GLOBAL TECH UPLIFT
Linking Intelligence

DELIVERABLE SNAPSHOT · CYBER AND MISSION

AI strategy, governance and pilot readiness

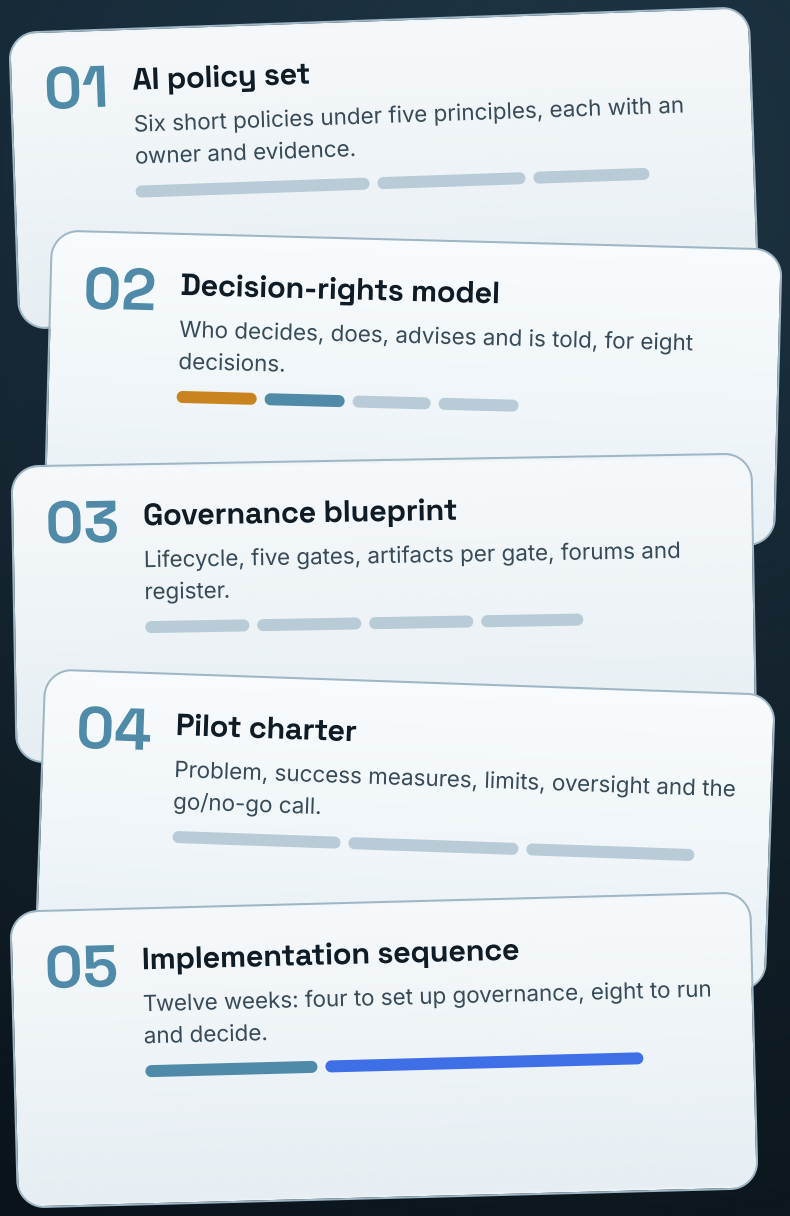
Build the workflow, not just the model.

Cyber and mission edition

Alert triage · DNS and proxy logs · Threat intelligence · Least privilege · Incident response · OT boundaries

What a GTU engagement of two to four weeks leaves behind: the policies, decision rights, operating blueprint, pilot charter and sequence an organization needs to start AI pilots with clear controls, then decide on evidence.

01 AI policy set
02 Decision-rights model
03 Governance blueprint
04 Pilot charter
05 Implementation sequence
\+ Templates you keep



Illustrative sample prepared for Ridgeline Utility Cooperative, the security operations center of a fictional regional utility. Names, figures and dates show the format and depth of the deliverables, not a client's results.

THE ENGAGEMENT

At a glance

A short, principal-led engagement that turns scattered AI activity into a governed pipeline, and ends with an approved first pilot.

WHY ORGANIZATIONS ASK FOR IT

- Security tools ship AI faster than policy follows, from alert triage to auto-response.
- Nobody clearly owns the go or no-go decision for an AI use case.
- Pilots start from a demo, without agreed measures of success.

