AI Governance & Risk Assessment Services

A four-part assessment ladder for the underserved middle, with a free community on-ramp.

Aegis Ethos Consulting offers four assessment services, each built to answer a different question and scoped for organizations without enterprise budgets or in-house data science teams. They are designed to work together: start fast, go deeper only where it is warranted, apply specialized fairness analysis where a specific tool calls for it, and test readiness before adoption rather than after. A free public program, the AI Town Hall, is the community on-ramp. Funder-ready packages and ongoing support retainers are described on page 7.

SERVICE	WHAT IT ANSWERS	INVESTMENT
RISK-90: AI Risk Triage™	Where should we look harder?	**$1,500 flat**
AUD-COMP: Comprehensive AI Governance Assessment	What does a full, evidence-based audit actually find?	**From $10,000**
Disparate Impact Standards Assessment	Does this specific AI tool treat groups differently?	**$4,000 per tool**
READY-AI: AI Readiness & Responsibility Assessment	Can we adopt and sustain AI responsibly?	**$9,500 flat**
The AI Town Hall	What is AI already deciding about us, and what can we do?	**Free**

Pricing effective October 2026. Final fee confirmed in a written scope before work begins.

> *Every engagement is delivered by a governance and risk-management professional with 25 years of banking-grade compliance experience, not an automated score with no one to explain it.*

AI RISK TRIAGE™ · RISK-90

A Fast, Structured Read on Your AI Risk Exposure

WHAT IT DOES

A focused session that screens your organization's AI tools against the NIST AI Risk Management Framework (Govern, Map, Measure, and Manage) to identify where governance gaps exist and where deeper review may be needed. It is built to answer one question fast: where should your organization look harder?

WHAT IT MEASURES

Govern
Accountability, policy, and human oversight for AI-assisted decisions.

Map
Whether your AI systems and their higher-stakes uses are identified and documented.

Measure
Whether outcomes are actually tested for fairness, accuracy, and consistency.

Manage
Whether there is a process to fix issues and monitor systems over time.

WHY IT'S DIFFERENT

- **Fast without being shallow.** A single session, structured around a proven framework, gives you a genuine read on exposure, not a vague conversation.
- **Practitioner-led, not a black box.** Delivered by a governance and risk-management professional with 25 years of banking-grade compliance experience, not an automated score with no one to explain it.
- **A clear next step, not just a score.** Every session ends with a written report and a specific recommendation: resolve it internally, go deeper with a full assessment, or take it to your own counsel.

WHAT YOU RECEIVE

A written Findings Summary with a plain-language Risk Heat Map across all four framework areas, plus specific recommended next steps, including which follow-on AEC service, if any, best fits your organization.

$1,500 Flat Fee · Delivered Virtually or On-site · Contact: consult@aegisethosconsulting.com

COMPREHENSIVE AI GOVERNANCE ASSESSMENT · AUD-COMP

Audit-Depth Review for When Screening Isn't Enough

WHAT IT DOES

A 30 to 60 day engagement that extends AEC's AI Risk Triage into full audit depth. Where a triage session flags where to look harder, this assessment does the looking: reviewing real outcome data across your AI systems, interviewing the people who use them, and examining vendor documentation to produce a written findings report with a specific remediation roadmap. It adds a structured root-cause layer, so a finding is traced to why it happened, not just that it happened.

WHAT IT MEASURES

Governance, evidence-verified
Accountability, policy, and oversight checked against documents and practice, not just what you tell us.

Multi-system outcome review
Real decision data reviewed across every AI system in scope, not a single snapshot.

Stakeholder interviews
Structured conversations with the people who build, use, and are affected by your AI systems.

Cross-framework obligations
Title VII, FTC Act Sec. 5, FHA, ECOA/Reg B, and the EU AI Act, checked against what is actually documented.

WHY IT'S DIFFERENT

- **Audit-depth, not just a self-report.** RISK-90 is a fast screen based on a single conversation. This assessment verifies findings against real documents, real outcome data, and real interviews.
- **Root-cause, not just symptoms.** A structured root-cause step traces each material finding to its underlying driver, so remediation fixes the cause.
- **Built for real capacity.** Thorough manual audit and arithmetic-level adverse-impact analysis; anything needing statistical modeling is flagged and routed to a specialist.
- **An actionable roadmap.** Every finding carries a recommended remediation, an owner, and a target date, plus a clear flag for anything to take to your legal counsel first.

