



GOYAVE CONSULTING

METHODOLOGY & PORTFOLIO

HUMAN-FIRST. END-TO-END. KNOWLEDGE-FIRST.

The method, and the five offers it runs in.

Jérôme Fenyo, independent AI-First consultant.

Paris, working across Europe.

EDITION 2026

GOYAVECONSULTING.COM



Scale is where AI programs stop. 95% of generative AI pilots deliver no measurable P&L impact (MIT NANDA, 2025), and fewer than 10% of deployed use cases pass the pilot stage (McKinsey, 2025). A use case can succeed and the organization stay exactly as it was.

Two conditions make a transformation hold at company scale: a way of working, and somewhere the knowledge accumulates. Then it has to land in three places: on the leaders who decide, on the business functions that produce the value, and on the application estate that keeps everything running.

WHAT THE WAY OF WORKING RESTS ON

01 HUMAN-FIRST THE TARGET IS THE CRAFT. NOT THE PROCESS. The subject is the judgement of experts and leaders: deciding, weighing trade-offs, negotiating, designing.	02 END-TO-END THE WHOLE VALUE CHAIN. FROM DAY ONE. Design the full scope immediately. Deliver it in slices.	03 KNOWLEDGE-FIRST THE APPLICATION IS REPLACEABLE. THE KNOWLEDGE IS NOT. Capture the business knowledge while building. The software can then be rebuilt in days.
---	--	--

Together they are the frame: the minimal structure an organization puts in place so that an AI transformation reaches production and holds there. Goyave Consulting covers the five, one offer each: AI-First, Craft Twin, the AI Leadership Compass (AIRC), the Agentic Sales Engineering Studio (ASES), and Intelligent Application Management (IAM). This document sets them out chapter by chapter, with the diagrams that carry them.



**THE REAL DELIVERABLE OF AN AI TRANSFORMATION
IS CODIFIED BUSINESS KNOWLEDGE.**

AUTHOR
Jérôme Fenyo
Founder & AI-First
Consultant
Goyave Consulting

CONTACT
jerome.fenyo@goyave-
consulting.com
goyaveconsulting.com
linkedin.com/in/fenyo

EDITION
2026
English
48 pages

CONTENTS

00	THE PRACTICE The argument, the five offers, and how a project runs. Three convictions, one method · Five offers, one frame · What the five do together · From immersion to a digital twin of the craft	4
01	AI-FIRST The method: where to aim, at what scope, and what to keep. None of the three works alone · Build and make it react · No baseline, no provable ROI · Ten points to settle · Making it a durable capability	11
02	CRAFT TWIN A second brain for a team, and the twin of a craft. A memory of the craft · Three layers, already built · Four depths · What an engagement does · What you keep	21
03	AILC · AI LEADERSHIP COMPASS Two days for an executive to take ownership of AI in their own role. Think, then do · Four assets · How a cohort runs	29
04	ASES · AGENTIC SALES ENGINEERING STUDIO An agentic studio for complex B2B deals, €5M to €500M and above. The full bid lifecycle · Agents under rules, humans at the gates · Where strategic deals are lost · How it fits together · A studio, not a sales tool	33
05	iam · INTELLIGENT APPLICATION MANAGEMENT Application management moving from break-fix to prediction and autonomy. Three shifts · From break-fix to autonomy · Seven use cases · Layer by layer, not all at once · Five categories of risk	41
06	ABOUT Twenty-five years of business lines and large deals, and what I do with them now. How I work	47



CHAPTER 00

THE PRACTICE

Jérôme Fenyo - independent AI-First consultant, Goyave Consulting. I transform the craft of leaders and experts, and build the systems myself, in their real environment. For large organizations.

01

Three convictions, one method

02

Five offers, one frame

03

How a project runs

THE ARGUMENT

THREE CONVICTIONS, ONE METHOD.

HUMAN-FIRST


Transformation applies to what experts and leaders actually do: decide, weigh trade-offs, negotiate, design, convince. That craft rests on tacit knowledge, which rule-based automation never reached. Automation still happens - as a consequence, not as the objective.

END-TO-END

Value sits in the connections between steps, not in the fragments. Splitting a transformation into isolated use cases optimizes fragments and produces pilots that never scale. The full scope is designed immediately; delivery stays progressive.

KNOWLEDGE-FIRST

The durable deliverable is the business knowledge captured and structured while building. AI collapses the cost of capturing it: no forms, no taxonomy, it is picked up in the flow of the work. Once that knowledge exists, the software can be rebuilt in days.



None of the three is new on its own. What is written down, and applied, is how they work together.

NOTES & OBJECTIONS

Generic copilots help novices more than seasoned experts. So the work is to capture how the best operate, then spread it across the team. This is not big bang. Scope is total from the start, value is delivered in slices along the chain. Captured knowledge only becomes an asset once it is grounded in verifiable sources and validated by the expert. Otherwise it is a hidden liability. Tacit knowledge is now argued as a competitive advantage in its own right (California Management Review, March 2026).

THE THREE TOGETHER PRODUCE

A DIGITAL TWIN OF THE CRAFT.

THE PORTFOLIO

FIVE OFFERS. ONE FRAME.

APPROACH & FOUNDATION

01

AI-FIRST

THE APPROACH

The craft of experts and leaders as the target, the whole value chain from day one, and captured knowledge as the deliverable that lasts.

DELIVERY

How a project runs: immersion rather than up-front scoping, prototype and pilot on one cadence, the harness from day one.

In active use.

02

CRAFT TWIN

THE FOUNDATION
UNDERNEATH

A foundation that turns what a team already produces - meetings, mail, chats, documents, data - into a memory its experts can query, and then into a twin of their craft.

In use on my own engagements.

LEADERS · FUNCTIONS · ESTATE

03

AILC

AI LEADERSHIP
COMPASS · EXECUTIVE
AI COACHING

Two days for senior executives to take ownership of AI. Day one, they write their doctrine. Day two, they build and govern a personal AI working system. Program running.