WHAT CHANGES IN TWO TO FOUR WEEKS

- Every known AI use has an owner, a risk tier and a place in the register.
- Leadership approves pilots against written criteria, at a regular forum.
- The first pilot starts with a charter and a decision date.

WHO TAKES PART

Executive sponsor · AI governance board · Business owners · Data owners · Model owners · Risk and compliance · Security · SOC lead · Detection engineering

HOW THE WEEKS RUN

WEEK 1
Discover
Interviews, AI-use inventory, data map

WEEK 2
Design
Principles, policies, decision rights, risk tiers

WEEK 3
Stand up
Board charter, intake form, register, templates

WEEK 4
Charter and decide
First pilot chartered and approved

WEEKS 5 TO 12 · OPTIONAL SUPPORT
Run the pilot and decide on evidence
Configure, test, go live in stages, evaluate against the charter, make the go/no-go call

2 to 4 weeks	8 to 12	3 workshops	1 board session	About 6 hours
Engagement length	Stakeholder interviews	Policies, decision rights, pilot design	Approves the governance model and first pilot	Time asked of each stakeholder

WHAT YOU RECEIVE

AI policy set
Principles plus six one-page policies with owners and evidence.
Page 3

Decision-rights model
RACI for eight decisions and escalation by risk tier.
Page 4

Governance blueprint
Lifecycle, gates, artifacts, forums and model register.
Page 5

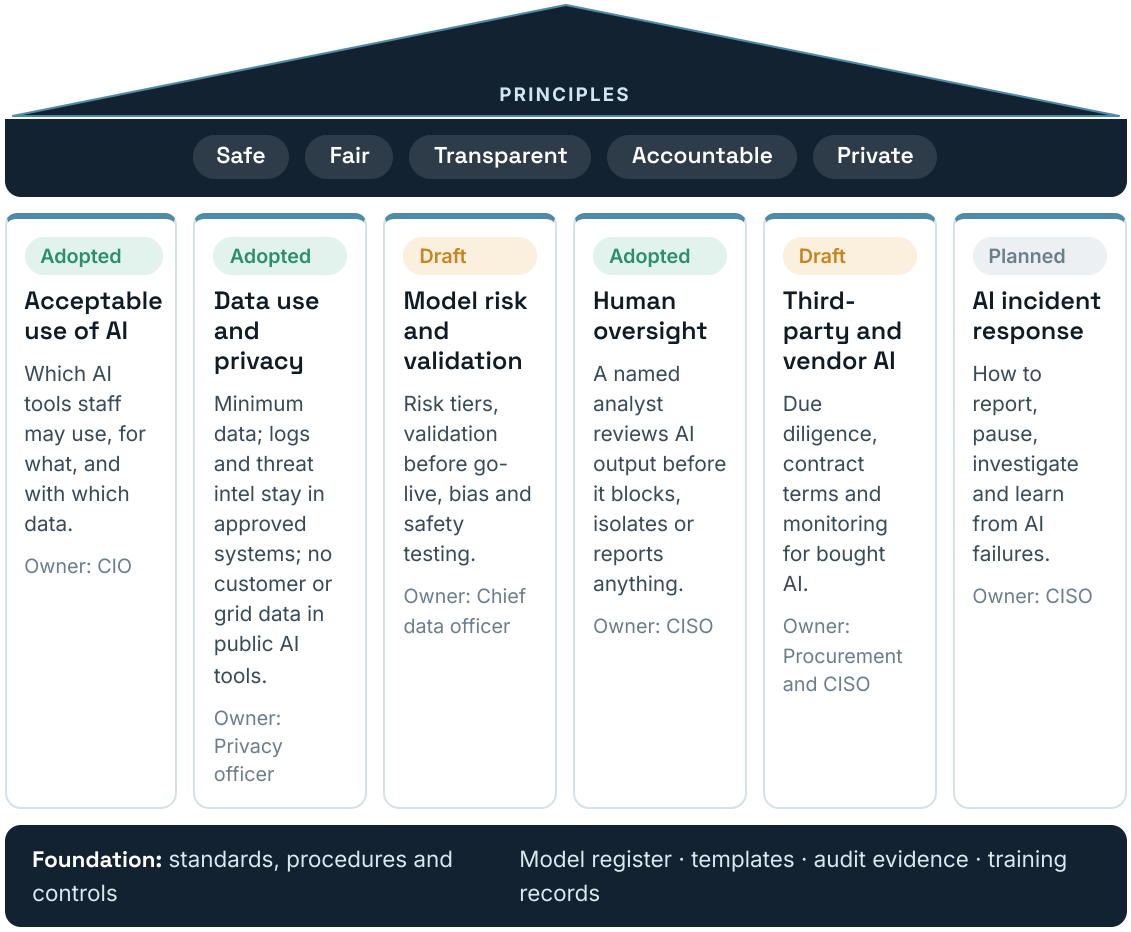
Pilot charter
One page a pilot cannot start without.
Page 6

Implementation sequence
The 12-week order of work and its milestones.
Page 7

DELIVERABLE 01

01 AI policy set

Six short policies under one set of principles. Each names an owner, a review date and the evidence that shows it is followed, so the policy works in practice and not only on paper.



