Evening Star AI | eveningstar.ai

WHITEPAPER | VERSION 2.0

The Case for Governed Agentic AI in Municipal Government

Delivering Better Services, Faster Decisions, and Defensible Operations

For city managers, CIOs, department heads, procurement leaders, and elected officials navigating the transition from fragmented legacy workflows to governed AI-enabled operations.

Prepared by	Publication Date	Version	Companion Publication
Evening Star AI	May 2026	2.0 (Revised Edition)	Governing Agentic AI in Public and Regulated Environments

eveningstar.ai | raygauger.com | rcgauger@gmail.com | eveningstarai@protonmail.com

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Municipal governments are being asked to deliver faster, more responsive public services under increasingly difficult operating conditions.

The Pressures Cities Face

- persistent workforce shortages,
- aging and disconnected software systems,
- rising resident expectations,
- expanding compliance obligations,
- growing public records burdens, and
- flat or constrained operating budgets.

Most municipalities are not struggling because public employees lack commitment or expertise. They are struggling because too much staff capacity is consumed by repetitive administrative coordination work:

- routing requests,
- validating submissions,
- re-entering data across systems,
- tracking deadlines,
- classifying records,
- reviewing compliance checklists, and
- manually coordinating workflows between disconnected platforms.

Agentic AI Introduces a New Operational Model

Unlike first-generation chatbots or traditional workflow automation, governed agentic AI systems can:

- interpret unstructured inputs,
- plan multi-step workflows,
- invoke approved tools and APIs,
- adapt actions based on intermediate results,
- maintain audit logs,
- escalate exceptions to humans, and
- operate autonomously within policy-defined boundaries.

This matters because municipal operations are fundamentally workflow-intensive environments.

The Core Argument

> The opportunity is not to replace municipal judgment.
> The opportunity is to reduce the administrative friction surrounding that judgment.

Cities that deploy governed AI systems effectively are likely to realize measurable advantages in service responsiveness, backlog reduction, staff capacity utilization, operational consistency, audit readiness, resident experience, and long-term organizational scalability.

What This Paper Presents

1. A practical explanation of what agentic AI is — and is not.
2. A grounded assessment of municipal operational pressures.
3. Realistic municipal use cases.
4. Governance and accountability requirements.
5. Known risks and mitigation strategies.
6. A phased implementation framework.
7. A readiness model for municipal leaders.

> ***Governed agentic AI is best understood not as a technology trend, but as a municipal operational capacity strategy.***

SECTION 1 — The Municipal Operating Reality

The Structural Capacity Problem

Municipal governments increasingly operate in a permanent capacity deficit. Over the last decade, local governments have faced a convergence of pressures:

- experienced staff retirements,
- difficulty recruiting specialized personnel,
- increasing digital service expectations,
- expanding records and compliance obligations,
- cybersecurity pressures,
- infrastructure maintenance backlogs, and
- fragmented technology ecosystems.

Many municipalities continue to rely on operational environments composed of disconnected enterprise systems, manual workflow coordination, email-driven approvals, spreadsheet tracking, paper-intensive intake processes, and limited API interoperability.

Even highly capable departments often spend substantial time on administrative coordination work rather than judgment-intensive public service work.

Workforce Constraints

Research from organizations including the National League of Cities (NLC), ICMA, and NASCIO consistently shows that local governments face persistent workforce recruitment and retention challenges, particularly in permitting, records management, procurement, finance, inspections, IT operations, and cybersecurity.

Many municipal leaders report that staffing limitations — not merely budget limitations — are now the primary operational constraint. In practical terms, this means:

- permit applications wait in queues,
- public records requests accumulate,
- inspection schedules become reactive,
- invoice processing slows,
- procurement timelines extend, and
- residents experience slower response cycles.

Legacy System Fragmentation

A typical municipality may operate a permitting platform, an ERP, a records management system, GIS infrastructure, a CRM or 311 platform, procurement software, asset management systems, email systems, and departmental databases — many of which were never designed to interoperate cleanly.

As a result, municipal staff frequently act as the integration layer between systems. Examples include:

- copying data between applications,
- manually validating records,
- routing documents via email,
- re-entering information already stored elsewhere,
- coordinating approvals across departments, and
- maintaining shadow spreadsheets to track workflow status.

This creates operational drag that scales poorly as service demand increases.

The Resident Expectation Shift

Residents increasingly compare government service responsiveness to private-sector digital experiences. While municipal governments operate under very different legal and accountability constraints than private companies, resident expectations have nevertheless shifted toward faster acknowledgment, real-time status visibility, digital-first interactions, mobile accessibility, 24/7 intake capability, and transparent communication.

The result is an expanding expectation gap between how municipal systems currently operate and how residents increasingly expect services to function.