04

ASES

AGENTIC SALES
ENGINEERING STUDIO ·
AGENTIC B2B SALES

An agentic studio for complex B2B deals, €5M to €500M and above: the full bid lifecycle, human validation at every gate, a structured deal memory. Currently being deployed.

05

iAM

INTELLIGENT
APPLICATION
MANAGEMENT

Application management moving from break-fix to prediction and autonomy: ticket automation, AIOps and self-healing, knowledge that stays alive.

White paper published, 2025.


Each offer has its own chapter. The order is the order they are read best in: what makes a transformation hold, then where it lands.

WHAT THE FIVE DO TOGETHER

TWO CONDITIONS. THREE PLACES.

Scale is where AI programs stop. 95% of generative AI pilots deliver no measurable P&L impact (MIT NANDA, 2025). Fewer than 10% of deployed use cases pass the pilot stage, and more than 80% of companies report no tangible effect on enterprise EBIT (McKinsey, 2025). Three measurements of one fact: a use case can succeed and the organization stay exactly as it was.

The question is what has to be true for a transformation to hold at company scale. Two conditions make it hold anywhere: a way of working, and somewhere the knowledge accumulates. Then it has to land in three places: on the leaders who decide, on the business functions that produce the value, and on the application estate that keeps everything running. Goyave Consulting covers the five, one offer each. Together they are the frame: the minimal structure an organization puts in place so that an AI transformation reaches production and holds there.



Approach and foundation. Leaders, functions, estate. One offer each.

THE FIVE, ONE BY ONE

AI-First is the way of working, and it comes in two halves. The convictions say what to aim at. Human-first: the target is the craft of experts and leaders, not the process. End-to-end: the whole value chain from day one, because value sits in the connections between steps. Knowledge-first: what is kept at the end is the business knowledge captured while building. The method says how a program runs: immersion rather than up-front scoping, a prototype and a pilot on one cadence, the harness from day one, integration treated as the critical path. Scope is total from the start, delivery stays in slices. None of the three convictions is new on its own. What is mine is not the idea: it is executing all three together, and writing down how.

Craft Twin is where the knowledge accumulates. At scale the binding constraint is retention: what a program learns leaves with the people who learned it. Craft Twin turns what a team already produces - meetings, mail, chats, documents, data - into a memory its experts can query, sourced and dated, and at its deepest level a twin of how a role decides. Four depths, you stop at yours. The application on top is replaceable, the knowledge is not. The three offers that follow plug into it, each at a different depth.

AILC lands on the leaders. At scale, the people who arbitrate decide whether the rest holds. Two

days: day one they write their own doctrine, day two they build and govern a personal AI working system they keep. The rule holds on both days: the executive decides, AI executes.


ASES lands on a business function. The effects that count at enterprise level appear when the craft itself changes. ASES is the first studio: specialized agents under written, versioned rules, human validation at every gate, a structured deal memory, on complex B2B deals. The same architecture fits any craft where experts make consequential decisions over a long cycle, and the studios that follow will be built on it · finance · human resources · procurement · supply chain.

iAM lands on the estate. Most of the IT budget and most of the people are there, and almost none of the AI conversation is. iAM covers the estate: break-fix to prediction and autonomy, ticket automation, AIOps and self-healing, knowledge that stays alive under operational volume. With the guardrail the white paper states: silent intelligence must not become opaque intelligence.

On a project, that frame has a name: the harness. Enough structure to keep the move into production within reach, little enough to leave experimentation free. This is the same idea, at the scale of a company.

HOW IT RUNS

FROM IMMERSION TO A DIGITAL TWIN OF THE CRAFT.



Scope is total from the start; value is delivered in slices along the chain.

IMMERSION	PROTOTYPE & PILOT	DIGITAL TWIN
Weak signals, real usage, available data - then contact with the craft itself. Not up-front scoping.	Two tracks on one cadence, three-week increments, slightly offset. They converge at delivery.	Business knowledge captured and structured, held by a specialized agent. The deliverable that lasts.


RUNNING UNDERNEATH, FROM DAY ONE

The full method is the subject of the next chapter.

THE DELIVERABLE THAT LASTS

THE REAL DELIVERABLE OF AN AI TRANSFORMATION IS CODIFIED BUSINESS KNOWLEDGE.

THE SOFTWARE ONLY USES IT.





CHAPTER 01 · IN ACTIVE USE

AI-FIRST

THE APPROACH · HUMAN-FIRST, END-TO-END, KNOWLEDGE-FIRST

A method for AI transformation: start from the craft, think the whole chain from day one, treat captured knowledge as the deliverable that lasts.

01

None of the three works alone

02

Build and make it react

03

Measuring value, and scaling

THE THREE CONVICTIONS

NONE OF THE THREE WORKS ALONE.

It says where to aim (the craft of leaders and experts), what scope to work at (the whole value chain, immediately), and what to keep at the end (the business knowledge captured on the way). It comes with a way to measure value, and with the conditions under which it fails.

Taken one by one, none of them is new. Augmentation over automation goes back to the 1950s, systems thinking to the 1960s, knowledge as an asset to the 1990s. What is mine is not the idea: it is executing all three together, and writing down how.

01

HUMAN-FIRST

THE TARGET IS THE CRAFT. NOT THE PROCESS.

02

END-TO-END

THE WHOLE VALUE CHAIN. FROM DAY ONE.

03

KNOWLEDGE-FIRST

THE APPLICATION IS REPLACEABLE. THE KNOWLEDGE IS NOT.

ONE PAGE EACH, WITH ITS OBJECTION

Each principle is stated, then argued, then given the counter-argument that is usually raised against it. The three sections that follow do that in order.

01


HUMAN-FIRST

THE TARGET IS THE CRAFT, NOT THE PROCESS.

