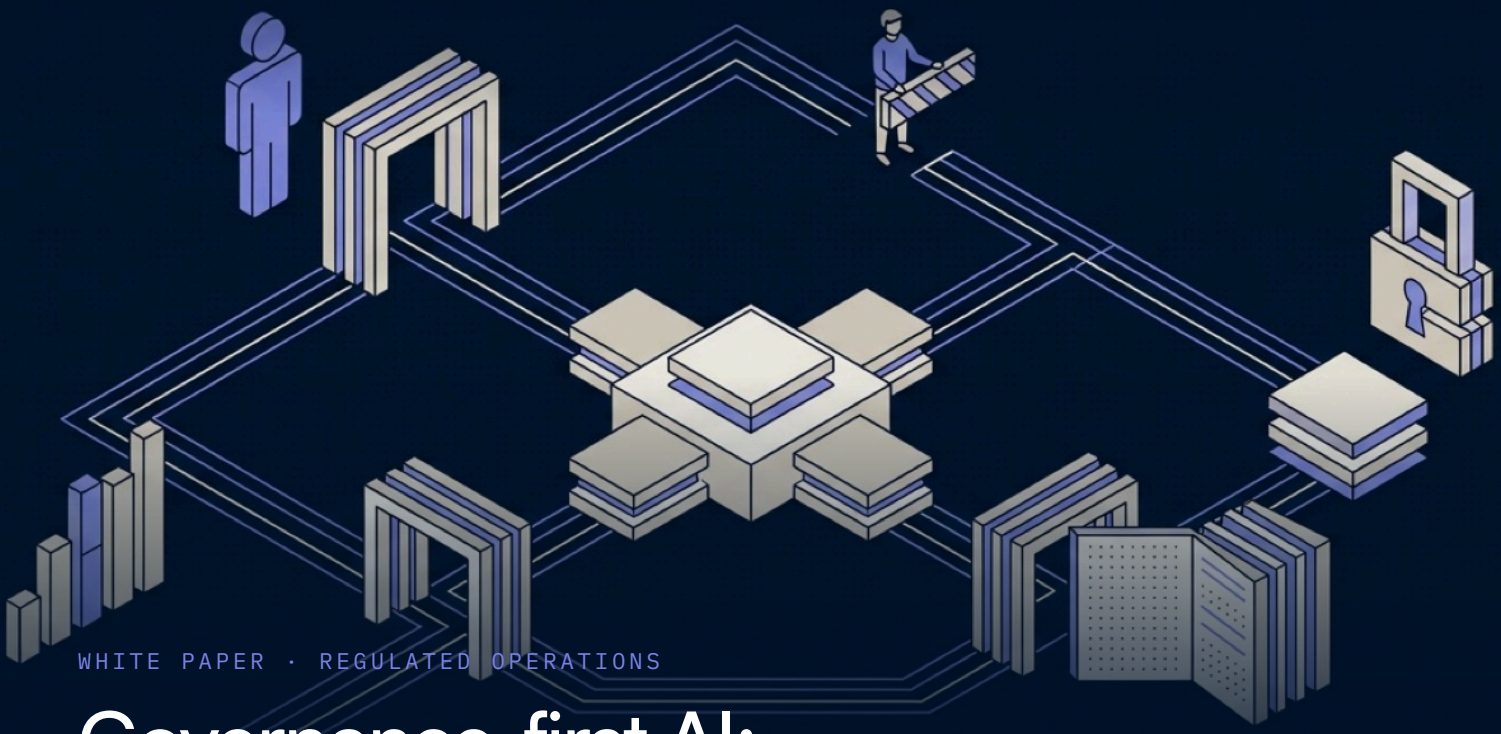

WHITE PAPER · REGULATED OPERATIONS

Governance-first AI: making automation safe in regulated operations.

The model was never the hard part. The value is safely automating the closed-system last mile of regulated work, with a person approving every critical action and a full audit trail behind it. This paper sets out the controls that make that safe to do.

FOR INSURANCE · FINANCIAL SERVICES · LEGAL & INSOLVENCY · LOGISTICS · GOVERNMENT · CONSTRUCTION

HELLO@NORTHSTACK.DIGITAL

THE ACCOUNTABILITY GAP

"The AI did it" is not an answer you can give a regulator, a client, or a court.

Most AI tooling is built to remove the human from the loop. In a regulated operation that instinct is backwards. The constraint was never the model's capability. It was the work no model can safely reach on its own: the MFA-gated portals, the legacy systems with no application programming interface (API), the inboxes and the spreadsheets where the real obligations live.

When a system acts in that environment without a person standing behind the action, the organisation has not removed risk. It has moved the risk somewhere it cannot see and cannot answer for. Autonomy without accountability is not efficiency. It is a liability that has not been priced yet.

