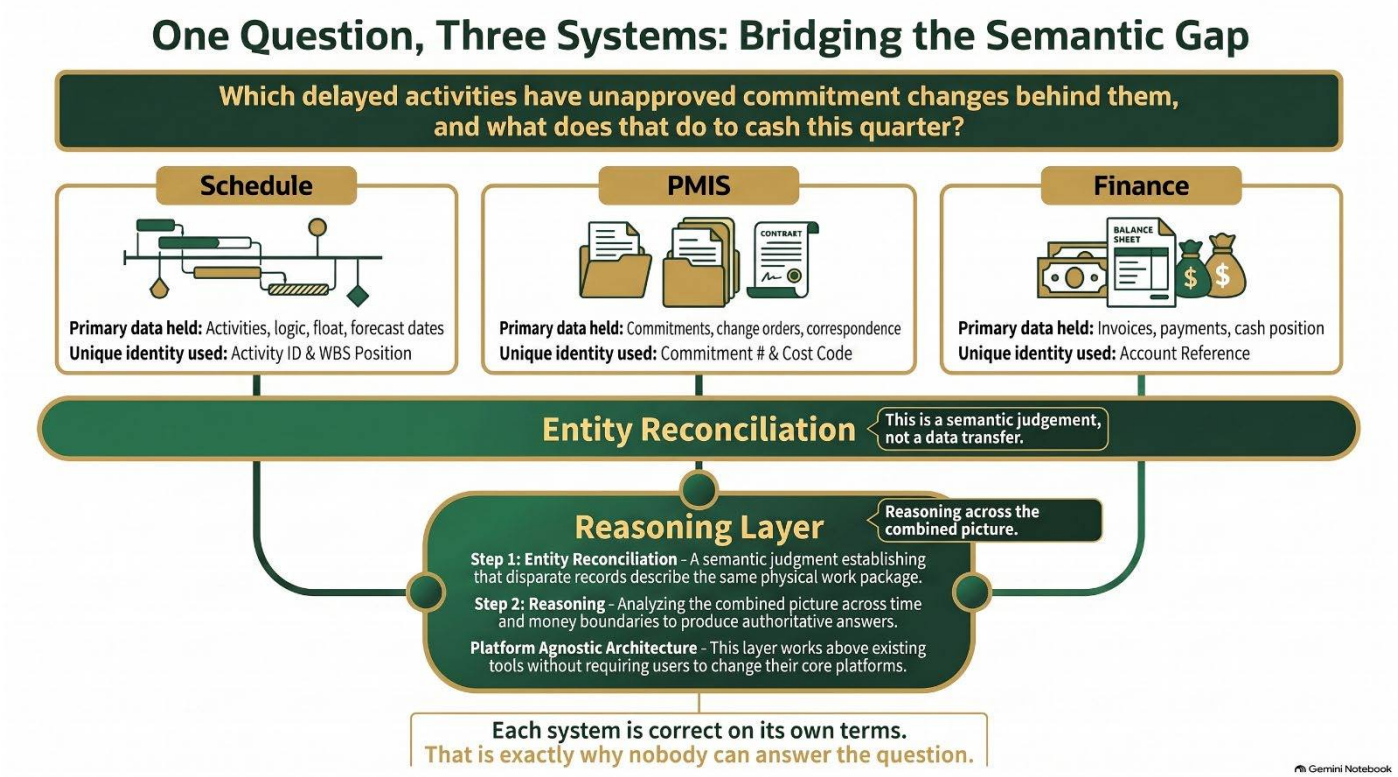
ARTICLE 11 OF 16 — AI FOR PMIS SERIES

Reasoning across platforms: when the answer lives in three systems

Here is a question a development director asked in a meeting recently, and it is worth examining because it is completely reasonable and almost impossible to answer.

Which of our delayed activities have unapproved commitment changes sitting behind them, and what does that do to the cash position this quarter?

Every component of that question exists in a system somewhere. The delay is in the schedule tool. The commitment changes are in the PMIS. The cash position is in finance. Nobody in the room could answer it, and the reason is not that anyone had failed to do their job.



Every component of the question exists in a system somewhere. Nobody in the room could answer it, and the reason is not that anyone had failed to do their job.

Why this is structurally hard

The instinct is to treat this as an integration problem that somebody has neglected. It is not, and understanding why matters before anyone commits budget to solving it.

Each system holds a different identity for the same real-world thing. The schedule knows an activity by its ID and its WBS position. The PMIS knows a commitment by its number and its cost code. Finance knows a vendor by an account reference. The same package appears in all three under three names, structured three ways, at three levels of granularity — because each system was designed for a different question and each is correct on its own terms.

Joining them is therefore not a technical exercise of moving fields between databases. It is a semantic exercise of establishing that this activity, this commitment and this account refer to the same thing. That judgement has traditionally been made by a person, which is precisely why the answer takes a fortnight and arrives as a spreadsheet.

The join between project systems is semantic, not technical. That is why integration projects deliver data movement and not answers.

What a layer above the platforms does

This is the second half of the partner layer described in the previous article, and it does two things in sequence.

First it reconciles entities across systems — establishing and maintaining the correspondence between the activity, the commitment, the vendor and the account. Then it reasons over the combined picture, which is only possible once the first step holds.