Open any consulting AI roadmap and the sequence is the same: map the processes, pick high-value use cases, automate, count the savings. That reflex is old. In 1990 Michael Hammer was already writing "don't automate, obliterate" - and reengineering ended up delivering headcount cuts rather than the promised performance. He later conceded he had underestimated the human dimension.

Automation still happens. It shows up as a consequence: manual checks and reconciliations start to disappear on their own, without ever having been the objective.

Human-first moves the target. Transformation applies to the craft of experts and leaders - deciding, weighing trade-offs, negotiating, designing, convincing - not to the company's processes. That craft rests on tacit knowledge: we know more than we can tell. It is exactly what rule-based automation never reached, and it is where the value sits.



Capture how the best operate, then spread it.

CAUTION

Measured on generic copilots, AI lifts novices far more than seasoned experts. Which is why the work is not to hand out a copilot. It is to capture how the best operate, amplify it for them, and spread it across the team.

PUBLIC EVIDENCE

Three randomized trials at Microsoft, Accenture and a Fortune 100 firm, 4,867 developers: an AI coding assistant raised completed tasks by 26% (Cui et al., February 2025). The gain concentrated on the least experienced, a split the authors themselves report as not statistically significant. A copilot lifts the people who know least.

02


END-TO-END

THE WHOLE VALUE CHAIN, FROM DAY ONE.

95% of generative AI pilots deliver no measurable P&L impact (MIT NANDA, 2025). The usual explanations - data, skills, governance - miss the root cause, which is methodological: the use case.

Value sits in the connections between steps, not in the fragments. Optimize each part in isolation and the whole does not improve - a principle as old as systems thinking. Split a transformation into independent use cases and you get pilots that work in isolation and produce nothing at company scale, because the value was in what you left out: the links.

End-to-end rejects both ways of splitting it, the functional one (a portfolio of use cases) and the operational one (a minimal product you grow). The complete scope is designed and built from the start, then converted progressively along the chain. What made splitting necessary was cost: you could not afford to build the whole thing. AI as a build tool removes that constraint.



The fragments, and the chain.

CAUTION

This is not big bang. Scope is total immediately; delivery stays incremental, and moves toward continuous.

PUBLIC EVIDENCE


BCG, 2026: end-to-end process redesign is what separates organizations reaching 60% cost reductions from those capturing less than 20%. In the same research, a European media group ran more than 100 siloed use cases and produced no measurable P&L impact. It settles the scope question. What holds afterwards is the next chapter.

**OPTIMIZE THE FRAGMENTS
AND THE WHOLE
DOES NOT IMPROVE.**

03

KNOWLEDGE-FIRST

THE APPLICATION IS REPLACEABLE. THE KNOWLEDGE IS NOT.



No forms, no taxonomy. Grounded in sources, validated by the expert.

The durable deliverable is not the platform or the copilot. It is the business knowledge captured and structured while building it. Knowledge management promised this in the 1990s and failed, for one reason: nobody codifies. Writing down what you know takes time, pays nothing, and gives away power.

AI collapses the cost of capture. Knowledge is no longer written by the expert: it is picked up in the flow of the work – decisions, trade-offs, interactions – in natural language, with no form and no taxonomy.

And the implicit part comes out when the expert reacts to something concrete: put a prototype in front of them and they reveal the rule, the exception, the trade-off.

Once that knowledge exists as a structured asset, the application becomes interchangeable. You can unplug it, connect another intelligence, and rebuild the full experience in days.

CAUTION

Captured knowledge is only an asset once it is grounded in verifiable sources and validated by the expert. Specialized AI tools still get it wrong in a significant share of cases. Ungrounded, unvalidated knowledge is a hidden liability.


PUBLIC EVIDENCE

At Bloor Research, an analyst built a working twin of his own judgement over three years, then extended it to a fifty-person firm. The asset is real. No large organization has yet done it at its own scale.

HOW A PROJECT RUNS

BUILD AND MAKE IT REACT.

The shift this rests on: with today's tools you can aim very early for a solution that executes in a real environment, on real, connected data, instead of going through long specification cycles. You no longer describe a need for someone else to build. You build, and you make it react. Everything below follows from that.



The method, six moves.

01

IMMERSION, NOT UP-FRONT SCOPING

Scoping builds a box you then have to work inside. Immersion goes the other way: weak signals first - recurring irritants, real usage and the workarounds, data already available - from which a value proposition is derived. Then contact with the craft itself. Use cases are identified, and accepted as provisional: some will be wrong the next day. That is a feature of the method, not a defect.

02

TWO TRACKS, ONE CADENCE

A prototype and a pilot are built in parallel on short increments, around three weeks, slightly offset. The prototype runs on stubs and pseudo-real data: it builds confidence, and lets business users check that the target is right. The pilot is the same thing converted to a real environment - connected, real data, usable. They are not two sequential phases; they converge at delivery.

03

THE HARNESS

The minimal frame imposed on a vibe-coding environment: enough structure not to compromise the move into production, not enough to block experimentation. Three things to describe, and no more: execution environment, packaging and deployment, integration interfaces. Its cost never appears while building - it hits hard on the day you have to reach the real environment. So it goes in on day one.

04

INTEGRATION AND PRODUCTION ARE THE CRITICAL PATH

Once building stops being the constraint, the limiting factor becomes data integration and release to production – and those are human and organizational issues, not technical ones. The time lost is spent finding who holds the data and under which protocol. A rigorous release process is a good thing, but it is designed to operate an application, not to experiment with one: lowering the bar does not make it faster. Connect production from the start, and consider a release path dedicated to experimentation.

05

A DIGITAL TWIN OF THE CRAFT

The central deliverable. Business knowledge captured and structured, with a target: a specialized agent that holds the knowledge, rules, data and methods of one domain, and assists the teams and the experts themselves. How it gets built is an offer of its own: Craft Twin, chapter 02. A classic agile cycle, which formalizes a specification up front, works against this: it would remove the point of collecting knowledge continuously.