SECTION 2 — What Agentic AI Is — and Is Not

Defining Agentic AI

The term 'AI' is now used to describe a wide range of technologies. For municipal leaders, it is important to distinguish among traditional workflow automation, robotic process automation (RPA), generative AI assistants, AI copilots, and agentic AI systems. These are not interchangeable.

Technology Type	What It Does	Key Limitation for Municipal Use
Traditional Workflow Automation	Follows fixed rules; executes predefined sequences; operates deterministically	Brittle; struggles with unstructured input, exceptions, and dynamic decision-making
Robotic Process Automation (RPA)	Mimics user interactions with existing UI; automates repetitive screen-based tasks	Fragile when UI changes; limited reasoning; no natural language understanding
Generative AI / Copilots	Summarizes, drafts, classifies, answers questions; interprets natural language	Primarily conversational; does not independently execute operational workflows
Agentic AI	Combines language understanding, planning, tool use, workflow orchestration, and autonomous execution within policy constraints	Requires governance architecture; more complex to deploy and monitor than previous types

What Agentic AI Actually Does

A governed municipal agent may:

8. receive an intake request,
9. determine request type,
10. retrieve parcel or records data,
11. validate completeness,
12. generate notifications,
13. update workflow systems,
14. escalate exceptions,
15. log every action, and
16. route unresolved issues to staff.

The Key Distinction

The key distinction is not conversation.
The key distinction is operational execution within defined governance boundaries.

Why This Matters for Municipal Operations

Many municipal tasks are not intellectually complex. They are operationally repetitive. That makes them strong candidates for governed AI orchestration. Examples include intake validation, routing, scheduling, checklist enforcement, records tagging, invoice matching, compliance monitoring, and notification coordination.

This does not eliminate the need for human oversight. Instead, it changes how staff time is allocated — from administrative coordination toward judgment-intensive public service work.

SECTION 3 — Municipal Use Cases

Guiding Principle

The strongest early municipal AI deployments are operationally narrow, highly repetitive, low-risk, reversible, measurable, and governance-friendly.

The operational outcomes described below are illustrative and should be understood as directional rather than guaranteed. Actual results depend heavily on implementation quality, workflow maturity, data quality, API availability, governance rigor, staff adoption, and organizational readiness.

3.1 Service Request Intake and Routing

Service Request Intake and Routing

311 / Citizen Services / City Clerk

TYPICAL MUNICIPAL CHALLENGE

- 311 and resident service teams receive requests through web forms, email, phone calls, mobile applications, council referrals, and walk-in interactions.
- Many requests are incomplete, duplicated, misclassified, or routed inconsistently.
- Staff time is often consumed by administrative triage rather than resolution.

POTENTIAL AGENTIC WORKFLOW

- Classify requests by type and category.
- Detect likely duplicates against open tickets.
- Validate required information and flag incomplete submissions.
- Assign routing categories and create CRM records.
- Acknowledge receipt and escalate ambiguous cases to staff.

EXPECTED OPERATIONAL BENEFITS

- Faster acknowledgment cycles and reduced manual routing effort.
- Improved queue consistency and resident communication.
- Improved reporting visibility across service lines.
- Staff capacity redirected from triage to resolution.

NOTE

- This is often one of the strongest first deployment candidates: high-volume, errors are usually reversible, and residents immediately notice service improvements.

3.2 Permit Intake and Completeness Review

Permit Intake and Completeness Review

Planning & Development / Building Department

TYPICAL MUNICIPAL CHALLENGE

- Permit technicians spend substantial time validating submission completeness, confirming document presence, checking parcel information, identifying incorrect permit types, and issuing deficiency notices.
- This creates avoidable queue delays before substantive technical review begins.

POTENTIAL AGENTIC WORKFLOW

- Compare submissions against permit-type checklists.
- Cross-reference GIS and parcel records.
- Identify missing documents and generate standardized deficiency notices.
- Route complete submissions to the correct reviewer queue.
- Escalate exceptions for staff review.

EXPECTED OPERATIONAL BENEFITS

- Shorter intake queues and reduced administrative burden.
- Improved reviewer throughput and more consistent deficiency notices.
- Improved applicant visibility into submission status.

NOTE

- Cities should avoid fully autonomous permit approvals. The strongest near-term value lies in administrative workflow acceleration rather than final regulatory determination.

3.3 Public Records and FOIA Intake

Public Records and FOIA Intake

City Clerk / Records Management / Legal

TYPICAL MUNICIPAL CHALLENGE

- Public records requests arrive in inconsistent formats and require manual routing, custodial determination, acknowledgment tracking, and deadline monitoring.
- Municipalities increasingly face records backlogs, rising request complexity, and increasing public scrutiny.

POTENTIAL AGENTIC WORKFLOW

- Classify requests and identify likely custodians.
- Create tracking records and acknowledge receipt.
- Monitor statutory deadlines and generate alerts.
- Flag requests requiring legal review.