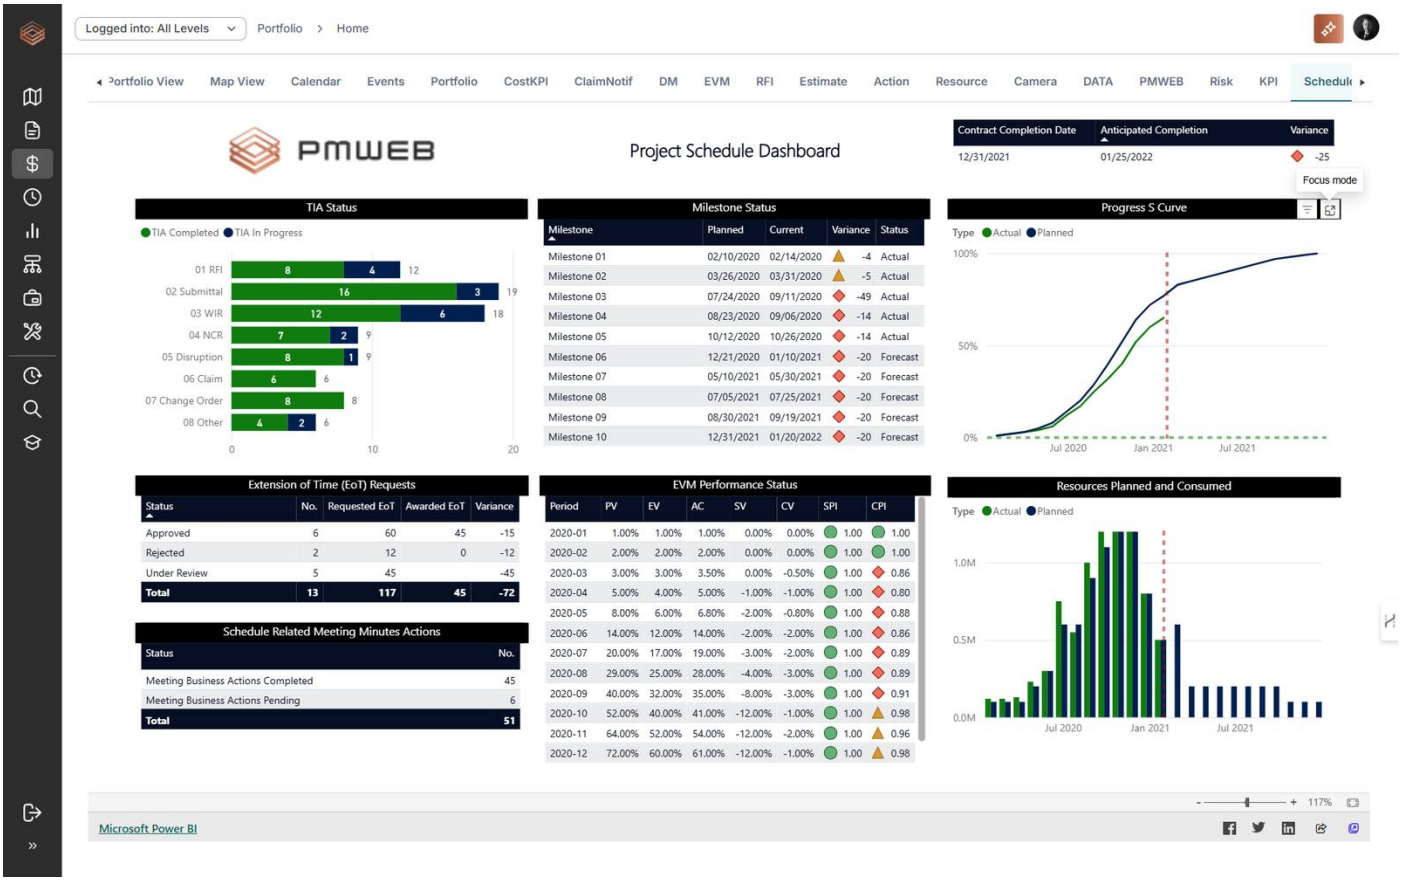
The ordering is the whole thing. Reasoning applied to systems that have not been reconciled produces answers that look authoritative and are quietly wrong, which is worse than no answer at all. Any product offered at this layer should be asked how it establishes the correspondence and what it does when the correspondence is ambiguous.

The case for developers and PMC firms

The organisations for whom this is most acute are those running a PMIS for commitments, cost and correspondence alongside a scheduling tool for time — which is to say most real estate developers and most project management consultancies in this region.

That combination is not an accident of procurement. It is the right architecture, because a PMIS and a scheduling engine are genuinely different tools and every attempt to make one do the other's job has produced a weak version of both. The consequence is that the two most important dimensions of a project, time and money, live in separate systems by design.

The questions that matter most cross that boundary. Which delays have commercial consequences. Which variations have schedule consequences. Whether the forecast completion and the forecast final account are telling the same story. These are not exotic analytical questions. They are the standard questions of project control, and they are hard to answer because of an architecture that is otherwise correct.



On evidence of scale

One point worth making about this category, because it is young enough that caution is warranted.

Ask for evidence from portfolios comparable to yours in size and complexity. Public infrastructure agencies running large capital programmes are a useful reference class, because their reporting obligations are demanding and their portfolios are genuinely heterogeneous — several delivery partners, several systems, several decades of history. A capability that holds up in that environment has been tested against real conditions rather than a clean pilot.

Anonymised examples are reasonable and should be expected. What you are assessing is whether the volume and structural complexity resemble yours, not whose logo is on it.

This does not require you to change platform

A deliberate point, and the reason this article sits where it does in the series.

A reasoning layer above the platforms is not an argument for replacing them. It works above a PMIS, a scheduling tool and a finance system regardless of which products those are, and the case for it is as strong for an organisation running Oracle Unifier and Primavera, or Procore, as it is for one running PMWEB.

That is the honest position and it is also the commercially useful one, because the question this layer answers is not which platform should we have chosen. It is what do we do about the fact that our data is correctly held in three places.

What to establish before committing

How does it connect — and specifically, does it require a project to move or copy data into a new repository, or does it read the systems where the data already lives? The first creates a fourth system to govern.

How is the entity correspondence maintained when a cost structure changes, which it will. A reconciliation that holds at implementation and degrades over eighteen months is a liability.

And what happens to permissions. A layer that reads three systems must respect the access model of all three, and it is worth testing that a user with restricted access in the PMIS does not gain a wider view through the analysis layer than they hold in the source. The principle PMWEB applies inside its own platform is the right one to require of any layer above it: AI should work under the permission model the owner already administers, not a second one. A reasoning layer with its own access model is a second one by definition, and the time to find that out is before the first board question that relies on it.

> **The question this layer answers is not which platform we should have chosen. It is what we do about the fact that our data is correctly held in three places.**

Next week, closing Part III: twelve questions to ask your PMIS vendor about AI, whichever platform you are running.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