06

THREE ROLES

An AI Product Builder, who prototypes in depth and goes down to the smallest detail – the particular rule, the display convention, the edge case. An AI Delivery Builder, focused on end-to-end, integration and interaction, who makes the object live and connected. Both share the same foundation: product owner and vibe coder. One person can do both. The third role, AI Program Lead, is different in nature and does not merge: it keeps the whole ecosystem coherent and keeps pulling the program back up to the right level.



Immersion first. The box comes out of the craft, not before it.

MEASURING VALUE

NO BASELINE, NO PROVABLE ROI.


The standard objection to end-to-end is that a use case has an isolable return and a whole chain does not. The isolable figure is precise, and usually measures a pilot that will never scale. The discipline sits elsewhere.

MEASURE THE BASELINE FIRST

The first predictor of a demonstrable return is not the size of the model. It is whether the baseline for the whole craft was documented before anything was deployed. Without a baseline, no ROI is provable. With one, end-to-end becomes measurable.

READ VALUE WHERE IT ACTUALLY APPEARS

The effects that count only show up end-to-end: a whole step that disappears, a cycle that goes from weeks to days, teams released from a manual reconciliation. An isolated use case does not produce them, so it never measures them.



Without a baseline there is nothing to read the value against.


FUNDING BY DOING

Under a frozen budget, the workable lever is not a single cheque against a promise. The savings from the first increments fund the next ones. It keeps the program honest and the sponsor free.

WHAT I CONCEDE

Internal gains are structurally harder to isolate than a supplier's invoice - released time is not monetised. And accounting rules keep internally generated knowledge off the balance sheet: its value is real and quantifiable otherwise, through avoided rebuild cost, retained expertise and reduced vendor dependency, but it is a strategic asset, not a P&L line. Saying so is part of the argument.

**MEASURE THE ZERO STATE
OF THE WHOLE CRAFT,
BEFORE ANYTHING IS DEPLOYED.**



STARTING A PROJECT

TEN POINTS TO SETTLE IN THE FIRST DAYS.


The list I go through when a program starts. Nothing here is optional, and none of it costs anything at that stage - the cost comes from discovering these points at deployment.

POINT	WHAT IT MEANS
ANCHOR	One meaningful performance indicator, and a sponsor able to hold coherence over time.
DIRECTION	Real end-to-end as the main criterion for trade-offs.
SETUP	Two tracks, prototype and pilot, on short increments.
IMMERSION	Weak signals and a value proposition, then contact with the craft.
KNOWLEDGE	Capture running from day one: clean storage, recordings, documentation after the fact.
DATA & INTEGRATION	Data owners and integration architect identified and engaged early.
HARNESS	Minimal standards to reach a real environment without friction.
RELEASE PATH	Decide whether experimentation needs a path of its own.
ROLES	Program lead secured; the two builder roles staffed.
DECISION	Who decides what. Late decisions are the main hidden cost of these projects.

BEYOND THE FIRST PROJECT

MAKING IT A DURABLE CAPABILITY.

Everything above describes how to run one program. Turning it into a capability of the organization means replicating the approach past the first success, with multidisciplinary business, product and tech teams - and one shared layer underneath.



Start, build, run.

START	BUILD	RUN
Explore and get started: irritants and opportunities, existing usage, available data, fast internal prototypes, value estimated before going further.	Industrialize: turn high-value uses into robust solutions, connect them to the data platform and the IT landscape, structure reusable standards and components.	Operate and evolve: supervise usage and output quality, measure business impact, adjust workflows over time, manage governance, compliance and cost.

THE SHARED LAYER UNDERNEATH

Standards, components and know-how pass from one program to the next. Without it, every project starts again from zero.





CHAPTER 02 · IN USE ON MY OWN ENGAGEMENTS

CRAFT TWIN

THE FOUNDATION · A SECOND BRAIN, AND THE TWIN OF A CRAFT

A foundation that turns what a team already produces - meetings, mail, chats, documents, data - into a memory its experts can query, and then into a twin of their craft. Three layers: a configured working environment, a knowledge system on vector and graph stores, automated capture of the flows. Four depths, and you stop at the one your case justifies. The components run on my own engagements. This is the first time they are packaged as an offer.

01

A memory of the craft

02

Three layers, four depths

03

What an engagement keeps

WHAT IT IS

A MEMORY OF THE CRAFT, NOT A DOCUMENT BASE.

An expert produces hundreds of artifacts a week: meetings, transcripts, mail, documents, decisions. That material is the substance of their expertise, and it is unusable. Scattered, never consolidated, not transferable.

01

**THE UNIT IS THE ROLE,
NOT THE PROJECT**

A role crosses several projects and contexts. It carries the instincts, the decision criteria, the network. Model by project and you build silos that die with the project. Model by role and the memory outlives the program that funded it.

02

WHAT IT PRODUCES

An expert asks in their own words and gets an answer that cites its source and its date. Derived knowledge is marked as derived and weighted below a verbatim record. An unresolved contradiction is shown with both sources, not settled quietly.

03

**SECOND BRAIN,
THEN TWIN**

The two terms get used as synonyms. They name two systems. A second brain is externalized memory: I can find what I know and reason over it. A twin is a live model of a role: it reacts the way the original does. Crossing over takes continuous capture, an explicit model of how the role decides, and an answer to what this expert would decide and why.

THE QUESTION IT ANSWERS

WHAT WOULD THIS EXPERT DECIDE, AND WHY?

HOW IT IS BUILT


THREE LAYERS. ALREADY BUILT.

None of this was written for the occasion. The three layers exist, they run, and they are why a specialized AI engagement moves fast.