ONE-PAGE POLICY, EXAMPLE

Human oversight policy

PURPOSE	Keep people accountable for decisions that AI informs.
APPLIES TO	Any AI that informs detection, response actions or reports to regulators and partners.
RULES	• A named analyst checks AI output before any block, isolation or report. • Staff can override any AI output; overrides are logged. • Tier 3 uses need two-person review during the pilot. • Oversight is tested before go-live, not after.
OWNER	Chief information security officer; reviewed every six months.
EVIDENCE	Reviewer roster, override log, quarterly audit sample of 50 cases.
STATUS	ADOPTED by the AI governance board, week 3.

DELIVERABLE 02

02 Decision-rights model

Who decides, who does the work, who is consulted and who is told, for the eight decisions that most often stall AI programs. Each decision has exactly one accountable role.

A Accountable: makes the call · R Responsible: does the work · C Consulted before · I Informed after

Decision	Executive sponsor	Governance board	Business owner	Data owner	Model owner	Risk and compliance	Security	SOC lead
Accept a use case into the pipeline (gate G0)	I	A	R	C	C	C	I	C
Assign the risk tier tier 1, 2 or 3	I	A	C	C	C	R	C	C
Approve access to data scope, de-identification	I	I	C	A	R	C	C	I
Start a pilot charter approved (gate G1)	C	A	R	I	C	C	C	C
Approve go-live gate G3	C	A	R	I	R	C	C	C
Pause or roll back any time, any tier	I	I	C	I	R	C	C	A
Buy an AI product vendor or embedded AI	I	A	R	C	C	C	R	C
Tell analysts and partners what AI does and does not do	A	C	R	I	C	C	I	C

ESCALATION BY RISK TIER

Tier 1 · Low — Decision in 5 days

Internal productivity, no personal data, easy to undo.

Approved by: business owner, logged in the register

Tier 2 · Moderate — Decision in 15 days

Staff-facing, uses personal data, a person reviews every output.

Approved by: AI governance board

Tier 3 · High — Decision in 30 days

Can block traffic, isolate systems or affect operations.

Approved by: board and executive sponsor, with independent validation

The SOC lead can pause any AI use at any tier; the pause is reviewed by the board within 48 hours.

DELIVERABLE 03

03 Governance blueprint

How a use case moves from idea to retirement: the stages, the five gates and what each gate needs, plus the forums and register that keep it running.

Intake — Use-case brief
Triage — Risk tier set
Design — Data and workflow
Validate — Test and pilot
Deploy — Go live in stages
Monitor — Drift, safety, value
Review — Continue or retire

G0 · Accept

After intake

- Use-case brief and named owner
- Expected value and users
- Entry in the model register

G1 · Design review

Before build

- Risk tier and impact assessment
- Data sheet and access approval
- Pilot charter

G2 · Validation

Before users

- Accuracy, bias and safety tests
- Oversight plan tested
- Security review

G3 · Go-live

Before scale

- Model card and monitoring plan
- Training completed
- Rollback plan rehearsed

G4 · Periodic review

Every 6 to 12 months

- Monitoring report and incidents
- Value against the charter
- Continue, change or retire

FORUMS

AI governance board · monthly · chaired by the executive sponsor. Approves tier 2 and 3 uses, owns the policies, reviews the register.

AI working group · weekly · triage, templates, register upkeep, support to owners.

Incident path · report within 24 hours; the SOC lead can pause; board review within 48 hours.

MODEL REGISTER (EXCERPT)

USE CASE	TIER	STAGE	OWNER	NEXT REVIEW
DNS anomaly alert triage	2	Validate (pilot)	SOC	Jan 2027
Phishing-report triage	2	Design	SOC	Dec 2026
Vendor auto-isolation playbook	3	Monitor	CISO office	Nov 2026
Meeting notes summarizer	1	Deploy	IT	Mar 2027

BOARD DASHBOARD

100% Known AI uses in the register

12 days Median intake to decision

0 Open incidents

1 Use past its review date

DELIVERABLE 04

04 Pilot charter

One page a pilot cannot start without: the problem, how success is measured, the limits, who oversees it and the decision at the end.