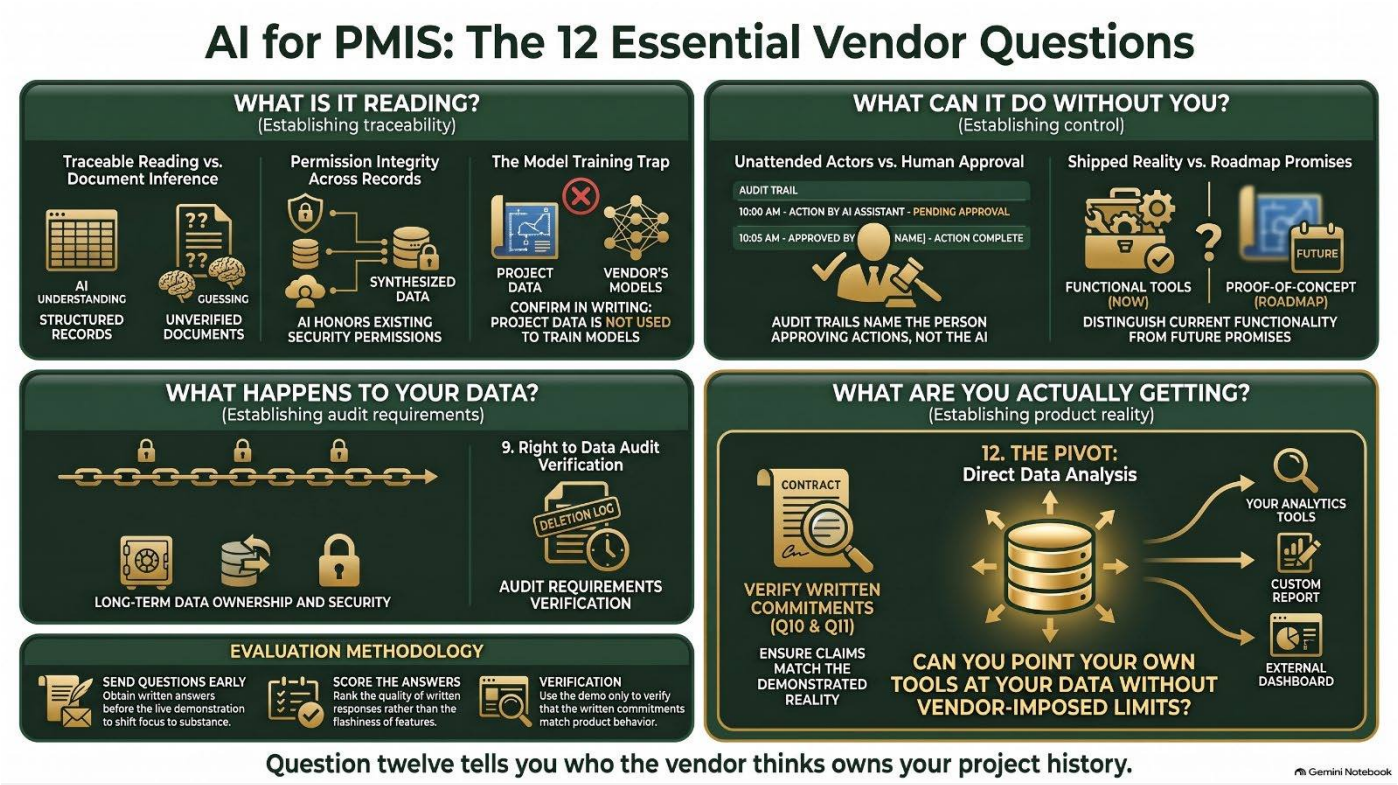
ARTICLE 12 OF 16 — AI FOR PMIS SERIES

Twelve questions to ask your PMIS vendor about AI

Every major platform serving capital delivery is shipping AI capability this year. If you have sat through more than two vendor presentations you will have noticed that the announcements are difficult to tell apart. The same adjectives. The same chat panel. Frequently the same demonstration, which is a natural-language question producing a well-formatted answer.

The differences that matter are not in the announcements. They are in the answers to questions the announcements do not address, and most of those questions have to be asked deliberately because no vendor volunteers them.

What follows is a framework, not an argument for a platform. It applies to whichever product you are running or evaluating.



Every major platform is shipping AI capability this year and the announcements are difficult to tell apart. The differences that matter are in questions no vendor volunteers.

Group one: what is it reading?

Everything downstream depends on this, and it is the group most often skipped because it sounds technical.

1. What data does the AI actually read, and is it structured records or inferred from documents? An answer drawn from validated records with defined relationships is a different class of answer from one assembled by reading files. Both can be correct. Only one can be traced.

2. What accuracy is claimed, and how was it measured? The honest answer from most vendors today is that no figure exists. That is acceptable. What is not acceptable is a percentage with no methodology behind it, and you should ask what was tested, on what sample, and against what ground truth.

3. Can reporting output be presented in the language your executive stakeholders read? For most organisations in this region the project record is in English and this is not a capture question at all. It is a presentation question, and it applies to the board and ministry reporting layer rather than to the working records beneath it.

Group two: what can it do without you?

4. Can it act without a human? Ask for a demonstration that ends in the system preparing something, then watch what happens. Prepared changes should remain visibly pending. If the action completes, you are being offered an unattended actor inside your control environment.

5. Is every AI action in the audit trail, and does the trail name the person who approved rather than the assistant that prepared? This is the question that will matter in a dispute two years from now, and it is almost never asked during evaluation.

6. Can capabilities be enabled individually? All-or-nothing forces a decision most organisations are not ready to make. Ask what the unit of control is — per capability, per user group, per project, per environment — because that determines how much of your rollout you actually govern.

Group three: what happens to your data?

7. Does the AI honour existing permissions, including where a result connects activity across records the user would reach separately? Synthesis is where permission models fail. A summary that combines several sources can reveal something none of them would have revealed individually.

8. Is your data used to train or improve models, and how long is AI-related data retained? Require this in writing. It is the question your audit committee will ask, and a verbal assurance given during a demonstration is not something an audit committee can accept.

9. Is there a published governance framework, with a date, that constrains the product roadmap? The difference between a vendor with a framework and one without is not the document. It is whether the answers above are design commitments or current behaviour, and current behaviour changes without notice.

Group four: what are you actually getting?

10. Is what you are being shown shipped, or planned? This deserves a direct answer and it is reasonable to ask for it in writing. Roadmaps in this category are frequently marked as proof-of-concept, and a version number on a proof-of-concept list is a target rather than a date. A vendor who distinguishes clearly between the two is telling you something good about how they operate.

11. Is AI usage metered, and what happens when you exceed the allowance? Several platforms are introducing token, query or per-seat limits. This matters more than it appears to, because a capability people are told to ration does not get adopted — teams revert to the workarounds the capability was meant to replace. Ask whether the limit is per seat or per tenant, because only one of those lets you turn a project team loose on it.

12. Can you run your own analysis on your own data?

The twelfth question

The last one is the pivot, and it is worth more attention than the other eleven combined.

Ask a vendor whether you can point your own tools at your own project history — whether there is a database with a documented schema, a reporting replica, direct query access. Then listen carefully, because the answer tells you who the vendor believes owns your project record.

A platform that exposes your data only through a rate-limited interface has made a decision. It may be a defensible decision, taken for sound architectural reasons in a multi-tenant environment. But the consequence for you is that every question your organisation ever wants answered has to travel through the vendor's product roadmap. Your ability to analyse fifteen years of your own delivery history is bounded by somebody else's release cycle.

> **Ask who can analyse your project history. The answer tells you who the vendor thinks owns it.**

This also determines what happens if you ever decide to leave, which is not a hypothetical over a ten-year horizon. The practical difference between a documented relational database and an API-only export is the difference between a migration that is engineering and a migration that is archaeology.

Applying it honestly

It would be convenient to run this framework against PMWEB and report a clean sweep. That is not what an evaluation framework is for, and any framework whose author scores perfectly on it should be treated with suspicion.

PMWEB's positions are strong on the control and permission questions. AI access follows existing platform permissions. The assistant will not finalise an action without explicit approval, and prepared updates stay pending. Capabilities can be enabled individually. Relationships are surfaced for analysis rather than created automatically. Customer data remains contained within the platform. On question 11, PMWEB's stated position is that usage carries no limit.

On question 2 there is no accuracy figure, because none has been published. On question 3 the presentation-layer position is a fair question to put and one we are putting. And on question 10 the distinction between what is shipping and what sits further out is one this series has been careful to observe, which is why several articles in it describe what the discipline needs rather than what any product does today.

Those are the honest answers. A vendor or partner giving you a complete set of positive answers to twelve questions of this kind is not describing a product. They are describing a brochure.

How to use it

Send the twelve before the demonstration rather than after it, and ask for written responses. This changes the meeting. A vendor who has answered in writing arrives prepared to discuss substance, and the demonstration becomes a verification of the answers rather than a performance.

Score the answers rather than the features. Two platforms with near-identical capability lists can produce very different scores here, and the difference will predict your experience over the following three years far better than the feature comparison will.

And record the questions a vendor cannot yet answer, without treating that as a failure. This field is moving quickly and honest gaps are normal. What matters is whether the gap is acknowledged or papered over, because the second tells you what the rest of the relationship will be like.