LAYER	WHAT IT IS	
THE WORKING ENVIRONMENT	The tools business teams work in, configured in the client's own tenant: Anthropic Claude (Cowork and Code), Google AntigraVity, OpenAI Codex, Microsoft Copilot. Configured, not just licensed.	
THE KNOWLEDGE SYSTEM	Source processing, cleaning, enrichment, storage and downstream use, on vector and graph stores. Private secured perimeter, or sovereign hosting.	
THE CAPTURE PIPELINES	Twelve connectors for meeting transcripts, mail, chats, messages and calendars, plus a resolver that follows the documents referenced inside them. Incremental, idempotent, read-only on every source.	
RAW, NOT RESUMMARIZED	REPLACE, DO NOT ACCUMULATE	PACKAGED, NOT HANDMADE
Capture stores the material unaltered: the verbatim record, the body of the message. A summary produced at capture destroys information for good. Summarizing comes later, and it can be regenerated.	Every sheet carries a logical identity and a physical fingerprint. Same fingerprint, nothing is written. New fingerprint, the previous version is marked superseded. Stale content stays in the journal for audit and is never served.	Building with my own hands does not mean rebuilding everything each time. It means I control what I deploy, and I deploy components that have already run.


A FIRM SELLING DAY RATES
BRINGS NO ACCELERATOR.

THE CHANGE GATE · FOUR OUTCOMES



Stale content is kept for audit. It is never served.

CAPTURE: INDEX FIRST, FETCH LATER · TWELVE CONNECTORS



Filtering at capture loses the material for good. Volume is controlled before, relevance is judged after.

HOW FAR YOU GO

FOUR DEPTHS. YOU STOP AT YOURS.

The foundation is not an all-or-nothing deployment. Each depth works on its own, and each one names the condition that justifies moving up. Most cases stop at the second.



Each depth works on its own. The one above adds to it.

WHEN TO GO DEEPER

The move from FLOW to SUBSTRATE has documented thresholds. Orders of magnitude to be recalibrated for each context, not acceptance criteria: more than 1,500 sheets or 20 MB of markdown, more than 50 long documents that resist distillation, more than 10% retrieval failure on a quality test, more than three search round-trips on average. Below them a vault and an agent do the job, at a fraction of the cost.

WHERE THE OTHER OFFERS PLUG IN

AILC stops at the first depth: a personal working environment with structured memory. ASES sits at the second: a deal memory fed by the flow of the pursuit. iAM needs the third: a knowledge base that stays alive under operational volume.

THE FOUR DEPTHS, SIDE BY SIDE

DEPTH	WHAT IT ADDS	WHAT IT PRODUCES	WHEN IT IS CALLED FOR
VAULT	A structured vault, distilled sheets with a documented header, an index, an agent that reads files.	Sourced answers in weeks. The full loop, proven small.	A corpus that exists and nobody can use.
FLOW	Automated capture of meetings, mail, chats, messages and calendars. Incremental, verbatim, read-only.	The memory stops going stale.	Knowledge is being produced now, and lost at the same rate.
SUBSTRATE	An append-only journal as the source of truth, vector and graph projections rebuilt from it, versioning, contradiction handling, hybrid retrieval.	Stale content replaced rather than stacked, entities resolved against one reference base.	Volume, or questions that cross contexts, have broken the level below.
TWIN	The decision model of the role: rules, trade-offs, criteria, and the reasons behind them. Provoked, not captured.	An answer to what this expert would decide, and why.	Critical expertise held by too few people, with a date on it.

CAUTION

Knowledge management promised this in the 1990s and failed, for one reason: nobody codifies. Writing down what you know takes time, pays nothing, and gives away power. What changed is the cost of capture, not the ambition. One failure mode from that era is still live: a memory with no freshness management answers confidently from dead data, and the failure is silent.

WHAT AN ENGAGEMENT DOES

ONE ROLE, THE WHOLE CHAIN.

The order matters more than the ambition. Three depths across three roles at once prove nothing that one role taken end to end does not prove faster.

01

INVENTORY BEFORE ANYTHING ELSE

Where the material sits, who holds it, under which protocol, and what is locked by classification or rights management. The obstacles are organizational, not technical.

02

THE REFERENCE BASE COMES EARLY

People, organizations, products, roles, aliases. Every extracted mention resolves against it. Without it, three spellings of the same person are three entities and the graph is decorative. Deferred, it forces reprocessing of everything already ingested. It is the highest-return move upstream.

03


EXPECT THE EFFORT WHERE IT SITS

Capture is visible and gratifying, and it is roughly a fifth of the work. The rest is preparing what does not arrive as a flow: documents, datasets, code, media. Planning it the other way around is the usual reason these programs run late.

04

THE TACIT PART IS PROVOKED

The reason behind a trade-off appears in no flow. Nobody writes down that they decided that way because the supplier had already failed twice. It comes out of interviews and annotated decisions. It is not automatable, and it is what separates a memory from a twin.



WHAT YOU KEEP

THE FOUNDATION STAYS. SO DOES THE KNOWLEDGE.

The application on top is interchangeable. What is worth owning sits underneath it.

OPEN FORMATS, READABLE WITHOUT ME

Markdown sheets with a documented header, an append-only journal, a reference base. The vector and graph stores are projections rebuilt from the journal: losing one costs a rebuild, not a corpus. None of it needs my tools to be read a year later.

IT CAN RUN INSIDE YOUR PERIMETER

Inference and embedding included, on local or sovereign hosting. For confidential perimeters that is a hard constraint rather than an option, settled before the first connector runs.

IDENTITIES STAY WALLED

A consultant or a supplier operates under separate professional identities, and the wall between them is contractual. It is enforced at the database level, before any search or graph traversal, never as a post-filter. Deduplication and entity resolution never cross it either.

READ-ONLY, ALWAYS


Capture never writes to a source. No message sent, no event changed, nothing deleted.

