NET-TECH CONSULTING

Real People, Real Protection, Right Now

AI Governance Fast-Track Sprint

Triage-first governance + security for teams defining an AI project

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Why AI Governance Fast-Track exists

Clarity + control before AI becomes operational dependency.

AI is moving faster than governance

AI features ship inside SaaS tools and copilots. If you can't answer "what's in scope and where data goes," you're already behind.

Risk is usually discovered late

Most teams find the gaps after an incident, customer questionnaire, or insurer asks for proof (logs, approvals, controls).

"Defining an AI project" is ambiguous

Tool adoption, internal automation, customer-facing features, and build/model pipelines need different guardrails.

Budget windows close quickly

If planning is happening now, a fast, structured sprint lets you lock scope, risk tier, and a 90-day execution plan.

NET-TECH CONSULTING
Real People, Real Protection, Right Now

What the sprint is

A triage-first assessment that turns “AI project” plans into a bounded, actionable program.

AI Fast-Track converts ambiguity into decision-grade outputs:

- Clear use-case + system boundary
- Data classification + risk tier
- Mandatory gates (Go / Conditional Go / No-Go)
- The right governance + security workstream (module selection)
- A 90-day execution plan with owners + evidence

Operator-built for mid-market teams

Framework-agnostic

Control-based work you can map to NIST AI RMF / ISO 42001 topics if your stakeholders require it.

Evidence-first

Outputs are designed to be reusable for executives, auditors, insurers, and customers.

Minimal lift

We do the structure + artifacts. You provide stakeholder input and access to relevant settings/logs (read-only is fine).

NET-TECH CONSULTING
Real People, Real Protection, Right Now

What you walk away with

Concrete artifacts you can circulate internally and execute against.

Core deliverables

- AI Project Definition Packet (scope, boundaries, data classes, initial risk tier)
- AI Governance 30-Point Assessment Checklist (baseline coverage)
- Risk Scorecard Summary (0–3 scoring across domains)
- AI Control / Gap Matrix (mapping-ready)
- 90-Day AI Action Plan (owners, sequencing, verification evidence)
- Executive Summary (1–2 pages for leadership/board updates)

Optional add-ons (if needed)

- AI Policy Starter Pack (AI Acceptable Use + Secure AI Use Standard + data handling rules)
- AI Vendor Due Diligence addendum + contract clause checklist
- AI Incident Response addendum + tabletop exercise
- Quarterly governance reporting + inventory refresh

How it works

Stage-gated flow: define → assess → route → plan.

Stage 0: Triage

60–90 min workshop
Use-case, data, integrations
Risk tier + gates

Stage 1: Baseline

30-point control review
Scorecard + gap matrix
Evidence requirements

Stage 2: Modules

Pick the right path:
Tools / internal / customer
Build-pipeline / mixed

Stage 3: 90-Day Plan

Sequenced plan with owners
Verification evidence
Exec-ready summary

Typical engagement cadence (fast and structured)

- Triage workshop → within days you get a Project Definition Packet (with Go/Conditional/No-Go gates)
- Baseline assessment + selected module(s) → scorecard, gap matrix, evidence model
- Delivery ends with a 90-day plan that your team can execute (or Net-Tech can support via vGRC/vCISO)

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Triage routes you into the right AI path

Pick the guardrails that match the system boundary and data class.

Path A — Third-Party AI Tools

Copilot / SaaS AI / LLM APIs.
Focus: vendor due diligence, SSO/MFA/RBAC, logging/retention, data leakage controls.

Path B — Internal Automation

Internal workflows & decision support.
Focus: human oversight, QA sampling, change control, operational metrics.

Path C — Customer-Facing AI

External outputs (chatbots/agents/features).
Focus: abuse testing, transparency/escalation, monitoring, AI incident readiness.

Path D — Build/Model Pipeline

RAG, fine-tuning, model hosting.
Focus: data provenance, evaluation plan, versioning, drift monitoring & retraining triggers.