INVESTMENT, PRICED BY SCOPE

SCOPE	INVESTMENT
1 to 2 AI systems or tools	**$10,000**
3 to 4 AI systems or tools	**$14,500**
5 or more systems or tools, or any lending, housing, or hiring data	**Scoped, from $18,500**

From $10,000 · 30–60 Day Engagement · Priced by scope · Contact: consult@aegisethosconsulting.com

DISPARATE IMPACT STANDARDS ASSESSMENT

Multi-Framework AI Fairness Screening, Built for the Underserved Middle

WHAT IT DOES

A structured evaluation of a specific AI tool (hiring software, tenant screening, credit scoring, or similar) against five federal and international standards, wherever you operate, governing discriminatory outcomes. It identifies where a tool's results may differ across groups of people, maps that finding to the specific law it falls under, and gives your organization a clear, defensible record of the review.

WHAT IT MEASURES

Title VII, Civil Rights Act
Employment decisions: hiring, promotion, termination.

FTC Act, Section 5
Unfair or deceptive practices in consumer-facing AI.

Fair Housing Act
Rental, sale, and tenant-screening decisions.

ECOA / Regulation B
Credit, lending, and loan-term decisions.

EU AI Act, Chapter III
High-risk system classification and equality obligations.

Eight protected-class dimensions
Each mapped to the law that actually protects it, not a one-size-fits-all checklist.

WHY IT'S DIFFERENT

- **Built for smaller pools, not enterprise scale.** Off-the-shelf fairness tools are calibrated for organizations screening thousands. This accounts for the smaller applicant, tenant, or borrower pools typical of municipalities, non-profits, and community-based businesses, and flags when a sample is too small to trust.
- **Practitioner-led, not a black box.** Delivered by a governance and risk-management professional with 25 years of banking-grade compliance experience, not an automated score with no one to explain it.
- **Disciplined about its own limits.** It flags risk and identifies what needs a closer look. Findings that reach a legal conclusion are yours to take to your legal counsel; results that need deeper statistical review are flagged for a specialist.

WHAT YOU RECEIVE

A findings report that identifies which of the five frameworks apply to your tool, where outcomes may differ across protected groups, and a clear set of next steps: what your team can resolve internally, and what you should take to counsel for review.

$4,000 per tool screened · Standalone or added to an AUD-COMP engagement · Contact: consult@aegisethosconsulting.com

READY-AI · AI READINESS & RESPONSIBILITY ASSESSMENT

Can Your Organization Adopt and Sustain AI Responsibly?

WHAT IT DOES

A forward-looking assessment that answers a different question from the legal-controls work: not whether your AI use is legally defensible, but whether your organization can adopt and sustain AI responsibly in the first place. It examines culture, capacity, governance design, vendor exposure, and human oversight, and produces a single Readiness & Responsibility profile. It is priced and sold separately from RISK-90 and AUD-COMP; a client may purchase either the legal-controls work, READY-AI, or both.

WHAT IT MEASURES

Organizational Readiness
Problem-fit, data, capacity, and lifecycle plan for adopting AI at all.

Leadership & Culture Fit
Leadership receptiveness and the cultural risk signature the organization carries.

Governance Framework Alignment
Whether AI use is mapped to the right multi-level frameworks.

Contract, IP & Liability Exposure
Vendor-contract terms, IP exposure, and where liability may be shifted onto you.

Responsible Use & Oversight Design
Harm profile, human-in-the-loop design, and over-reliance or skill-erosion risk.

WHY IT'S DIFFERENT

- **Forward-looking, not backward-looking.** The legal-controls assessments audit against external standards after the fact. READY-AI asks whether the organization is positioned to adopt AI well before exposure is created.
- **Especially built for the vendor-contract trap.** Underserved-middle organizations often sign AI vendor contracts without the leverage or counsel to negotiate them. READY-AI surfaces liability-shifting clauses and IP exposure and flags them for review.
- **It reports; it does not advise.** Every finding is written observationally, what was observed and where the gaps are, never as a business recommendation. Findings that touch a legal standard or a contract term are flagged for you to take to your counsel.

WHAT YOU RECEIVE

A single Readiness & Responsibility profile organized by section, led with the one synthesis message the organization most needs to understand, each finding stated observationally, with legal- and contract-touching findings flagged for your counsel. READY-AI reports what was observed; your organization decides what to do.