CAUTION

Personal flows cannot be captured on the same terms as a meeting transcript. Mail, private messaging and chats carry third-party personal data, and secrets travel through them: service credentials arrive by mail more often than anyone admits. They are flagged as sensitive at capture and never pulled down. The perimeter is declared source by source, before a connector runs.



Name one role, and we'll start there.



CHAPTER 03 · PROGRAM RUNNING

AILC

AI LEADERSHIP COMPASS · EXECUTIVE AI COACHING

Two days for a senior executive to take ownership of AI in their own role. Day one, they write their doctrine. Day two, they build and govern a personal AI working system. Cohorts of six to eight, no lectures: they produce, and they leave with what they produced. Co-built and co-run with David Horvat, former CIO at Allianz. The program has its own site, ai-ldc.com - in French.

01

Think, then do

02

Four assets



03

How a cohort runs

THE TWO DAYS

THINK, THEN DO.

The rule the program runs on: work with AI, and have it do as little of the thinking for you as possible. The executive decides, AI executes.



Two days, two outputs.

DAY ONE · THINK

EXECUTIVE AI PLAYBOOK

Each executive writes their own position on what AI does to their function - anchored in their reality, not in a market overview.

A business AI playbook of around twenty pages, plus a one-page executive summary.

DAY TWO · DO

EXECUTIVE AI ACCELERATOR

They configure, prototype and govern a personal AI working system that can be handed to their team.

One or two prototyped use cases, and a configured AI workspace they keep.

THE EXECUTIVE DECIDES, AI EXECUTES.

WHAT THEY LEAVE WITH

FOUR ASSETS.

Usable the Monday after – in a committee, or when taking up a new role.

ASSET	WHAT IT IS
01 BUSINESS AI PLAYBOOK	Their doctrine: vision, convictions, one-year roadmap, 90-day plan.
02 EXECUTIVE SUMMARY	The one-pager to put in front of an executive committee.
03 ONE TO TWO USE CASES	Prototyped on their own moments of truth. Evidence of execution, not slides.
04 A CONFIGURED WORKSPACE	A personal AI working environment with structured memory, which stays theirs.

WHAT CHANGES

THE ARBITRATION MOVES TO THE PERSON WHO OWNS THE OUTCOME.

Today an executive approves AI proposals that others have built and framed. After two days they hold a written position of their own, and use cases prototyped on their own moments of truth. Approval becomes judgement.

The decisions happen either way. The question is whose material they rest on.



THE FORMAT

HOW A COHORT RUNS.

DURATION Two full days - one for the playbook, one for the accelerator.	COHORT Six to eight executives. Group coaching among peers.
METHOD AI-assisted production, debate, peer review. No lectures.	WHAT TO BRING A laptop and a dedicated AI environment, set up with each participant beforehand.
OPTIONAL Follow-up over three to six months: office hours and refresher sessions.	WHO IT'S FOR No technical prerequisite. Works for an executive coming alone among peers, or for a leadership team deciding to move together.

BOOK A COHORT OF SIX TO EIGHT
jerome.fenyo@goyave-consulting.com · ai-ldc.com





CHAPTER 04 · CURRENTLY BEING DEPLOYED

ASES

AGENTIC SALES ENGINEERING STUDIO

An agentic studio for complex B2B deals, €5M to €500M and above. It covers the full bid lifecycle through specialized agents working under explicit rules, with human validation at every gate and a structured deal memory. The same architecture transfers to other business functions.

01 The full bid lifecycle	02 Agents under rules	03 The studio in diagrams
-------------------------------------	---------------------------------	-------------------------------------

WHAT IT COVERS

THE FULL BID LIFECYCLE.

Not a proposal generator. A studio that carries a deal from the first signal to the signature, and keeps what it learned.

STAGE	WHAT HAPPENS
QUALIFICATION	Score the opportunity across feasibility, desirability, winnability and viability. Produce a go or no-go that can be defended.
SOLUTION	Win themes, solution architecture and delivery model, challenged against what the client actually decides on.
PRICE-TO-WIN	Commercial model, competitive price position, margin trade-offs, negotiation playbook.
STORYLINE	The executive narrative and the orals, built from the client's decision criteria rather than from the solution.
REVIEW GATES	Opportunity, solution, proposal and contractual gates: materials, compliance checklist, decisions recorded.
DEAL MEMORY	Everything above stored as structured memory, so the next deal starts from what the last one learned.

FROM THE FIRST SIGNAL TO THE SIGNATURE.

Six stages, one continuous rail. The cadence that runs them comes later in this chapter.

PUBLIC EVIDENCE

Vodafone, March 2026: an agent produces an RFI/RFP draft ninety percent finished in minutes, against the ten to thirty days a telco response normally takes. It reads policies, past responses, the product catalogue and pricing, then hands over to the account manager. A proposal generator, and a good one. Six stages sit around it.

HOW IT IS BUILT

AGENTS UNDER RULES, HUMANS AT THE GATES.

The same three convictions as everywhere else: the craft of the best bid people is the target, the whole cycle is in scope from day one, and the deal knowledge is what remains.

01

SPECIALIZED AGENTS, EXPLICIT RULES

The studio is not one assistant. It is a set of roles – qualification, content, planning, orchestration, pricing, challenge, analysis – each with a defined scope and each operating under written rules. Those rules are readable, arguable and versioned. Nothing runs on an implicit prompt.

02

HUMAN VALIDATION AT EVERY GATE

The agents prepare, structure, challenge and produce. The decision stays human at every gate: qualification, solution, price, storyline, submission. This is not a safeguard bolted on afterwards; the review gates are the backbone of the design.

03

DEAL MEMORY

Bid knowledge is normally lost twice: when the team disbands, and when the next bid starts from a blank template. The studio keeps a structured memory of the deal – decisions, trade-offs, pricing logic, what convinced and what did not – so knowledge compounds from one deal to the next.

