

How ANI changes the shape of enterprise IT

A point of view on Artificial Neural Intelligence: tools, data, AI and people working as one connected system, with a named person accountable for the outcome. It is an argument from reasoning and early delivery experience, not a proven result.



Andrew Nolan
Chief Executive Officer



Jack Dinneen
Chief Product and Technology Officer

CONTENTS

Abstract

Key findings

1. What we mean by Artificial Neural Intelligence, and what we don't

2. Why now: agents arrived, the silos stayed

3. From neural networks to neural intelligence

4. Four things, working as one

5. The nervous system, and where people stay in charge

6. What this changes for an enterprise

7. How Anitech puts ANI to work

8. What we don't claim

9. Where to start

Figures

References

About the authors

ABSTRACT

Agents arrived in the enterprise faster than the enterprise connected itself. Most organisations now own capable models and a growing number of agents, yet the work that matters still crosses functions by being carried between people. This paper sets out a point of view we call Artificial Neural Intelligence (ANI): tools, data, AI and people working as one connected system, with the outcome at the centre and named humans accountable for it. It explains what the term means and does not mean, why the timing matters, how the four parts fit together, where people stay in charge, what changes for an enterprise and what Anitech does about it today. It also says what we have not proven. The companion paper on ETA is the design; this one is the argument.

KEY FINDINGS

- 1 Capable models and per-function agents have not connected the organisation. Work that spans functions is still carried by people, and nobody holds the process as a measurable thing.
- 2 ANI is a point of view about how the parts of an enterprise should work together. It is not a product, a platform or a model, and no vendor can sell it as a single item.
- 3 Four things have to work as one: the tools you already run, your data connected through shared meaning rather than copied, AI on top, and people who change how the work is done from inside.
- 4 Autonomy should be earned in small steps, with named people deciding at the points that matter. Measurement has to come before automation, or nobody can say what the AI changed.
- 5 Our evidence is thin. We have one two-week pilot with one client, self-reported, with quality outcomes not yet measured. The argument stands on its reasoning, and we are working to earn the proof.

1 What we mean by Artificial Neural Intelligence, and what we don't

Most large organisations now have capable AI models, and many have bought agents for individual functions: a coding assistant for engineering, a drafting assistant for sales, a triage agent for support. Each is useful. Together they have not changed how the organisation works, because the important work

crosses functions and the crossing is still done by people sending messages to each other. We call the alternative Artificial Neural Intelligence (ANI).

Artificial Neural Intelligence (ANI)

The idea that no part of your business should run as an isolated project: connected, your systems, data and teams behave like one nervous system, with your outcome at the centre.

The definition is about the organisation. It says nothing about a particular model, vendor or architecture. ANI describes a state an enterprise can be in, and a direction of travel, and it carries a test: can the parts of the business that touch one outcome behave as a single system, and is a person accountable for that outcome? We have delivered enterprise software since 2006, and in our experience the difficulty sits mostly in the design of the work and in who is accountable for it. It is rarely in the size of the model.

The word collision

In the AI literature, ANI already stands for artificial narrow intelligence, the term for systems that do one task well, as opposed to artificial general intelligence. A reader who meets our acronym cold could reasonably assume we mean that. We do not. We always spell ours out, and we chose the word neural on purpose. A neural system is defined by its connections and by what passes along them, and not by how clever any single cell is. Narrow intelligence, in the sense of a capable agent confined to one function, is the problem this paper describes. If the collision costs us some clarity, we accept that.

What ANI is not

- It is not a model. No model, however good, is the thing we mean.
- It is not a single product you buy. ETA is how we make it real in product form, but the idea does not depend on it.
- It is not a promise of artificial general intelligence or of machines running a business unattended.
- It is not a replacement for your systems of record. The tools you run today stay.
- It is not a claim that organisations need more AI. Many have plenty. They need it connected and governed.

2 Why now: agents arrived, the silos stayed

Not long ago the question in most boardrooms was whether to use AI. The better question now is what changed because of it, and most organisations are better equipped to answer the first than the second.