$9,500 Flat Fee · Standalone or paired with the legal-controls stack · Contact: consult@aegisethosconsulting.com

THE AI TOWN HALL

Free Community Education, and the Front Door to Every Service in This Packet

FREE · NON-PARTISAN · EDUCATIONAL — A FREE PUBLIC EVENT; NOBODY IS ASKED TO BUY ANYTHING TO GET VALUE FROM IT.

WHAT IT IS

A 90-minute community conversation held in familiar spaces: libraries, community centers, chambers of commerce, faith spaces, and non-profit gatherings. People leave understanding three things: where AI is already deciding things for them (jobs, housing, credit, insurance), what rights and laws still protect them, and what concrete steps they can take next.

WHO IT IS FOR

Residents
Renters, job seekers, borrowers, and parents; the people AI often decides about, usually without their knowledge.

Small business owners
Main-street employers who may already use AI tools that can carry legal responsibility.

Non-profit leaders
Directors who screen volunteers, handle sensitive data, or receive AI-scored grant decisions, often with no legal or technical staff.

Local officials & civic staff
Council members, clerks, and school board members who will vote on or administer AI tools.

WHY IT MATTERS, AND HOW IT CONNECTS TO AEC

- **It meets a real gap.** Residents and small employers are already subject to AI decisions and have almost nowhere to turn for plain-language guidance.
- **It is genuinely free and non-partisan.** Focused on rights, literacy, and practical action rather than any product or political position, with a take-home page in the community's main languages.
- **It is the front door.** Many attendees leave realizing they already use AI tools that score, rank, or filter people, and that no one is advising them on whether those tools are fair, compliant, or defensible. That is exactly the gap the assessments in this packet fill.

Free to Attendees · 90 Minutes · Hosted in your community space · Contact: consult@aegisethosconsulting.com

Practitioner-led community education. AEC identifies areas of potential risk; it does not render legal opinions or provide legal representation, and findings that need a lawyer are flagged for you to take to counsel.

Funder-Ready Packages, the Community Rate, and Ongoing Support Retainers

FUNDER-READY PACKAGES

One itemized budget that a foundation or grant program can fund, with the written proposal funders ask for. Every package is built from the list prices, with no bundle discount, and comes in two levels.

- **Funder-Ready Package — up to $20,000.** A fixed, not-to-exceed budget sized for most grants.
- **Comprehensive Funder Package — scoped, above $20,000.** For funders who can support the full ladder or a higher audit tier, confirmed in a written scope with a not-to-exceed total before work begins.

EXAMPLE PACKAGE	INCLUDES	LIST TOTAL
Core	RISK-90, READY-AI, 1 Disparate Impact screening	**$15,000**
Expanded	RISK-90, READY-AI, 2 Disparate Impact screenings	**$19,000**
Audit-led	AUD-COMP (Tier 1), READY-AI	**$19,500**
Comprehensive	RISK-90, READY-AI, AUD-COMP (Tier 1), 1 Disparate Impact screening	**$25,000**

COMMUNITY RATE

Qualifying community organizations may be eligible for a reduced rate on RISK-90, READY-AI, and Disparate Impact screening. Eligibility is based on written criteria, not negotiation. Ask when you request a quote.

ONGOING SUPPORT RETAINER

For organizations that want a steady, defined amount of help after an assessment. No embedded role: a set number of hours each month, used for the work below.

LEVEL	HOURS / MO	FEE / MO	INCLUDES
Essentials	**Up to 6**	**$2,000**	Monthly standing call, one AI tool or vendor-change review, and a regulatory watch note.
Standard	**Up to 12**	**$4,000**	Everything in Essentials, plus a quarterly RISK-90 refresh and tracking of open findings.
Extended	**Up to 20**	**$6,500**	Everything in Standard, plus leadership briefings, staff sessions, and two-business-day turnaround.

Three-month minimum, billed monthly. Unused hours do not roll over. Additional hours are billed at $350 per hour, quoted in advance. AEC identifies and flags findings that touch a legal standard; taking them to counsel is the client's responsibility. AEC does not provide legal representation.

Funder packages up to $20,000, or scoped above · Retainers from $2,000 / month · Contact: consult@aegisethosconsulting.com