04

CONTINUOUS CHALLENGE

A dedicated role does nothing but attack the deal: assumptions, blind spots, the credibility of the delivery story, the price position. Being challenged early by a machine that has read everything is cheaper than being challenged late by the client.

NOT BOLTED ON AFTERWARDS

THE REVIEW GATES ARE THE BACKBONE OF THE DESIGN.

WHY THE STUDIO EXISTS

WHERE STRATEGIC DEALS ARE LOST.

Twenty-five years of large pursuits, and the same four causes. None of them is a weak solution.

01 The solution is almost never the cause. Deals are lost on a poor reading of where the real power sits.	02 Price is set too late, or without competitive discipline.
03 Execution degrades under pressure: the storyline and the client's perception go with it.	04 Without rigorous orchestration, the critical moments are left to chance.



THE STUDIO IN DIAGRAMS

HOW IT FITS TOGETHER.

Four diagrams: what the studio is for, what it can do, how it runs day to day, and which agents do the work.

PURPOSES OF THE STUDIO · FIVE PURPOSES


WIN THE DEAL	Favor winning the deal over the theoretical optimum of the solution.
CREDIBLE DELIVERY	Ensure a credible post-signature delivery, risks identified and the solution under control.
EXCELLENT DELIVERABLES	Produce deliverables that hold up, for the client as much as for internal review.
STEER THE DEAL	Keep the deal steered end-to-end, including decisions, trade-offs and the investment.
CRITICAL ANALYSIS	Keep a continuous critical analysis of the strategy, the price, the solution and the influence map.

SEVEN CAPABILITIES, ONE LOOP · MEMORY FEEDS THE NEXT DEAL

UNDERSTAND & SET UP Read the opportunity, open the deal, set the artifacts up properly.	PROACTIVE ENGAGEMENT Suggest the next move rather than waiting to be asked.
ASSISTED PRODUCTION Draft the deliverables, versioned, for humans to take further.	DEAL STEERING Keep the plan, the trade-offs and the risks under control, day to day.
CRITICAL EVALUATION Attack the deal continuously: assumptions, price, blind spots.	ORCHESTRATION Synchronize humans, agents and tools; hold priorities and dependencies.
STRUCTURED MEMORY Build memory per deal, sector and client – wins and losses both.	AT THE CENTER: THE STUDIO The seven capabilities close a loop around it.

DEAL LIFECYCLE, THE OPERATING CADENCE

DAILY AND WEEKLY.



One continuous rail, from the first signal to the signature.

FROM SIGNAL TO OPPORTUNITY

Turn weak signals into structured opportunities and initialize clean artifacts.

ALPHA / BETA VERSIONS

Generate complete alpha and beta versions, and give the human contributors something to react to.

STRATEGY & PRICE ANALYSIS

Challenge the strategy, the solution and the price. Reveal the blind spots.

MEMORY & IMPROVEMENT

Structure the memory per deal, learn from wins and losses, by sector and client.

ACTION PROPOSALS

Propose proactive actions, suggest the workshops, push for controlled risk-taking.


DAILY & WEEKLY FOLLOW-UP

Steer the deal day by day and week by week, tracking trade-offs and risks.

AGENT & TOOL COORDINATION

Coordinate humans, agents and tools; manage priorities and dependencies.

SEVEN ROLES AROUND THE HUMAN BID TEAM.



Deal analyst – analyzes qualification and competitive position.

Price-to-win shaper – develops pricing strategies and models.

Content provider – generates client-facing content.

Deal challenger – validates assumptions and plans for risks.

Deal planner – develops strategies and solutions.

Bid orchestrator – manages timelines and task assignments.

Deal initializer – triggers and sets up new deals.

AT THE CENTER: THE HUMAN BID TEAM

Every gate is theirs. The agents prepare; the team decides. Each role has a defined scope and written, versioned rules. Nothing runs on an implicit prompt.

PUBLIC EVIDENCE

Fujitsu, March 2025: an orchestrator coordinating specialised agents made sales proposal creation 67% more productive. The agents retrieve and synthesise knowledge scattered across internal sources. The architecture holds at that scale. The case describes no validation step for what the agents produce. That step is the centre of this page.

BEYOND SALES

A STUDIO, NOT A SALES TOOL.

Nothing in the architecture is specific to selling. Specialized agents, explicit rules, human validation at defined gates, structured memory: the same structure fits any function where experts make consequential decisions over a long cycle.

FINANCE	HUMAN RESOURCES
PROCUREMENT	SUPPLY CHAIN

Each new studio reuses the same foundation: agents, rules, validation, memory. What changes is the craft it learns.



An agentic system: named roles, and a human who decides.



Bring a deal and we'll go through it together.



CHAPTER 05 · WHITE PAPER PUBLISHED, 2025

iAM

INTELLIGENT APPLICATION MANAGEMENT

Application management is where most of the IT budget sits, and where most of the AI conversation is not. iAM covers its automation: incidents and tickets, proactive maintenance, root-cause analysis, living knowledge, and the move from technical availability to measured experience.

01 Three shifts	02 Seven use cases	03 Five categories of risk
---------------------------	------------------------------	--------------------------------------

WHAT CHANGES

THREE SHIFTS.

The name of the white paper comes from what this looks like when it works: operational intelligence that is always on and never noticed. With one guardrail, stated in the paper - silent intelligence must not become opaque intelligence.

REACTIVE TO PROACTIVE

Break-fix waits for the incident. Prediction and prevention change what the team spends its days on: anomaly detection ahead of failure, and remediation that runs on its own where it can be trusted.

SLA TO XLA

Technical availability says little about whether the business could work. Experience level agreements measure the experience itself, alongside the classic indicators - and change what the operating model optimizes for.

INTELLIGENCE AT THE CORE

Not a chatbot on top of the service desk. AIOps, automatic ticket classification, correlation, predictive maintenance: the intelligence sits in the operating chain, not beside it.



Availability reads flat while the measured experience moves.