We think many enterprises sit on a four-layer ladder. Systems of record sit at the bottom, each owned by a function. Above them are point-to-point automations that move fields between tools, and above those, increasingly, an agent per function. Above that there is nothing: no shared context, no shared rules, no shared record of the work. The ETA paper draws the ladder in full. The point here is that every silo has gained an agent and the agents do not talk to each other, so people still connect them.

Systems hold artefacts, nothing holds the process

Take one ordinary commitment. A deal is signed, finance sets up the customer and raises the purchase order, an invoice goes out, operations mobilises the work. That is one commitment crossing three functions and, in most companies, five or more systems. Each system holds artefacts, such as a contract, a purchase order, an invoice and a staffing record. None holds the process that links them. The elapsed time lives in the gaps, in a message to a colleague, a spreadsheet, someone remembering to chase.

Today

Each function runs its own tools, data and agents. People carry the work across the gaps.



Connected

One shared meaning and one record. The work moves between functions without a person in the middle.




Figure 1. Siloed functions, each with its own tools and agents, against the same functions connected so that work and context move between them with people in control.

Ask who owns the CRM and you will get an answer, and the same for finance, support and delivery tooling. Ask who owns the time from signature to cash and often there is no answer, or the answer is a committee. Everyone owns the machinery, and almost nobody owns the number. An organisation that cannot measure how its cross-functional work performs cannot tell whether an agent made it better.

Models have also improved faster than organisations can absorb them, so the constraint has moved. It is whether the model has the right context, is permitted to do the right thing, and whether anybody can show afterwards what it changed. Public research hints that context is the hard part: an agent that scored 91.2% on the older Spider 1.0 text-to-SQL tasks solved 21.3% of tasks built from real enterprise workflows in Spider 2.0 (Lei and others, ICLR 2025, reference 1). That is a 2024 benchmark and not a business, so we offer it as a direction and not a measure of today.

We make the diagnostic argument at more length in our article [Agents arrived, the silos stayed](#).

3 From neural networks to neural intelligence

The phrase we use on our own site, from neural networks to neural intelligence, describes three steps. Each is a real stage of maturity.

1. The model. A trained neural network predicts, drafts, classifies and writes code. By itself it has no memory of your business, no permissions and no accountability. Most enterprise AI activity today is still here.
2. The agent. Wrap a model in tools, instructions and a goal and it can take bounded actions. Most organisations have taken this step, function by function. An agent is still one cell, and it sees what its function sees.
3. The connected organisation. Link the functions so that context, work and evidence move between them. People remain the decision-makers, agents act inside boundaries, and the record of what happened is a by-product of the work. This is what we mean by neural intelligence: intelligence that comes from the connections as much as from the components.




Figure 2. From neural networks to neural intelligence: a model, then an agent, then a connected organisation in which tools, data, AI and people work as one.

Biology gives a useful, if imperfect, comparison. A single neuron is simple, and what we call intelligence emerges from how many simple parts are connected and what signals travel between them. We do not claim an enterprise is a brain. We claim the comparison guides better than treating each function as a machine to optimise alone. It also explains an uncomfortable fact about AI spending. A faster coding assistant shortens one step, but if requirements arrive unclear and review queues are long, the organisation waits at the same places as before.

4 Four things, working as one

ANI has four parts. Each is a project in its own right and each is familiar. What is new is designing them as one thing, with every layer pointing inward at what the business needs to do next.




Figure 3. The four parts of ANI: your tools, your data connected, AI on top, and our people inside. Any one is a project. Together they behave like a nervous system.

Your tools

The tools you already run stay: your CRM, finance platform, HR system, delivery tooling and the long tail of smaller applications. Each represents years of configuration and trained habits. The first principle of ANI is additive, so whatever we build must work across several vendors, leave systems of record as the place where the official record lives, and not require a business to change its core applications before it sees any benefit.

Your data, connected

In most enterprises the same customer, order or person is described differently in each tool. The usual answer is to copy everything into one store under one schema. We distrust that approach: it is slow to maintain, it goes stale, and it creates a second copy of sensitive data that somebody has to secure.

