EMASES

Strategic Thesis

Vision for B2B Platforms in the AI Era

Our worldview, what we have built, what we are researching, and the commitment that binds it together.

Author
Vinh, Founder, EMASES

Version
2.2, September 2026

Foreword. Who This Is For

This document is written for finance, valuation, and operations leaders evaluating a long-term platform partner. If that is you, it should take about ten minutes and leave you knowing exactly how we think, what we have actually built, and what we have not yet proven.

EMASES is a young company. We do not have decades of references or a recognisable brand inside every procurement department. What we have is a clear thesis about where the AI economy is heading and why governed, vertical platforms win in it, a working studio, and an unusual habit: being precise about the difference between what is proven and what is still direction. The document is organised exactly that way. Worldview first (Parts 1 to 3), then the Studio: what exists today (Part 4), then the Lab: what we are researching (Part 5), then the one rule that governs everything we publish (Part 6).

Part 1. The Gap, and the AI-Spray Trap

In 2026, the largest hyperscalers are spending roughly $660 to $720 billion on infrastructure, most of it for AI. Against that, only 29% of organisations report significant returns from generative AI, and just 11% of enterprises piloting AI agents have moved any into production. The infrastructure layer is racing ahead; the value-delivery layer is stuck. The next phase of value capture belongs to whoever bridges that gap, and it will not be the model providers or the hyperscalers. It is bridged at the application layer, by vertical, governed platforms operated by domain experts. That is the category EMASES operates in.

Why is the value layer stuck? Because most enterprises, under board pressure to show AI progress on quarterly timelines, are deploying agents and copilots on top of systems that were never designed for them, without first addressing decades of accumulated data fragmentation and governance debt. We call the pattern AI-spray: bolting AI onto a broken foundation. The evidence is consistent. Around 70% of organisations discover their data infrastructure is inadequate only after launching ambitious AI initiatives. Over half name data quality as their primary obstacle. Gartner expects more than 40% of agentic AI projects to be cancelled by 2027.

AI debt is the cost of every agent deployed on a foundation that was not designed for it. It is being accumulated silently across most large enterprises right now. The bill will come due.

Two further facts shape our strategy. First, vendor-led AI deployments succeed roughly twice as often as internal builds: enterprises are good at operating software, not at building it, and boards are redirecting budgets accordingly. Second, the same research that documents the failures shows the counter-pattern succeeding: domain-specific agents on clean, vertical foundations are growing over 60% annually and consistently outperforming general-purpose agents in measurable impact. The market is bifurcating, and the bifurcation favours exactly the architecture we build.

Europe sharpens the case. EU enterprise AI adoption is roughly a quarter of US levels, so the headroom is still in front of us. The EU AI Act became fully applicable in August 2026, which rewards governed platforms in regulated domains. And European buyers are structurally more careful, which makes "vendor-led, governed, deployed in your own environment" the winning motion here, not "move fast and break things".

Part 2. How We Build

2.1 In your own Microsoft tenant

In B2B, the buyer is not a technology enthusiast. The buyer is a CFO or a compliance officer whose career risk is asymmetric: a successful AI project earns modest credit; a failed one with a data breach is career-ending. Their dominant decision criterion is trust, not capability. So every EMASES product deploys into the client's own Microsoft tenant. You own the environment, you own the data, and your existing Microsoft governance, identity, and audit apply unchanged. That is a structurally different conversation than handing financial data to one more SaaS cloud, and it is why we chose Microsoft as our foundation: it is the substrate our clients already run on, already trust, and already pay for.

2.2 Architecture before features

The dominant incumbents in enterprise software cannot rebuild their core architecture: their data models carry twenty years of commitments, and any deep change breaks thousands of integrations. Their architectural debt has become their business model. We start fresh, and we spend the freshness deliberately: on every project, the data model and governance architecture are designed before any user-facing feature, because the data model is the only part of a system that is genuinely expensive to change later.

We do not pursue a perfect, future-proof architecture; no architect anticipated the automobile, and Paris survived it because its structure could be restructured. Our platforms are built so the data model, the business logic, and the interface can each be replaced independently as the world changes. That is what lets us make a commitment larger vendors cannot honestly make: we will still be useful to you in ten years, even though nobody can predict what you will need in ten years.

The same discipline is what AI-ready means in practice. The data is structured, the schema is consistent, permissions are explicit, the audit surface is built in. When AI lands on a foundation like that, it lands at specific friction points and it works. It also stays cheap and predictable: because the platform computes every figure before the AI writes a word, an AI output is one small call with a near-fixed cost, not an agent swimming through raw data with an unpredictable bill. Cost per outcome is a consequence of architecture, not of model choice.

2.3 One architect, AI execution, and the question you are already asking

EMASES products are designed end-to-end by a single domain-expert architect, with execution delegated to AI workers under a governed build discipline: every change planned, reviewed, and recorded in an append-only build record. Three years ago this scope required a funded engineering team; today it does not, and the client-side effects are concrete: faster delivery, one coherent architecture instead of fifty weakly-connected modules, and a cost structure incumbents cannot match.