> **Score the answers rather than the features. Two platforms with near-identical capability lists can produce very different scores here, and the difference will predict your next three years better than the feature comparison will.**

That concludes Part III. Part IV turns to the use cases organisations build for themselves, and to the question of how you start without moving fifteen years of history.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

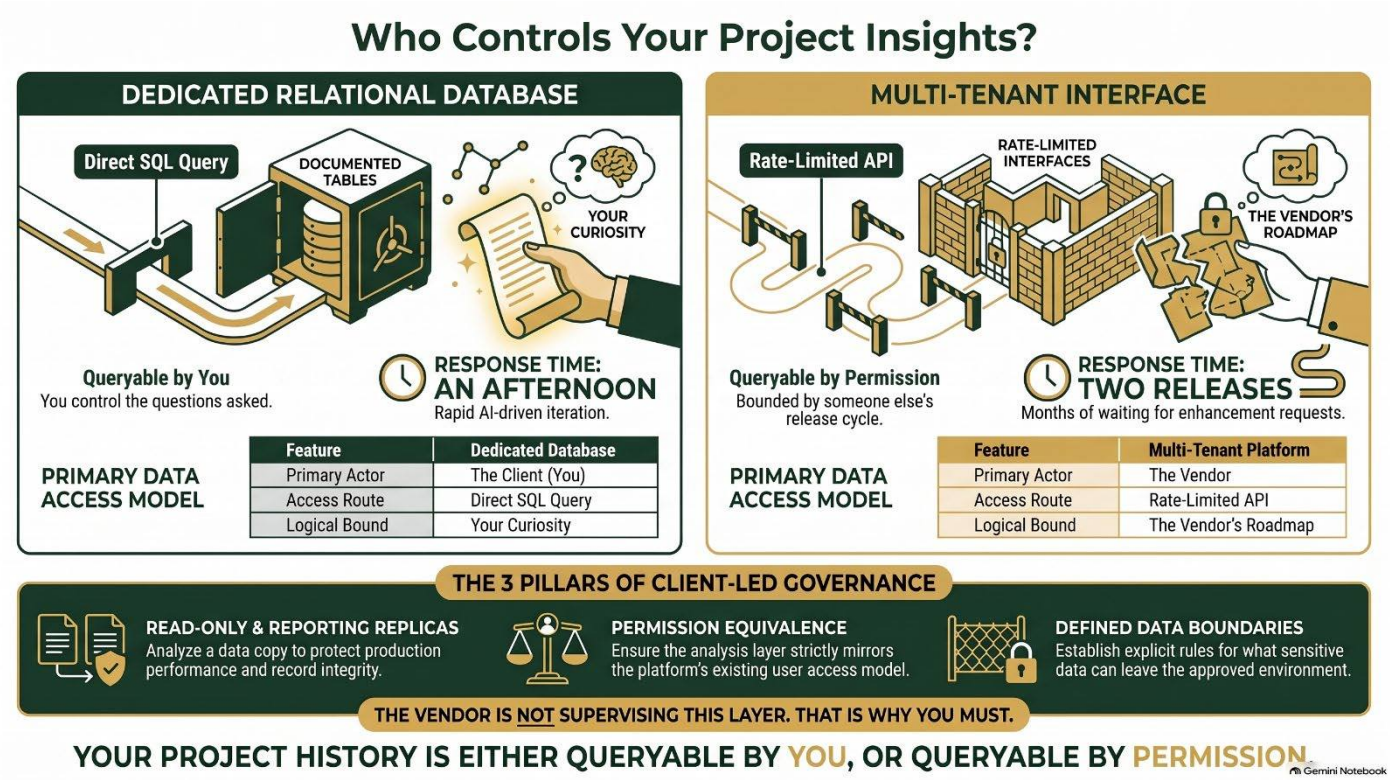
ARTICLE 13 OF 16 — AI FOR PMIS SERIES

The use cases your vendor will never ship

Every PMIS client has had this experience. The question you genuinely need answered is not on the report menu, and nothing on the menu is close enough to adapt.

So you raise an enhancement request. It is scoped, prioritised against everyone else's, scheduled, and arrives two releases later in a form that is not quite what you asked for. By then the commercial situation that prompted the question has resolved itself one way or the other.

The nine articles before this one dealt with capability supplied by a vendor or a partner. This one deals with the part nobody sells you, and it is the part where organisation-specific advantage actually accumulates.



The question you genuinely need answered is not on the report menu, and the enhancement request arrives two releases later in a form that is not quite what you asked for.

An architectural fact with commercial consequences

PMWEB holds project data in a dedicated MS SQL database with a documented relational schema. That sentence sounds like an implementation detail and it is not.

Compare it with the alternative. A multi-tenant platform where every client's data shares infrastructure cannot generally expose that data directly, so it offers an interface instead — an API, usually rate-limited, returning what the vendor decided to expose in the shape the vendor decided to expose it. That is a reasonable engineering response to a real constraint. But the consequence for the client is specific: every question you want answered has to be expressible through somebody else's interface.

A dedicated relational database with a known schema does not impose that. Your commitments, change orders, certifications, correspondence and workflow history are in tables you can read, with relationships you can traverse, in a query language that has been standard for forty years.

Your project history is either queryable by you, or it is queryable by permission. That is an architectural fact, not a feature comparison.

What that makes possible

Point a general AI tool at properly governed structured data and the economics of analysis change completely.

You describe the question in plain language. The tool works against your actual schema, proposes the query, returns the result, and shows you the logic it used. You look at the answer, realise the question was slightly wrong, and refine it. Three iterations later you have something correct. The elapsed time is an afternoon.

What makes this work is not the sophistication of the AI. It is that the data underneath it is structured, validated and related — the argument this series opened with, arriving at its practical conclusion. Value lists mean a category means one thing. Validation at entry means the dates are dates. Defined relationships mean a change order knows which commitment it varies. An AI given that substrate produces a query it can explain, and a query you can check.

An analysis that proves useful stops being an experiment and becomes something the PMO runs monthly. It never went through anyone's roadmap.

Where this must be governed

This capability is powerful precisely because the vendor is not supervising it, which is exactly why the client has to. Four conditions, and they are not optional.

Read-only access. Analysis reads; it does not write. A client-built analysis layer that can modify records has bypassed every control this series has spent twelve articles arguing for.

A reporting replica rather than the production instance. Analytical queries and transactional workloads compete, and the first time an exploratory query degrades performance for two hundred users the capability will be withdrawn and will not come back.

Permission equivalence. This is the one most often missed and it is the most serious. If the PMIS restricts a commercial manager to their own programme, an analysis layer that reads the whole database gives that person a view they were deliberately denied. The analysis layer must not become a route around the access model, which means access to it is governed as carefully as access to the platform.

A defined data boundary. Establish before you start what may leave the organisation's approved environment and what may not. Commercially sensitive positions, personal data and anything under a confidentiality obligation to a counterparty need an explicit position, taken by someone with authority to take it.

These four conditions are the client's version of the questions buyers now put to vendors. PMWEB's own checklist asks whether AI respects the permissions already in place, whether it reads your data or trains on it, whether there is an audit trail a dispute would survive, and who signs. When the organisation builds the layer itself, nobody else answers those questions for it. Permission equivalence answers the first. The data boundary answers the second. A log of the analyses run, the queries behind them and who acted on the result answers the third. And any output that informs a decision should be treated like any other submission: a named person reviews it and owns what follows. AI advises, the operator decides, and in this layer the organisation is both the vendor and the operator.