We map instead of copying. The shared layer is a set of agreed meanings, what a customer, a contract, a case or an approval is in this business, and how each tool's fields correspond to them. The data stays where it is and the meaning is shared. It is scoped to the workflows that matter and grown from there. Meaning is what lets a person ask how long it takes from signature to cash and get an answer two functions agree on, and it is what we expect to let an agent behave sensibly. The ETA paper describes the mechanism.

AI on top

With tools kept and data connected through shared meaning, AI has something to work with: agents that read across functions, flag what needs attention, draft work for a person to confirm and carry out defined, low-consequence steps. We put it on top for a reason. AI underneath everything, free to rewrite the data beneath it, is hard to trust and hard to audit. AI on top, reading widely and acting narrowly through a few controlled routes, can be given more responsibility safely over time.

Our people, inside

The fourth part is the one technology vendors usually leave out. The hard part of connecting an organisation is rarely technical. It is deciding what a customer is, who approves what, which workflow comes first, and how to redesign a process rather than automate a bad one. Ontology workshops, process redesign and change management are organisational work, and no software release replaces them. So our people work inside the client's organisation, in small senior teams, changing how the work gets done as well as building tools for it.

Inside our own walls, we train people on a ladder that starts with working safely with AI tools and climbs to designing the harnesses around them, graded on real work. We check AI output adversarially: state the verdict, argue the strongest case against it, rank the findings by how fatal they are, and say which claims were measured and which were modelled. Agents write drafts and code. People stay accountable. That is the whole job of this fourth part, and it is why we describe the offer as agentic solutions, powered by humans.

5 The nervous system, and where people stay in charge

We use the nervous-system image as a thinking tool for deciding what should happen automatically and what should not. It gives three ideas a useful vocabulary. It is a metaphor and not an architecture, and the figure shows the three ideas around the outcome and not a set of components.




Figure 4. Signals, reflexes and judgement around the outcome. Signals come in from tools and systems, reflexes are governed, pre-approved automatic responses, and judgement is people deciding at named gates.

Three ideas borrowed from the nervous system, and what each means in an enterprise.		
Idea	In the body	In an enterprise
Signals	Information travelling between parts	What actually happened: a deal signed, a ticket reopened, a story waiting for review, with who, when and what it cost. Recorded as a by-product of the work.
Reflexes	Fast, automatic responses that do not wait for the brain	Low-consequence, well-understood steps an agent may do without asking, inside an approved boundary, with every action logged and reversible.
Judgement	Deliberate decisions	Anything with real consequence, ambiguity or a person's interests at stake. A named person decides, and the decision is recorded against what they saw.

Most of the quality of an ANI system comes from drawing the line between reflex and judgement correctly, and from moving it carefully. That is a governance question first and a technical one second.

Named human gates

In our AI-led delivery work nothing is built until it has passed four named human approval gates: G1 Requirement, G2 Requirement approval, G3 Design approval and G4 Story approval. At every gate the reviewer and approver are named people on the client side, and not a function or a shared mailbox. A person merges every change, and the client alone decides when user acceptance is done.




Figure 5. Four named human gates before build, a person who merges every change, and client-only sign-off of user acceptance.

A few standing rules sit behind the gates. An agent never merges its own work. The model has no write path to a system of record. Approvals bind to the exact thing approved. These rules are structural: we do not rely on a model remembering them, and we do not claim to have solved prompt injection. The aim is that a successful attack on a model cannot reach anything that matters, because the route to anything that matters passes through a person. Our article [Rules our agents work under](#) sets them out in full.

Earned autonomy

The line between reflex and judgement should not be fixed forever, or drawn once for a whole agent. We think autonomy should be earned per kind of step, through four tiers: shadow, supervised, delegated and autonomous. Promotion needs written evidence and a named owner. Demotion is automatic when the evidence slips. The ETA white paper describes the tiers and criteria, and they are our design and not an industry standard.