A careful buyer's next question is the right one: what happens if that one person is unavailable? The model is built to answer it. Products ship as managed solutions running inside your tenant, so daily operation does not depend on EMASES being online. The platform is documented to be maintainable rather than memorised, and the append-only build record means the state of every component is reconstructible by a third party. Your data is yours and fully exportable at any time, including at the end of the partnership. And the commercial model funds continuity: recurring revenue pays for a permanent development stream, not a one-off project. We would rather name this risk and show how it is handled than pretend it does not exist.

Part 3. Why We Will Still Be Here

The question a client should ask any young vendor is not "how big are you" but "why will you still be here, and still be better, in ten years". Our answer rests on three things.

- Domain depth that cannot be hired. Years of operational experience inside regulated valuation and finance workflows, gained where the numbers actually go wrong. Competitors can hire developers and copy screens; operational scar tissue compounds with every deployment and cannot be bought.
- Scarce technical ground. Production-depth expertise in Microsoft Power Pages is held by at most a few hundred practitioners globally. We stand on ground few can follow us onto, in both its traditional and modern forms.
- The development stream itself. The value of an EMASES platform is not frozen in any deployed version; it lives in the continuous stream of research, upgrades, and methodology development that flows to clients for as long as the partnership runs. A static copy of any product decays; the stream compounds. This is why our commercial model is recurring, and why the Lab in Part 5 is a permanent function, not a marketing page.

WHAT IS PROVEN

Part 4. The Studio: What We Build and How We Deliver

Everything in this part is demonstrably true at the time of writing. Where a capability is still in validation, we say so.

4.1 How we deliver

- In your tenant, as managed solutions. You own the environment and the data; the proprietary methodology layer remains with EMASES. Data residency, sovereignty, and audit are handled by your existing Microsoft governance.
- Architecture first. Data model and governance designed before any feature, on every project.
- AI as workers, domain expert as architect, under a governed build discipline with append-only records of every change.
- Flexibility at configuration time, rigidity at run time. You shape the platform to your structures during setup; once live, it enforces the governance it was configured with.

4.2 Hotel Val Hub

A governed hotel valuation operating platform on Microsoft Power Pages and Dataverse. It connects property and historical P&L evidence, Jobs and Job Assets, portfolio valuation, projections, DCF, CapEx, Market Performance and asset benchmarking, valuation history and comparison, Senior Review, work queues, imports, exports, and permissions in one governed workflow. The operating principle is a clean split: the platform handles repeatable mechanics and governed computation; valuers retain judgement and signature.

Status, stated plainly: Hotel Val Hub is built and in external validation with senior practitioners from the hotel valuation industry. It is not yet commercialised. We would rather show you a working platform and an honest status than a polished claim.

4.3 YOYP Budget & Expenses Hub

A governed budget and expense platform for finance and FP&A teams, metadata-driven on Dataverse, running natively inside your Microsoft tenant. A finance admin reconfigures structure (items, centers, formulas) in minutes through the admin UI, not through a consultant statement of work. It shares the same extended P&L framework that underpins Hotel Val Hub: one proprietary framework, multiple products standing on it.

4.4 The commercial model

Pricing is anchored to the work our platforms carry, not to seat counts: for valuation, it follows jobs and assets, because the platform replaces work budget, not software budget. Revenue is recurring, because the value delivered is a continuous stream. Support is packaged and defined, never open-ended. And you always retain full ownership and export rights over your data, including at the end of a partnership. We consider that last point a selling feature, not a concession.

Part 5. The Lab: What We Are Researching

EMASES Lab is our R&D function. It exists because a client who buys an EMASES platform is buying into a development stream, and this is where that stream is visible. Everything in this part is direction, not capability: nothing here is sold, promised, or dated.

- AI-assisted P&L intake and mapping. The next product layer in Hotel Val Hub is a human-reviewed intake flow for client P&Ls. The intended pattern is to classify the source, propose mappings into the governed P&L structure, flag ambiguity, and require valuer approval before persistence. This is planned direction, not a current capability.
- The field-to-office valuation pipeline. A valuer receives a data room; platform agents convert the material into the platform's structures. The valuer inspects the property, capturing notes and voice; agents transpose them onto the same governed data model. Back at the office, the job is prepared and waiting for validation, and the report is ready the moment the job is locked. The professional's day concentrates on the two things only a professional can do: judgement and signature.
- From hotels to operational real estate. The extended P&L framework under Hotel Val Hub was designed from the start to describe operational real estate generally. The research direction is a valuation platform covering further operational asset classes on the same governed core.

Part 6. The Commitment: No Claim Before Evidence

Everything EMASES publishes follows one rule: no claim before evidence. What we present as capability is demonstrably true at the time of writing. What is not yet proven is labelled as direction, as Part 5 is, and when something changes status, the label changes with it, in both directions. Our clients are professionals whose signatures carry legal weight; they extend trust carefully and withdraw it permanently. We would rather show less and be believed than show more and be doubted.

No claim before evidence. Not a slogan. The operating rule for every sentence we publish.

Closing

Most enterprise software is sold as a defensive purchase: a way of not getting fired for choosing the safe vendor. We think the moment is ripe to sell something else: an offensive purchase, a way of building real operational advantage against competitors still anchored to legacy systems and accumulating AI debt on top of them. The clients who recognise that opportunity are the ones we want to work with.

If our worldview resonates with how you think about your own business, we should talk.

EMASES • Microsoft-Native B2B Platforms • emases.eu • hotelvalhub.emases.eu