Mixed is normal

Rule: treat each use-case/system boundary separately so controls and evidence stay clean (and scope doesn’t drift).

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Mandatory gates (Fast-Track rules)

Three checks that prevent “risky by default” AI deployments.

1) Data classification gate (WISP-anchored)

If Confidential/Regulated data is in scope, AI use requires explicit approval + risk review + logging. Regulated data may be prohibited without written authorization.

2) Vendor gate

If third-party AI is involved, route through vendor tiering + due diligence + exception handling (and document the decision).

3) Auditability gate

If the tool can’t support baseline admin control + access logging, it can’t be used for sensitive workflows. Proof matters for insurance, audits, and customer reviews.

Outcome: Go / Conditional Go / No-Go decisions — recorded with evidence.

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Evidence model (audit / insurance / customer friendly)

We don’t just recommend controls — we define the proof.

Control area	Typical evidence artifacts
Identity + access	SSO/MFA/RBAC configuration, admin roles, access review notes, audit logs & retention
Data handling	Classification tags, permitted/prohibited usage, retention rules, DLP/egress controls (where applicable)
Vendor risk	Tiering decision, due diligence packet, contract clauses (data use/training/sub-processors), exception approvals
Testing + monitoring	Evaluation plan, misuse/abuse test notes, output QA criteria, monitoring metrics + alert thresholds
Incident readiness	AI incident addendum, tabletop agenda/notes, escalation contacts, follow-up actions

Result: evidence that travels — executive-ready, auditor-ready, and usable for customer security reviews.

NET-TECH CONSULTING
Real People, Real Protection, Right Now

What we need from your team

Minimal lift — just enough input to bound scope and validate controls.

People + context

- Primary stakeholders for the use-case (business + IT/security)
- What decision(s) the AI system will influence
- Any contractual or regulatory constraints (e.g., GLBA, PCI, HIPAA, MSAs)

Tools + data

- List of AI tools/vendors and intended integrations (if known)
- Data types involved (and highest classification)
- Access to relevant admin settings/logging options (read-only is fine)

We can start even if parts are unknown — triage is designed for "we're still defining the AI project."

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Workshop format to get started

Goal: align on scope, risk tier, and the right path module(s) for your AI initiative.

1) Confirm the AI use-case

What decision/workflow is changing, who is impacted, and what success looks like.

2) Define system boundary

Tools, integrations, read/write access, and where data enters and exits.

3) Set gates + pick the module

Data class, vendor/build mix, Go/Conditional/No-Go rules, and the best path (A–D or Mixed).

4) Lock next steps

Owners, timeline, evidence expectations, and the 90-day plan output format.

Pre-work (optional)

If you can share your current AI tools list and a rough use-case description beforehand, we'll come to the workshop with a draft boundary and question set.

NET-TECH CONSULTING
Real People, Real Protection, Right Now

Where MSP services and Zero Trust fit

Governance sets the rules. Controls enforce them. Operations keep them from drifting.

Layer 1 — AI governance + gates (this sprint)

Defines what's allowed, what's blocked, what needs approval, and what evidence must exist.

Layer 2 — Technical controls (security + identity)

Enforce access, logging, and data handling. This is where tools like ThreatLocker can reduce data-leakage and application risk on endpoints.

Layer 3 — Managed operations (MSP / ongoing program)

Keep controls tuned, monitor drift, refresh evidence, and maintain a quarterly governance cadence.

We can start with governance now and align MSP/security services to the 90-day execution plan (if desired).

Next steps

Let’s keep this planning-stage momentum and leave the workshop decision-ready.

Checklist

- ☐ Confirm workshop attendees (business + IT/security)
- ☐ Share AI use-case summary + current AI tools list (if available)
- ☐ Identify highest data class in scope (Public / Internal / Confidential / Regulated)
- ☐ Confirm whether third-party AI vendors are involved (and any contracts already in motion)
- ☐ Align on success criteria and desired evidence format for leadership

Contact: Scott Munden/Zachary Kinder | scott.munden@net-tech.us | zachary.kinder@net-tech.us