One part of the design is worth stating here, because it surprises people. An override rate trending to zero looks like success. We treat it as an alarm, because it can mean the reviewer stopped reading, and a control nobody exercises is no longer a control. We write about it in [An override rate of zero is an alarm](#).

6 What this changes for an enterprise

Measurement comes before automation

Most programmes start by automating something and hope to measure afterwards. We think that is backwards. Without a baseline, nobody can say what the automation changed, and every later claim about savings is open to dispute. The first job is to count: how long the outcome takes end to end, how many people touch it, how many hand-offs there are, and what share of steps needs a person. Most organisations cannot answer those questions for any cross-functional workflow. A value case that survives a finance review leads with cycle time, error reduction and released capacity. We do not sell headcount reduction, and we say so in advance, because it changes what a client can safely measure and who is willing to take part.

Governance moves into the tools

Governance in most organisations lives in meetings and in people's heads. That works while the volume of change is human-sized. It stops working when software, drafts and decisions arrive faster than people can read them. We put governance where it can be enforced: in named gates, in workflow rules that block incomplete work from advancing, and in the pipeline that produces evidence. Review depth follows risk, so the effect is less ceremony and stronger control together.

The shape of delivery teams

When agents do the volume work, the unit of delivery changes size. In our AI-led engagements a pod is two people: a Forward Deployed Engineer, who oversees the work and the plans, and a Product Engineer, who prepares the next piece of work. Together they run three Harnesses covering requirements, engineering and quality. People move from producing artefacts to approving, overseeing and preparing. We describe the pod in [A two-person pod, two days](#).

None of this means conventional delivery is finished. Many engagements are better run by established teams, and many run best as a mix. A pod can take the backlog that suits it while an existing team continues with the rest, under the same client governance and the same approvers. The client holds the named approval roles and owns acceptance either way. A hybrid is where most enterprises will spend the next few years, and the route to a connected organisation runs through what already exists.

7 How Anitech puts ANI to work

Anitech is an IT services company that has delivered enterprise software since 2006. We put ANI to work in three ways, and say how far each has got.

Harnesses

Our [Agentic Harness Orchestration](#) products are three Harnesses, one for each stage of delivery. A Harness is the governed process around a model: what the agent may do, how its work is checked and what is recorded.

Status of the three Harnesses, October 2026.		
Harness	What it does	Status today
Requirements	Researches a title-only requirement across a client's own sources and drafts it with a citation for every claim. Unknowns become open questions.	Live with one client.
Engineering	Plans and builds stories under human approval, with quality gates that pass only on executed evidence.	Built. Going into first client use in October 2026.
Quality	Generates and runs test cases from a story.	In development. A QA agent is being trialled.

This does not cover the entire software lifecycle today. Only requirements is live, so the accurate description is a plan with one working part.

ETA, ANI in product form

[ETA](#) is the version of ANI you can deploy. It is designed to sit on top of the estate a client already has, map the tools to shared meaning without copying data, put a governed harness into each function, and record what it changes. It begins by measuring and only then automates. ETA is a design with no pilot results, and we publish none. The [ETA white paper](#) sets out its rules, architecture and limits.

AI-led delivery services

The third route is delivery. Our application development, dedicated teams and AI consulting services use the Harnesses and the operating model in this paper where they fit a client's work, under the client's governance and alongside conventional delivery where that fits better. They are how the fourth part of ANI, our people inside, reaches a client.

One pilot result

☰ THE PILOT, IN ONE SENTENCE

In a two-week pilot in August 2026 with a North American education-technology client, a two-person pod took 49 user stories to opened merge requests in about 48 elapsed hours, including the human approval gates, after seven working days of setup.

The figures are self-reported from the pilot's run logs (reference 2). The merge requests were opened and had not yet been merged, and quality outcomes have not yet been measured. It is one client, one scope and one pilot, and the seven days of setup are part of the cost. We do not generalise it into a multiplier, because nothing we have supports one. It shows that a small pod with named gates could move a large batch of work to the point where a person reviews it, in a short time, with the approvals intact. It does not show that the work was good. The next stage is to find out.