What to establish first

Before any of this, one practical question needs a specific answer from your provider: what form of access do you actually have to your own data?

A reporting replica you can query directly is the strongest position. A restorable backup you can stand up in your own environment is workable. Scheduled exports are weaker but usable. API-only access means Layer C exists for you in a much reduced form, and you should know that before you plan around it.

Ask the question explicitly, get the answer in writing, and establish it before you need it rather than during the quarter when a board question makes it urgent.

The point of the layer

None of this replaces what the platform ships or what a partner adds. It sits alongside them and it answers a different class of question — the ones specific enough to your organisation that no vendor would ever build them, and valuable enough that you would have paid for them if anyone had offered.

The organisations that will get the most from AI in project delivery over the next five years are not the ones that picked the platform with the longest AI roadmap. They are the ones that can ask their own data their own questions, and the difference between those two groups is being decided by architectural choices being made now.

The organisations that will get the most from AI in project delivery are not the ones that picked the platform with the longest roadmap. They are the ones that can ask their own data their own questions.

Next week: what those questions actually look like, for four different kinds of organisation.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

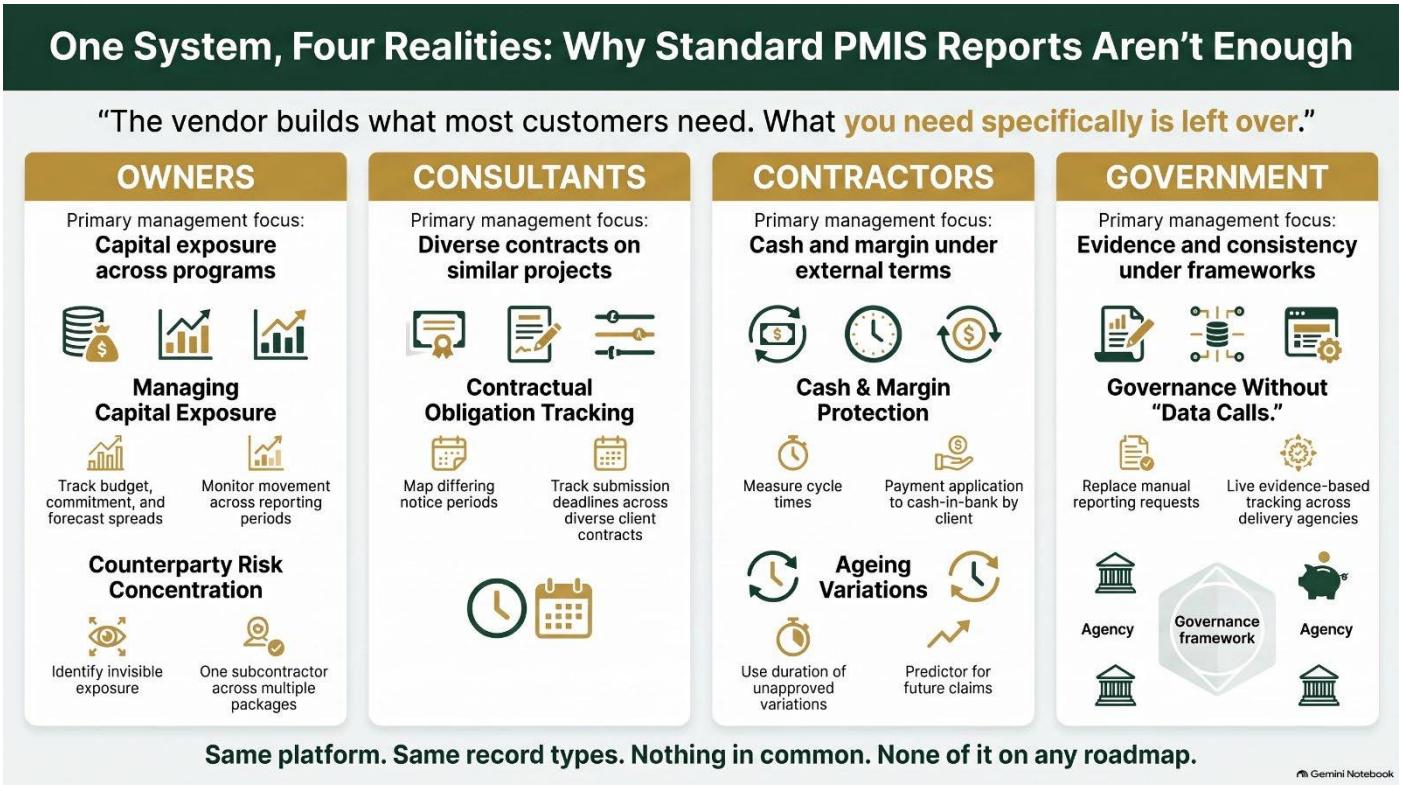
ARTICLE 14 OF 16 — AI FOR PMIS SERIES

Four organisations, four sets of questions

The previous article argued that organisations can build analysis on their own project data rather than waiting for a vendor to ship it. The obvious response is fair enough, but like what?

Here is the useful part. Take the same PMIS, holding the same kinds of record, and put it in four different organisations. The questions each one urgently needs answered have almost nothing in common.

That is precisely why no vendor builds them.



Take the same PMIS, holding the same kinds of record, and put it in four different organisations. The questions each one urgently needs answered have almost nothing in common.

The owner and the developer

An organisation delivering its own assets is managing capital exposure across programmes that are at different stages, and the questions are about position rather than progress.

Where does the portfolio actually stand between approved budget, committed value, certified work and forecast final cost — and how has that spread moved over the last four reporting periods? The individual numbers exist on every dashboard. The movement between them over time is what tells you whether estimates are holding, and it is rarely presented.

How concentrated is exposure by counterparty across the whole portfolio, including where the same entity appears as a subcontractor on packages held by different main contractors? Concentration is invisible at project level by definition, and it is the exposure that matters when a counterparty gets into difficulty.

Where is uncommitted contingency actually sitting, by programme and by phase, and how does that compare with the risk still open on each? Contingency is usually reported as a balance and is only meaningful against the exposure it exists to cover.

For a developer selling the asset, how do handover and release milestones line up against actual completion status — which is a question that crosses from the PMIS into the sales pipeline and is never answered by either system alone.

The PMC and the consultant

A consultancy delivering for multiple clients has a different problem. The projects are similar; the contracts governing them are not.

What obligations fall due in the next sixty days across every contract, given that each client's terms differ? Notice periods, submission deadlines, reporting obligations and approval windows vary by contract, and the standard answer is a spreadsheet maintained by whoever is most diligent.

How long does correspondence sit before it is answered, by project and by counterparty? Response latency is one of the earliest available indicators of a project heading toward claims, and it is measurable from the correspondence record long before anyone uses the word.

Which projects generate disproportionate variation volume relative to their value, and is that a client characteristic or a project characteristic? A firm that knows this prices differently and staffs differently.

What is the deliverable load across the portfolio over the next quarter, and where do the peaks collide? Resource planning across concurrent client commitments is a scheduling problem the scheduling tool does not solve, because it lives across projects rather than within one.