SILENT INTELLIGENCE MUST NOT BECOME OPAQUE INTELLIGENCE.


The three shifts are sequenced by what the data and the tooling can support, not by ambition. The next page shows where they lead.

PUBLIC EVIDENCE

Mutua Madrileña runs intelligent observability with event-driven automation across more than sixty middleware platforms: service tickets down fifty percent. At ABB, four thousand service-desk tickets are handled automatically and close to forty thousand hours of processing removed. That is the first shift, and part of the third. The second one, SLA to XLA, is untouched.

FROM BREAK-FIX TO AUTONOMY

THREE STAGES, ONE GUARDRAIL.



BREAK-FIX The incident arrives, a human triages it, a human fixes it. Availability is the only number that is measured.	PREDICTION Signals are observed and correlated. Anomalies surface ahead of failure; root cause takes hours rather than days.	AUTONOMY Self-healing runs where it is trusted, knowledge stays alive, and experience is measured alongside availability.
---	--	---

The guardrail, at every stage: silent must not become opaque.



WHAT IT COVERS

SEVEN USE CASES.

They are not alternatives to pick from. They depend on each other, and they are sequenced by what the data and the tooling can actually support.

USE CASE	WHAT IT DOES
01 INCIDENT AND TICKET AUTOMATION	Detection and classification, intelligent routing, clustering and correlation of related tickets.
02 PROACTIVE AND PREDICTIVE MAINTENANCE	AIOps: predictive analytics on operational signals, self-healing and autonomous remediation.
03 ACCELERATED ROOT-CAUSE ANALYSIS	Correlation across sources, generated summaries, agents that walk the topology to find the origin.
04 AUGMENTED KNOWLEDGE MANAGEMENT	A knowledge base that stays alive: dynamic capture, semantic search, no manual curation debt.
05 INTELLIGENT USER SUPPORT	Virtual agents on L1 self-service, experience measured along the way.
06 GENAI MAINTENANCE	Code assistance on maintenance work, including legacy nobody wants to open.
07 HYPERAUTOMATION	End-to-end orchestration of the operational workflows, once the pieces above hold.



WHAT AN ENGAGEMENT DOES

LAYER BY LAYER, NOT ALL AT ONCE.

The order matters more than the ambition. Data foundation first, then observability, then knowledge, then agents. Skipping a layer is how programs end up with intelligent automation on top of unreliable signals.

01 Assess the readiness of the current operating model, tooling and data


02 Build the data foundation the rest depends on

03 Deploy AIOps and observability, then self-service knowledge

04 Add agentic automation where the ground is stable enough for it

05 Move the contract and the indicators from SLA toward XLA

06 Train the internal teams, so the capability stays with them



WHAT GETS IN THE WAY

FIVE CATEGORIES OF RISK.

A third of the white paper is about what goes wrong. It is the part most vendor material skips.

01

DATA QUALITY AND FOUNDATIONS

Fragmented tooling and weak governance are the usual blockers. Including a point rarely made: inference economics. Reasoning loops cost money, so budget caps belong in the design.

02

MODEL RELIABILITY AND BIAS

Inherited bias, generative behavior, lifecycle management. Shadow mode first, human in the loop where the consequence is real.

03

INTEGRATION AND LEGACY

Adapters, and the fact that limited observability means blind AI. Controlled rollout rather than a big switch.

04

PEOPLE, SKILLS, CHANGE

New roles appear - AI-enabled operations engineers - and they are designed with the teams, not announced to them.

05

SECURITY, PRIVACY, COMPLIANCE

GDPR and sector regulation, agent exposure, responsible use and IP.

THE FULL ARGUMENT

THE SILENT INTELLIGENCE

43 pages, v1.0. The state of the art of application management moving to intelligent and autonomous operations. Download it at goyaveconsulting.com/resources.

ABOUT



JÉRÔME FENYO
Founder & AI-First Consultant

TWENTY-FIVE YEARS
LEADING BUSINESS LINES
AND LARGE DEALS. I
STILL BUILD WITH MY
OWN HANDS.

I start from the craft of leaders and experts - how they decide, weigh trade-offs, negotiate, design - and build the system around it. The work is to capture how the best operate, then spread it across the team.

Before Goyave: eighteen years at Capgemini - head of the Insights & Data business line, then VP for solutions and presales of application services - and seven years at Atos. Large transformation and outsourcing deals, portfolio governance, data and AI.

On the same engagement I work as consultant, product owner and vibe coder. That is what keeps the method consistent: it gets written while it is being used.

HOW I WORK

LARGE ORGANIZATIONS FIRST The problems I work on are worth addressing at scale: an application estate, hundreds of people practicing the same craft, decisions that pass through several levels. Below that, a team buys a tool and moves on - which is often the right answer. I do not take startups, small businesses, or single-team pilots.	ASSETS, THEN SCALE The method is written down and executed. The foundation is built and running. The first studio is in pilot. What arrives on day one was built before the engagement started. Scale comes from technology partners who deliver on those same assets, so a larger programme adds capacity and keeps the same method.	WHAT STAYS WITH THE CLIENT The captured knowledge and the foundation it sits in stay with the client, in open formats. The team that runs them afterwards can read them and extend them. Continuity sits in the deliverable and in the people who inherit it, which is what makes the transformation survive across team changes.
---	--	--

Based in Paris. Working across Europe. 6-18 month engagements · Independent consulting. English & French.



CONTACT

LET'S TALK ABOUT THE CRAFT YOU WANT TO TRANSFORM.

jerome.fenyo@goyave-consulting.com

goyaveconsulting.com

linkedin.com/in/fenyo

Based in Paris. Working across Europe.

6-18 month engagements · Independent consulting.

English & French.



GOYAVE CONSULTING

© 2025-2026 Goyave Consulting™. All rights reserved.

"Goyave Consulting" and "Goyave Solutions" are trade names used by Jérôme FENYO.