



SOFTED · SEPTEMBER 24, 2026

Owned Intelligence Is an Estate

Governing a portfolio of models and the agents acting across them

Stacy Gonyou

Director of AI Strategy and Solutions, IntePros

Jason Hunt

Vice President of AI Solutions, IntePros Federal

The unit of measure just split in two



Adaptive thinking

The model iterates and self-corrects before it answers. That phase costs more than the prompt.

The hidden log


Providers rarely show the reasoning. Every agent in a workflow runs the same curve.



WHY OWN, AND WHAT YOU ARE OWNING

Spend is compounding. Return is not.

SPEND



\$2.7T Gartner forecast, worldwide AI spend in 2026, up 49.5%

RETURN

37%

of organizations see any positive EBIT contribution from AI, flat year over year

6%

qualify as high performers

GETTING TO PRODUCTION

Share of organizations

Evaluate **60%**

Reach a pilot **20%**

Reach production **5%**

Sources: Menlo Ventures, ~500 US enterprise decision makers. Gartner, “Gartner Forecasts Worldwide AI Spending to Grow 49.5% in 2026.” McKinsey, State of AI 2026 (n=1,719). MIT Project NANDA, “The GenAI Divide,” 2025 (preliminary, not peer reviewed).

What a domain is, and why you will have several

A domain is a bounded body of work with its own data, its own business rules, its own definition of a correct answer, and usually its own cost center.

Same question: “Can I see that record?”

HR MODEL

No.

Policy-backed

HELP DESK MODEL

Yes.

Almost always

Several domain models, coordinated by an agentic orchestration layer. That is structure, not sprawl.

On bounded, repeatable, high-volume tasks with clear evaluation criteria, a domain-constrained model matches or exceeds frontier models on accuracy, and beats them on cost, latency and consistency.

Retrieval often beats fine-tuning for injecting broad knowledge. Fine-tuning wins on task execution, structured output and cost per call.

START HERE

Choosing your first domain

Takeaway 1 of 3

1

Pick the cost center with the most repeated decisions.

2

Write down what a right answer looks like there.

MOST OFTEN OVERLOOKED

3

Check you own the data that teaches it.

4

Name who owns that judgment today.

SO YOU OWN THEM. NOW THE WORK STARTS.

Owning the model is the easy part.
Controlling how it **behaves** is the work.

JOB ONE

Technical governance

JOB TWO

Organizational governance

Built for places where data can't leave the building

Air-gapped

- No tolerance for data leakage
- Only approved users get in
- Each team's data stays separate

THREE NON-NEGOTIABLES

01 Encryption

Data is encrypted in transit and at rest

02 Full logging and audit trails

Every data access, model query, and admin action is logged

03 Encrypted model parameters

Your Trained models can't be copied or pulled apart

So here is what we actually built

PII is removed or masked before training and inference.

CONTROLS

- TLS 1.3, in transit, AES-256 at rest
- Role-based access with multifactor authentication
- Every data access, model query, and admin action is logged
- Trained model weights are encrypted

COMPUTE SPLIT

- GPU**
Initial training
- CPU**
Handles most inference, so ongoing costs stay low

PUBLISHED FLOOR

- 8 cores · 32 GB RAM · 0.5 TB SSD · x86_64**
- GPU to train, CPU to run
- Tactical edge**
- Air-gapped**
- Secure enclaves**
- Fully offline**

THEN DO THIS

Know how your data actually moves

Takeaway 2 of 3

- 1 Inventory what data the domain actually runs on.
- 2 Trace where it moves and who can see it.
- 3 Check the documentation is current enough to trust.
- 4 Write down what you can't answer yet.

THE ORGANIZATIONAL PROCESS

Everyone is working on the technology. Almost nobody is working on the organization.

THE SECOND JOB

A named human owns each model's output

Financial return is measured, not assumed

Agents are managed as a digital workforce

“AI isn't failing. AI literacy is.”

Gartner ThinkCast, “Start With AI Literacy: Why Most Leaders Are Missing the Point,” episode description, July 23, 2026

WHAT'S MISSING

No published outcome data, measurement method or evaluated program showing enterprise literacy drives return.

Enablement works, and it cuts both ways

758 consultants, three groups: no AI · AI access · AI access plus training


INSIDE THE AI’S BOUNDARY



The trained group significantly outperformed both others.

OUTSIDE THE BOUNDARY

Points below the no-AI control



Training raised trust. People checked the output less.

AI LEVELS THE FLOOR

43% vs 17%

Lower vs top performers, consultant study

34%

Gain for novices across 5,000 support agents; negligible for experts

Sources: Dell’Acqua et al., “Navigating the Jagged Technological Frontier,” HBS WP 24-013; Organization Science 2025 (pre-registered).
Brynjolfsson, Li & Raymond, QJE 2025.

Mindset before process

Change is an operating state, not a project. By the time a rigid process is documented, the capability has moved.

89

studies

Supervisor and leadership support is one of the strongest predictors of whether training becomes daily behavior.

Source: Blume et al., "Transfer of Training: A Meta-Analytic Review," Journal of Management, 2010.



Raising our own AI IQ

How the IntePros enablement program is structured

01

**Find the people
others go to**

Open call and
manager nomination

02

One domain each

The lead writes their
own goal, measure,
baseline and target

03

**Mindset, then
architecture,
then production**

In that order

04

**Each lead ships
into production**

Not through a course

05

**Baseline,
monthly read,
close-out**

Measurement built in
from day one

AND THEN DO THIS

Find who runs this, domain by domain

Takeaway 3 of 3

1

Assess your organization and find your domain owners

Not the org chart. Who do colleagues already go to when they're stuck? Open call, plus ask managers.

2

Give each owner a domain and let them set the goal

Their own goal, measure, baseline and target. If the goal is yours, the curiosity isn't theirs.

WHERE MOST GO WRONG

3

Have them architect the domain

The agents it needs, a named human accountable for each, how they're governed, and when each expires for review. Existing agents get pulled in and given an owner.

4

Build it into how the domain actually runs

Mindset, then architecture, then production. Ship into the main process, not a pilot beside it. Baseline, read monthly, close out.

NEXT STEP

AI Insight Session

Bring a problem from inside your own organization. Thirty minutes, and whoever from our practice it needs.

dgijsbers@intepros.com David Gijsbers