The contractor

A contractor is managing cash and margin under terms somebody else wrote, and the questions follow the money.

What is the elapsed time from payment application to certification to payment, broken down by client, and how has it moved? Cycle time by counterparty is a cash forecasting input and a commercial negotiation input, and most contractors have a general sense of it rather than a measurement.

Where does margin erode between the approved position, the certified position and the paid position, by project and by work package? The erosion is usually discovered at final account, when the causes are historic.

What is the concentration in the supply chain, and what does performance history say about the subcontractors carrying the most exposure? This is the same counterparty question the owner asks, one tier down.

What variations are open, at what value, and how long have they been open? Ageing of unapproved variations is a direct predictor of the claims position and it is derivable from records that already exist.

The government entity and the PIF company

An organisation delivering under a national framework has obligations the other three do not, and its questions are about evidence and consistency.

Where does the portfolio stand against the framework's reporting requirements right now, without a data call — because the data call is the problem. The ability to answer at any moment from the system of record, rather than by asking twelve delivery entities to submit, is the difference between governance and administration.

How consistently is practice applied across delivery agencies? When several entities run the same processes, variation in how they are actually used is invisible unless somebody measures it, and it is the leading indicator of where assurance should be directed.

Can an audit question arriving with no notice be answered from evidence, in the same week? The question is rarely about whether the answer exists. It is about whether it can be assembled before the meeting.

And how does portfolio progress map against national programme targets, at a level of aggregation that means something to a board and can still be drilled into?

The pattern

Look across the four and one thing stands out. Not one of these is a feature request. Every one of them is answerable from data the organisation already holds, in a system it already runs, and none of them appears on any vendor's roadmap because they are specific to how each organisation is structured and what it is accountable for.

> **None of these is a feature request. All of them are answerable from data you already hold.**

That is the case for the client-built layer, stated concretely. The vendor builds what most customers need. The partner builds what the category needs. What your organisation needs, specifically, is left over — and it is usually the part that would change a decision.

Where to start

Pick the question your organisation has repeatedly wanted answered and repeatedly not answered, because the effort was never justified for a single occasion. That question is almost always the right first one, because the fact that it keeps recurring is evidence that it matters and evidence that nobody has solved it.

Answer it once, properly, with the governance conditions from the previous article in place. If the answer changes a decision, it becomes a monthly output. If it does not, you have spent an afternoon finding that out.

> **Pick the question your organisation has repeatedly wanted answered and repeatedly not answered, because the effort was never justified for a single occasion. That question is almost always the right first one.**

Next week: the objection that stops more PMIS decisions than price, and why it rests on a false analogy.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

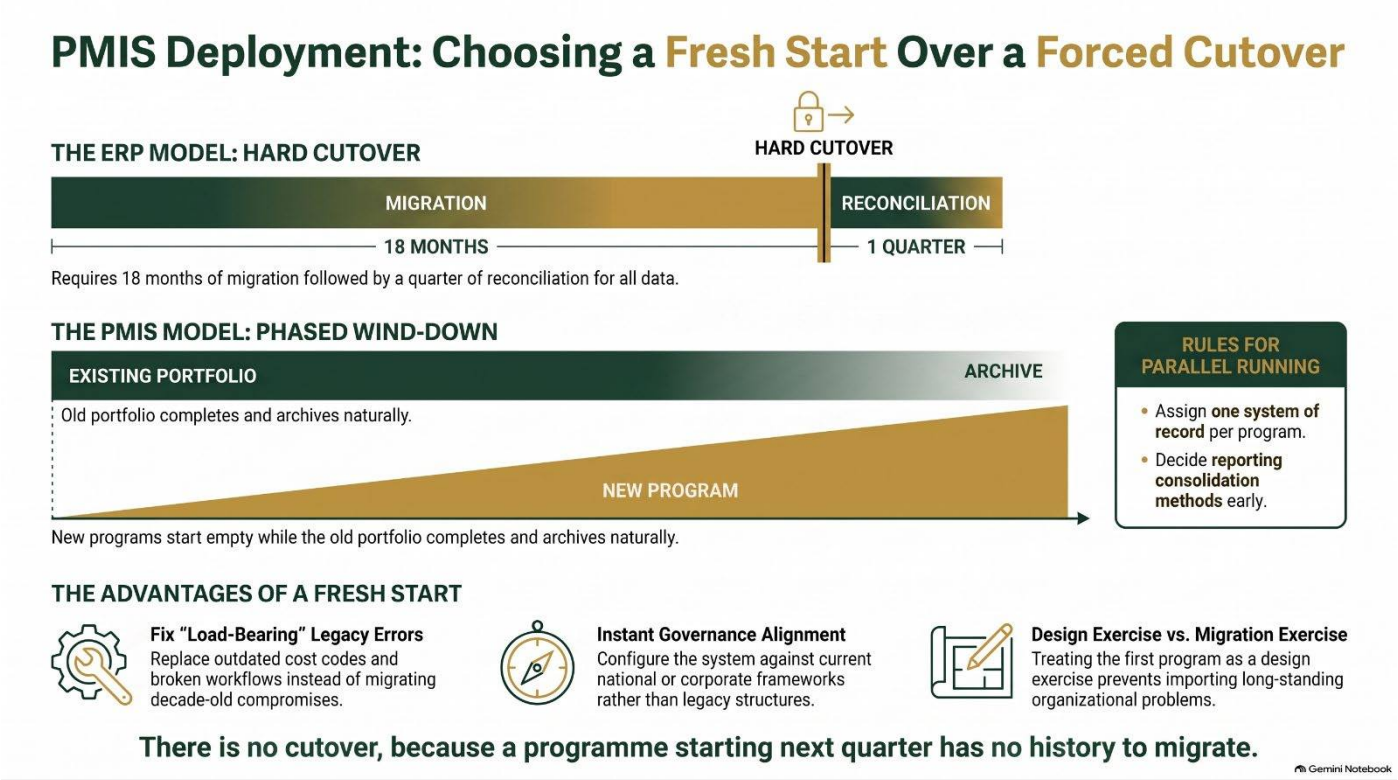
ARTICLE 15 OF 16 — AI FOR PMIS SERIES

You do not need a migration to start

In fourteen years of these conversations, the objection that stops a PMIS decision is almost never price. It is this:

We have fifteen years of history in the current system. We cannot move it. So we cannot move.

The objection is stated as a fact and treated as one. It is actually an assumption, and it rests on a comparison with a different kind of system entirely.



The objection is stated as a fact and treated as one. It is actually an assumption, and it rests on a comparison with a different kind of system entirely.

The analogy that does the damage

Most senior people forming this view have lived through an ERP implementation, and they are reasoning from it correctly.

An ERP is an organisational system. It holds the general ledger, the chart of accounts, the vendor master, payroll. There is one instance of the truth and everything depends on it, so replacing it means a cutover — a date on which everything moves, preceded by eighteen months of migration and followed by a quarter of reconciliation. Anyone who has done that once is right to be cautious about doing it again.

A PMIS is not an organisational system. It is scoped to projects and programmes.

That single structural difference changes everything about how you adopt one, and it is routinely missed because both are large systems bought by the same committee.