Pilot: DNS alert triage assistant | Tier 2 | Sponsor: CISO | Owner: Detection engineering manager | Weeks 5 to 12 | Approved at gate G1, week 4

1 Problem and users

Analysts see about 3,000 DNS and proxy alerts a week. Most are noise, and slow beaconing hides among them.

Users: 9 analysts on three shifts.

2 Workflow, today and in the pilot

TODAY Alert → Queue → Read all → Close or escalate

PILOT Alert → Score, group and explain, draft ticket → Analyst reviews → Close, escalate or respond

No block or isolation happens on the AI's word.

3 Success measures

Median time to triage	42 min	**≤ 15 min**
High-severity alerts triaged within 15 min	58%	**≥ 90%**
Analyst hours on triage a week	110	**≤ 70**
Tickets accepted with light edits	new	**≥ 60%**
Confirmed threats missed	0	**0, hard stop**

4 Scope

- **In:** DNS and proxy alerts on the corporate network.
- **Out:** operational technology networks, automated blocking.

5 Data and access

- Six months of alerts, tickets and outcomes.
- Runs inside the SIEM tenant; no logs leave approved systems.
- Least-privilege access, reviewed at G2.

6 Risks and mitigations

- Missed slow beacon: long-window check plus analyst review.
- Poisoned inputs: only trusted feeds reach the model.
- Over-trust: evidence and confidence shown with every score.

7 Human oversight

- Every escalation is reviewed by an analyst.
- The SOC lead samples 50 closed alerts a week.
- Pause authority: SOC lead.

8 Go / no-go at week 12

- **Go** if triage ≤ 15 min, high-severity alerts within 15 min ≥ 90%, no confirmed threat missed and analyst rating ≥ 4 of 5.
- **Stop** if any confirmed threat is missed, or fewer than 30% of tickets are usable.