8 What we don't claim

A point of view that cannot say where it might be wrong is marketing. These are the limits as we see them today.

- ANI is not proven. It is a view of how enterprise IT is likely to develop, supported by reasoning, a small body of our own delivery experience and public research that points the same way.
- We have no ETA pilot results. ETA is a design with a delivery method behind it.
- Our one pilot is self-reported, from one client, with merge requests opened and not merged and quality not yet measured. It is a signal and not a finding.
- Only one of our three Harnesses is live, with one client. The Engineering Harness goes into first client use in October 2026, and the Quality Harness is still in development.
- Joining a single case across several systems that share no common identifier is hard, and we have not shown we can do it without people mapping the keys by hand.
- We have not eliminated risk from agents. Prompt injection is a real class of attack, and our design limits what a successful attack can reach. That is mitigation by structure and not a guarantee.
- ANI does not fit every workflow. Work with low volume, no reliable signal of whether a result was right, or consequences that cannot be reversed is a poor place to start. Decisions about people are never where we begin.
- We do not promise cost reductions, headcount changes or a date by which any of this reaches your organisation.

9 Where to start

You do not need to buy anything to test whether this argument applies to you. Each of these steps is useful even if you go no further.

1. Ask the cheapest question first. Ask three senior colleagues how many people touch one deal from signature to cash, and how long it takes. If they answer from data and agree with each other, you are ahead of most organisations. If not, you have found the missing number.
2. Choose one cross-functional outcome with a named owner. Pick a workflow with enough volume to measure, a reliable signal of whether it went right, inputs that can be retrieved and actions that can be reversed or bounded. Avoid decisions about people.
3. Count before you change. Record the elapsed time, the people involved, the hand-offs and the share of steps that needs a person, for a few weeks, so the next investment can be judged against a baseline.
4. Write the boundary down. Name who approves what, which systems an agent may read, and which it may never write to.
5. Trial AI-led delivery on a contained backlog alongside your existing team, under your own approvers, with acceptance staying with you.

We are happy to help with any of these, including the first, which needs no engineering at all. The companion paper on ETA describes how we would measure a workflow. If you would rather talk it through with the two of us first, [get in touch](#).

FIGURES

[Figure 1](#). Siloed functions, each with its own tools and agents, against the same functions connected so that work and context move between them with people in control.

[Figure 2](#). From neural networks to neural intelligence: a model, then an agent, then a connected organisation in which tools, data, AI and people work as one.

[Figure 3](#). The four parts of ANI: your tools, your data connected, AI on top, and our people inside. Any one is a project. Together they behave like a nervous system.

[Figure 4](#). Signals, reflexes and judgement around the outcome. Signals come in from tools and systems, reflexes are governed, pre-approved automatic responses, and judgement is people deciding at named gates.

[Figure 5](#). Four named human gates before build, a person who merges every change, and client-only sign-off of user acceptance.

REFERENCES

- [1] Lei, F. et al. Spider 2.0: Evaluating Language Models on Real-World Enterprise Text-to-SQL Workflows. ICLR 2025. arXiv:2411.07763. <https://arxiv.org/abs/2411.07763>
- [2] Anitech Solutions. Pilot result: two-person pod, 49 user stories to opened merge requests. Self-reported from the pilot's run logs, August 2026. Anonymised at the client's request. Merge requests opened, not yet merged. Quality outcomes not yet measured.
- [3] Anitech Solutions. A governed operating layer for work that crosses functions: the Enterprise Technology Accelerator (ETA). October 2026. Companion to this paper. </white-papers/eta>

ABOUT THE AUTHORS



Andrew Nolan

Chief Executive Officer

Owens the client relationship, the outcome and the standard the work is held to.



Jack Dinneen

Chief Product and Technology Officer

Owens product and technology: the Harnesses, ETA and what actually ships.

How this paper was written. Drafted with AI assistance from Anitech's internal working papers and edited by Jack Dinneen and Andrew Nolan.