What programme scope actually means

A programme starting next quarter has no history to migrate. There are no legacy commitments on it, no prior correspondence, no closed variations. It begins empty whatever system you put it in.

So the sequence available to you is this. Deploy on the new programme. Let the existing portfolio complete on the existing system. Run both for as long as the old portfolio takes to finish, which is a defined and finite period because projects end.

There is no cutover date. There is no weekend when everything moves. There is no reconciliation quarter. The old system winds down as its projects close, and at some point it is holding nothing that anyone needs to open, at which point it becomes an archive.

> **An ERP has to be cut over because everything depends on it. A PMIS does not, because projects end.**

The opportunity most organisations miss

Here is the part that is genuinely valuable and rarely discussed, and it argues for a fresh start on its own merits rather than as a compromise.

Every mature PMIS configuration contains decisions everyone knows are wrong. A cost code structure inherited from a system two generations back. A WBS that made sense for the first programme and has been stretched across four since. Value lists with entries nobody uses and entries everybody needs that were never added. Workflows routing approvals to a role that was reorganised out of existence in 2021, now handled by an exception everybody knows about.

None of it can be fixed while the system is live, because changing a coding structure with eleven years of history behind it means restating the history, and nobody has the appetite. So the wrong decisions are load-bearing, and they get carried forward into the next system by anyone who migrates configuration along with data.

A new programme on a new platform is the only realistic opportunity to fix them. You configure against the delivery framework you actually intend to follow — which for many organisations in this region now means a national or corporate governance framework that did not exist when the current system was set up — rather than against the accumulated compromises of the last decade.

Organisations that treat the first programme as a migration exercise import their problems. Organisations that treat it as a design exercise do not.

The parallel period, run properly

Two practical points, because a parallel run done carelessly creates its own problems.

Be explicit about which system is authoritative for what. Every record type has exactly one system of record during the parallel period, determined by which programme it belongs to. Ambiguity here is what produces the double-entry that gives parallel running a bad reputation.

And decide reporting consolidation early. During the overlap, portfolio-level reporting has to draw from both, and the organisations that handle this well decide up front whether consolidation happens by exporting from the old system into the new, or by reporting at a level of aggregation both can support. Deciding it in month seven, under pressure from a board pack, produces a manual spreadsheet that then becomes permanent.

What this does to the decision

Reframed this way, the decision in front of an organisation is not whether to undertake a migration. It is whether the next programme starts on a system that reflects how they now intend to work.

That is a considerably smaller decision, it can be taken by people with authority to take it, and it does not require anybody to sign off on moving fifteen years of history before seeing whether the platform delivers.

Which is not to say that the history is worthless or unreachable. There are good reasons to bring some of it forward, and there is now a genuinely better way to do it.

> **The decision in front of you is not whether to undertake a migration. It is whether the next programme starts on a system that reflects how you now intend to work.**

Next week, closing this series: where AI actually reduces the cost of migrating legacy project data, and the three things it will not do for you.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

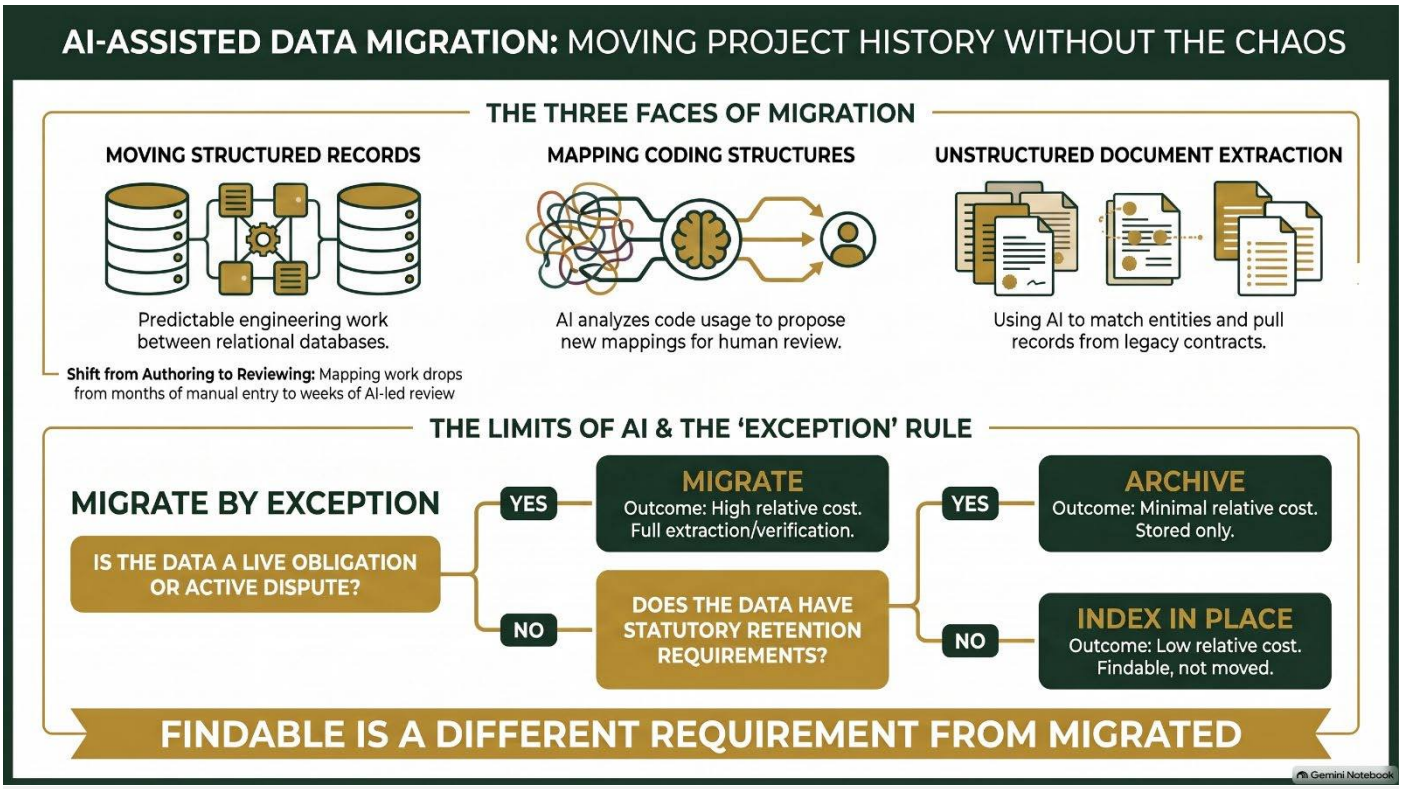
CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.

ARTICLE 16 OF 16 — AI FOR PMIS SERIES

AI-assisted migration of legacy project data — and what it will not do for you

The previous article argued that you can start a PMIS without moving your history. This one deals with the history — because some of it does need to come forward, and the cost of bringing it has changed.

It has changed less than the marketing around AI suggests, and more than most people implementing systems assume. Both halves of that matter.



Most migration projects go wrong because they are scoped as one activity. They are three, with different costs and different failure modes.

Three problems that get called migration

Most migration projects go wrong because they are scoped as one activity. They are three, with different costs and different failure modes.