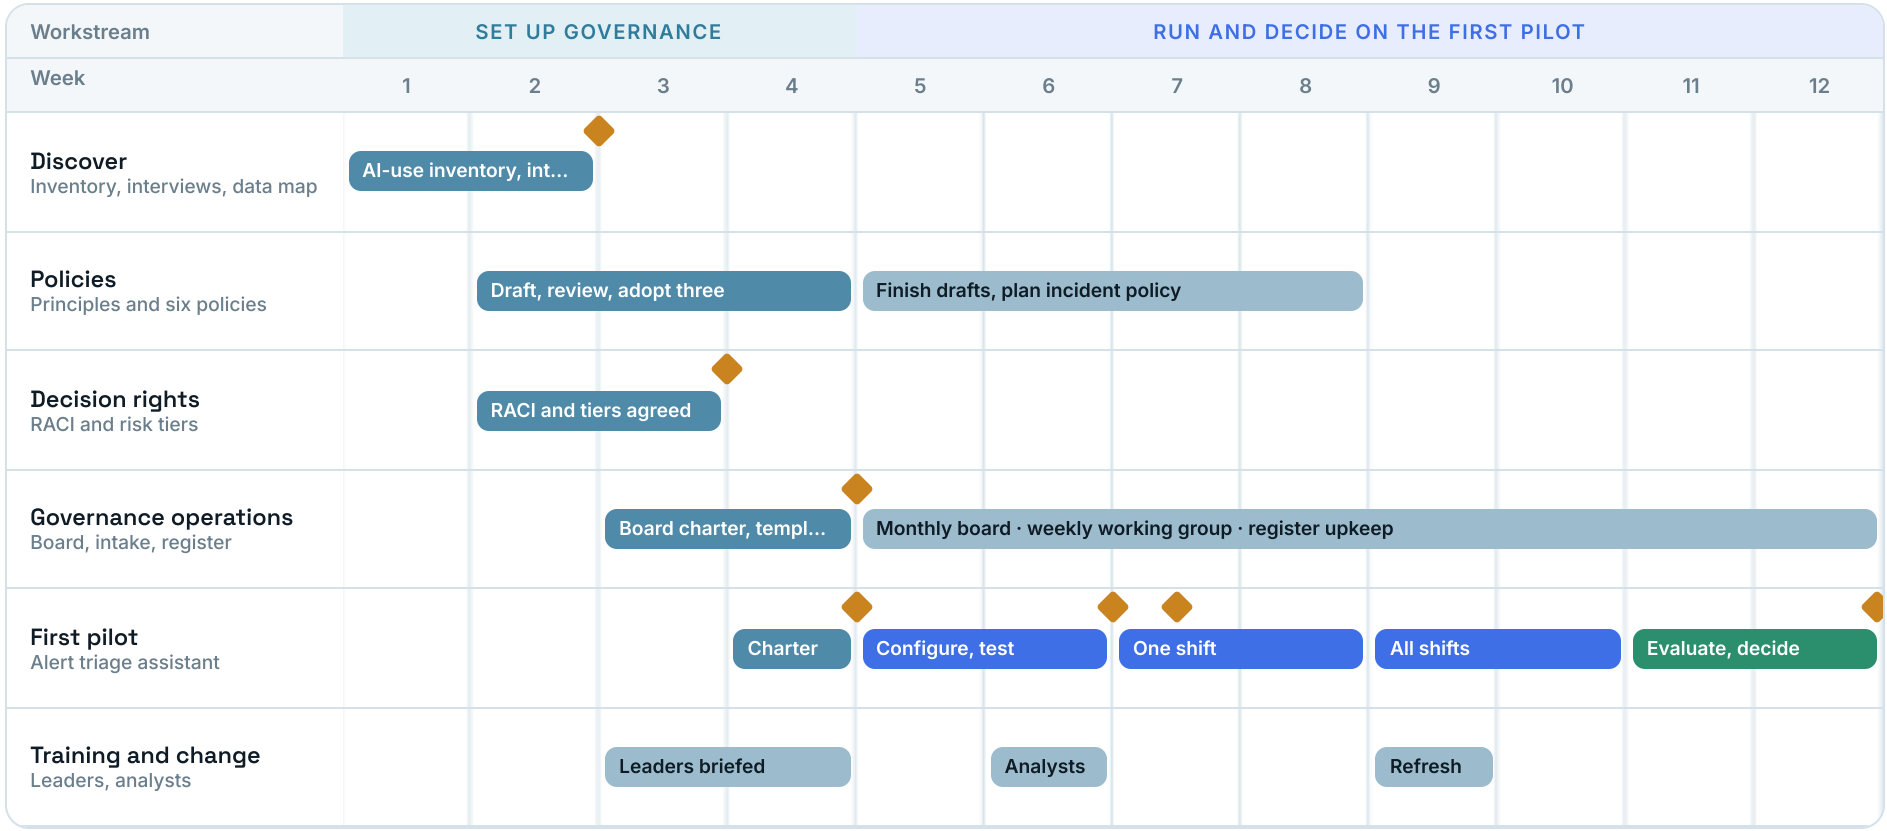
9 Team and timeline

- Weeks 5 to 6: configure and test (G2).
- Weeks 7 to 8: one shift (G3).
- Weeks 9 to 10: all shifts.
- Weeks 11 to 12: evaluate and decide (G4).

DELIVERABLE 05

05 Implementation sequence

The order of work over twelve weeks: four weeks to set governance up, eight weeks to run the first pilot and decide on evidence. Amber diamonds are decision points.



Decision points: inventory complete (wk 2), decision rights approved (wk 3), board launched and charter approved, G1 (wk 4), validation G2 (wk 6), go-live G3 (wk 7), go/no-go G4 (wk 12)

Every known AI use in one register

With an owner, a risk tier and a review date.

A first pilot decided on evidence

Against the measures written into its charter at week 4.

The next two pilots chartered

Queued through the same gates, with owners and decision dates.

WHAT YOU KEEP

Files, templates and next steps

Everything is yours to edit and reuse. Templates come pre-filled from the engagement, so the next use case goes faster.

AI policy set
Principles and six policies, one page each, with owners and evidence.

Document

Decision-rights matrix
RACI for eight decisions, risk-tier escalation rules.

Spreadsheet

Governance blueprint
Lifecycle and gates, forum charters, model register.

Slides and sheet

Templates
Intake form, risk-tier worksheet, impact assessment, model card, monitoring plan.

Documents

Pilot charter
One-page charter plus a measures sheet with baselines.

Document and sheet

Implementation plan
The 12-week sequence with milestones and owners.

Plan

Principal-led
Dr. Gaurav Tuli leads the work directly, from interviews to the board session.

Governance first
Risk controls and decision rights are agreed before anything scales.

Built with your people
Policies and gates are written with the teams who will run them.

Measured
Every pilot carries its own measures and a decision date.

Start with a free 30-minute discovery call

Build the workflow, not just the model.

globaltechuplift.com · info@gtu.li

This snapshot is an illustrative sample. The organization, people, figures and dates are fictional and show the format of GTU deliverables. Real deliverables are built with your teams and your data. Nothing here is legal advice.