A regulator asks who authorised a decision and on what basis. A client asks why their case was handled the way it was. A court asks for the record. None of these accept "the model decided" as a defence. Each of them expects a named person, a reason, and evidence.

What a black box costs you

01 An action you cannot trace to a person you cannot defend, however good the outcome looked at the time.

02 A decision you cannot explain is a decision you cannot stand behind under scrutiny.

03 A process you cannot audit is a process you have to take on trust, which is the opposite of a control.

04 An error you cannot reconstruct is an error you are doomed to repeat, because nothing recorded how it happened.

The fix is not less automation. It is automation built so that accountability survives it. Every action is attributable, every decision explainable, every step recoverable. That is what we mean by governance-first: the controls come before the build, not after the incident.

THREE LINES THAT NEVER MOVE

Some actions are never worth automating, at any efficiency.

Most of our governance is calibrated to the work in front of us. Three lines are not. They hold on every engagement, in every sector, regardless of what a client asks for or what a model can do.

RED LINE 01

We never automate payments.

A machine never moves money on its own authority. The financial loss from a single mistaken or manipulated payment dwarfs any time saved across thousands of correct ones. The asymmetry is permanent, so the line is permanent.

RED LINE 02

We never automate bank-detail changes.

Changing where money is sent is the single most exploited fraud vector in regulated operations. Automating it hands an attacker a fast, quiet, unsupervised path to the payout. A human verifies every such change.

RED LINE 03

A hard human gate before any regulator submission.

Anything filed with a regulator carries the organisation's name and its liability. A named person reviews and authorises before it leaves. The system prepares the submission. A person signs it.

Why they never move

These are not default settings waiting to be tuned, and they are not negotiated per client. They sit outside the commercial conversation on purpose. A red line you can buy your way past is not a red line. It is a price.

The reasoning is the same in all three cases: the worst outcome is catastrophic and irreversible, the time saved is marginal by comparison, and a competent person in the loop removes nearly all of the risk for almost none of the cost. When the maths is that lopsided, the answer does not change with the engagement. We say no in writing, and we hold it.

THE CONTROL SET

Six controls turn a capable system into an accountable one.

The red lines say what we will not automate. These six controls govern everything we do automate. They are designed in from the first rung, not bolted on after a system is live, because a control added after an incident is already too late.

CONTROL 01

Human-in-the-loop on high-stakes

Any action with material consequence stops for a person before it commits. The system does the preparation and presents the decision. The human makes the call. The threshold for "high-stakes" is set with the client, in writing.

CONTROL 02

Per-action approval for agents

Where an agent carries out a sequence of steps, it does not get a blanket mandate. Each consequential action is approved on its own merits, so authority is never delegated wholesale to a process running unattended.

CONTROL 03

Field-level confidence scores

Every extracted or generated value carries a confidence score, field by field. Low-confidence fields route to a person automatically. Reviewers spend their attention where the system is unsure, not on what it already knows.

CONTROL 04

Complete, exportable audit trail

Every action, input, decision and approval is logged and exportable. When a regulator, client or court asks what happened and who authorised it, the record is already there, complete, and in a form you can hand over.

CONTROL 05

Role-based access control

People and systems see and do only what their role permits. Authority to approve a high-stakes action is held by the people accountable for it, and that boundary is enforced by the system rather than by convention.

CONTROL 06

Failure and incident playbook

Every deployment ships with a defined response for when something goes wrong: how it is detected, who is notified, how it is contained, and how it is reviewed afterward. The plan exists before the incident, not during it.

FROM PRINCIPLE TO PRACTICE

Governance you can verify, not governance we assert.

Audit-first, before any build

Every engagement opens with a fixed-fee operations review. We map the operation, score each workflow, and hand back a costed, sequenced plan, including the work we tell you not to automate. No build is priced before the hard part is proven on your real data. We say no in writing.

A human role at every rung

We climb the North Stack Ladder one rung at a time. The human role does not disappear as the system matures. It moves up: from doing the work, to checking it, to supervising the system that does it.

`Mapped → Assisted → Automated → Integrated → Orchestrated`

Kept systems, not rip-and-replace

We automate around the systems you already run, including the ones with no API. We do not demand a platform migration as the price of progress. Marshall Peters kept their insolvency CRM and still cut document-review time by about 70%.

Proof

~900 operational minutes/day returned to one client (≈£120,000/year)	**~60%** less manual effort — PMD Finance
>50% faster quotes, live in ~6 weeks — PMD Finance	**~70%** less document-review time — Marshall Peters
~45% faster reporting — Marshall Peters	**~3 days** faster board actions — Norcis

Further clients include Ziani and Hippo. Delivered by a UK studio of five: counsel formerly with the UK Parliament, operators from Lloyds, HSBC and the Association of British Insurers (ABI), applied-AI research with two patents, and a practising medical doctor.

START WITH THE AUDIT

Prove the hard part on your real data before a single line is built.

HELLO@NORTHSTACK.DIGITAL