Moving structured records from one database to another. Mapping one organisation's coding structures onto another's. And extracting governed records from decades of unstructured documents.

Conflating them produces a single estimate that is wrong in both directions — too high for the first, far too low for the third — and a plan in which the hardest work is discovered late.

The first is engineering, and it is not where the cost sits

Moving structured records between relational databases is well-understood work with mature tooling. Staged loading, referential validation, reconciliation against source totals, iterative correction. It is not trivial and it is entirely predictable.

It is also considerably easier when the target is a dedicated MS SQL database with a documented schema, for the same reason that mattered in the earlier article on client-built analysis. You can load in stages, validate relationships before committing, reconcile against control totals, and correct in place. The tooling is standard and the people who know how to use it are not scarce.

Compare that with a target reachable only through an API. The load becomes a sequence of calls, subject to rate limits and to whatever the interface exposes, with error handling and restart logic built around the constraint. Same data, same destination, materially different exercise. Not impossible — genuinely more expensive, and slower to correct when something is wrong.

The second is where AI helps most, and it is not what people expect

The expensive problem in most migrations is not moving records. It is deciding what the old codes mean in the new structure.

A legacy WBS built for one delivery model has to map onto a new one. A cost code structure with eleven years of accumulated exceptions has to be expressed in the target's coding. A vendor master with duplicates, historic names, merged entities and three spellings of the same company has to be reconciled. This work has traditionally been done by two or three experienced people over months, and it is genuinely difficult because it requires knowing both what the old codes meant in practice and what the new structure is for.

AI is unexpectedly good at proposing these mappings — reading both structures, examining how codes were actually used rather than how they were defined, and proposing correspondences with the reasoning shown so a person can accept or reject each one.

The work does not disappear. It changes from authoring to reviewing, and that is roughly an order of magnitude. Two people producing a mapping over four months becomes two people reviewing a proposed mapping over three weeks, with better coverage because the proposal examined every code rather than the ones anybody remembered.

> **The mapping work changes from authoring to reviewing. That is the largest single saving available in a migration, and it is the least discussed.**

The third is where AI helps most visibly

Extracting governed records from legacy correspondence, contracts and submittals is the capability described earlier in this series, applied backwards in time. The same properties matter: entity matching so the extracted content attaches to the right record, and source traceability so it can be verified by exception rather than in full.

And the same scoping discipline applies. Classify everything, extract from the fraction that warrants it. An archive project scoped as extraction of the whole estate does not finish.

What it will not do

Three things, stated without hedging, because the honest version of this argument is more persuasive than the complete one.

It will not decide what is worth migrating. This is the decision where most migration budgets are actually lost, it requires commercial and legal judgement about which historic positions still carry consequence, and no tool will make it for you. Organisations that skip this decision default to everything, which is how a migration becomes a programme.

It will not validate its own output. A proposed mapping is a proposal. Extracted values are candidates. The verification step is real work performed by people who understand the material, and a migration plan without a resourced verification stage has moved the risk rather than removed it.

This is the principle PMWEB's Intelligence Control Framework states for the platform, AI advises and the operator decides, and it applies to migration with particular force, because a mapping accepted in month two governs every record loaded after it. Keep a record of what the AI proposed, who accepted or rejected each proposal, and on what basis. Whoever later has to explain why a historic cost landed in the code it did will read it.

And it will not reconstruct a record that was never properly kept. No extraction recovers an approval given verbally in a site meeting in 2019. No mapping reconstructs the intent behind a cost code that three people used three different ways. If the history is incomplete, migration produces a faithful copy of an incomplete history, and the gaps are the same gaps.

Migrate by exception

The recommendation this all leads to is simple to state and difficult to hold to.

Move the open commitments, the live obligations, the disputed history, and anything that will be needed as evidence in a matter that is still running. That is a small fraction of most archives and it is the fraction that has consequence.

Classify and index the rest in place, so it is findable when somebody needs it, which will be rarely and unpredictably. Findable is a different requirement from migrated, it costs a fraction as much, and for closed projects it is almost always sufficient.

Archive what remains against your retention obligations and stop paying attention to it.

Organisations that migrate by exception finish, usually within a period a sponsor will tolerate. Organisations that attempt completeness produce a very good plan, a series of extensions, and eventually a decision to stop that gets framed as a change of priorities.

Where this series ends

Sixteen articles ago this began with a claim: that the quality of any AI answer in capital delivery is set by the data underneath it, and that a PMIS and a document repository are not the same substrate.

Everything since has been a consequence of that. The capabilities arriving in PMWEB 2026 are aimed at getting more of the record into the system. The analysis worth wanting has to name the data it read. The monitoring worth having connects an event to records that were structured enough to be connected. The client-built layer exists because structured data can be queried by the organisation that owns it. And migration is tractable because the target is a documented relational structure rather than a black box.

And the governance argued for at the start of the series (permissions inherited rather than duplicated, an audit trail that names a person, a gate that does not move) is now something PMWEB has written down as a published framework under a five-word principle: AI advises, the operator decides.

The organisations that will get the most from AI in project delivery are not the ones with the longest vendor roadmap. They are the ones whose project record is complete enough, structured enough and accessible enough to be worth reading.

That was true before any of this and it will be true after. AI has simply raised the return on having done it properly.

The organisations that will get the most from AI in project delivery are the ones whose project record is complete enough, structured enough and accessible enough to be worth reading. AI has simply raised the return on having done it properly.

Thank you for reading this series. I am glad to discuss any of it — particularly the questions in articles four and twelve, which are the ones worth taking into your next vendor conversation whichever platform you are running.

About the Author

Bassam Samman, PMP, PSP, EVP, GPM is the Chairman and Chief Executive Officer of CMCS FZ-LLC, the PMWEB Authorised Implementation Partner for the Gulf Cooperation Council and broader Middle East and North Africa region. CMCS holds the PMWEB Partner of the Year 2025 designation and serves capital project owners, contractors and consultants across the UAE, Saudi Arabia, Qatar, Kuwait and India.

Bassam has over four decades of experience providing project management, project controls services and project management information systems to more than two hundred capital projects with a combined value exceeding one hundred billion US dollars. He is the author of four published books on project management digitalisation and has lectured at more than three hundred and fifty international events. He holds a Master's degree in Engineering Administration (Construction Management), with Faculty Commendation, from George Washington University and a Bachelor's degree in Civil Engineering from Kuwait University, with executive education at Harvard Business School and London Business School.

About PMWEB

PMWEB is an enterprise-class, web-based capital construction project management platform trusted by organisations worldwide including Harvard University, the United Nations, and Atlanta Airport. Built on Microsoft .NET technology, PMWEB manages the entire project lifecycle within a single integrated platform — from portfolio planning and cost management through document control, procurement, scheduling, asset management, and business intelligence. With 150+ pre-configured SSRS reports, drag-and-drop dashboard design, multi-currency support, and open REST/SOAP API integration, PMWEB delivers enterprise capability with implementation timelines measured in months, not years.

About CMCS

CMCS (Collaboration, Management and Control Solutions) is the regional PMWEB implementation partner for the GCC and MENA regions, awarded PMWEB Partner of the Year 2025. CMCS specialises in digital project management solutions for real estate developers, infrastructure owners, and construction companies across the Middle East and North Africa.