EXPECTED OPERATIONAL BENEFITS

- Faster initial response cycles and improved tracking consistency.
- Improved records discoverability and reduced administrative backlog pressure.
- Better compliance with statutory response timelines.

NOTE

- Municipalities must recognize that AI-generated records, prompts, logs, and reasoning traces may themselves become discoverable public records depending on jurisdictional law. Governance and retention policies are therefore essential.

3.4 Procurement and Contract Compliance

Procurement and Contract Compliance
Procurement / Finance / City Attorney

TYPICAL MUNICIPAL CHALLENGE

- Procurement teams frequently manage clause verification, insurance checks, vendor validation, debarment screening, contract milestone tracking, and compliance documentation.
- These tasks are often repetitive but operationally important.

POTENTIAL AGENTIC WORKFLOW

- Validate required clauses against policy checklists.
- Screen vendors against debarment and insurance requirements.
- Monitor contract milestones and track renewal deadlines.
- Generate exception reports for staff review.

EXPECTED OPERATIONAL BENEFITS

- Reduced administrative review time and improved consistency.
- Improved visibility into contract obligations.
- Earlier detection of compliance issues.

NOTE

- Human legal review should remain mandatory for nonstandard clauses, disputed language, and high-risk procurements.

3.5 Finance and Invoice Processing

Finance and Invoice Processing
Finance / Accounts Payable

TYPICAL MUNICIPAL CHALLENGE

- Accounts payable teams often spend substantial time matching invoices, validating purchase orders, checking vendor information, routing exceptions, and identifying duplicates.

POTENTIAL AGENTIC WORKFLOW

- Extract invoice data and compare against PO records.
- Validate vendor information against approved vendor registry.
- Identify discrepancies and route exceptions for staff review.
- Flag duplicate submissions.

EXPECTED OPERATIONAL BENEFITS

- Reduced manual matching effort and improved exception handling.
- Faster processing cycles and improved fraud detection support.

NOTE

- Straight-through processing should be restricted to low-risk transactions, defined tolerance thresholds, and policy-approved categories.

SECTION 4 — Governance Is the Enabler

Governance Is Not Optional

Municipal AI deployments operate in environments shaped by public accountability, open records law, audit requirements, procurement rules, cybersecurity obligations, and democratic oversight.

Governed deployment is not a secondary concern. It is the prerequisite for sustainable deployment.

Core Governance Principles

1. Human Accountability Must Remain Clear

AI systems may support workflows. Municipal accountability cannot be delegated. Cities should maintain clear ownership structures, escalation pathways, and documented approval authority.

2. Decision Logging Is Essential

Every agent action should generate structured logs documenting the triggering event, workflow objective, data sources used, tool calls, outputs generated, escalations, and final actions.

Without Decision Logging

Auditability degrades.
Public trust weakens.
Legal defensibility suffers.
FOIA responses cannot be supported.

3. Risk-Tiered Permissions

Not all AI actions carry equal risk. Municipalities should classify actions into risk tiers that determine the level of human approval required.

Risk Tier	Example Action	Human Approval Required
Tier 1	Drafting acknowledgment emails	No
Tier 2	Routing requests to department queues	Conditional
Tier 3	Financial transaction approvals	Usually
Tier 4	Regulatory determinations affecting citizens	Yes — always

4. Data Boundaries Must Be Enforced

Municipal agents should operate using least-privilege access, sandboxed tool permissions, segmented data access, and role-based authorization. Agents authorized for one data classification must not be able to access another.

5. Governance Must Be Operationalized

Governance frameworks should not exist only as policy documents. They should include testing procedures, red-team exercises, incident response workflows, review cadences, and measurable controls.

Companion Resource

The full governance framework — including decision log schemas, risk-tiered permission tables, pre-deployment checklists, and operating rhythm templates — is available in the companion publication:

Governing Agentic AI in Public and Regulated Environments (Evening Star AI, 2025/2026)

Available at eveningstar.ai

SECTION 5 — Known Risks and Failure Modes

AI systems create meaningful opportunities. They also introduce meaningful risks. Successful municipal deployment depends on understanding both.

Hallucinations and Incorrect Outputs

Generative systems can produce incorrect information, misclassify inputs, fabricate citations, or generate inaccurate summaries. Mitigation strategies include:

- confidence thresholds,
- deterministic validation,
- retrieval-grounded workflows,
- human review gates, and
- constrained tool permissions.

High-consequence decisions should never rely solely on unconstrained model outputs.

Prompt Injection and Security Risk

Municipal AI systems may become targets for prompt injection, malicious inputs, credential abuse, or workflow manipulation. Municipal deployments should therefore include input sanitization, API-layer controls, zero-trust access models, monitoring, logging, and adversarial testing.

Vendor Lock-In

Cities already face substantial dependence on long-lived enterprise systems. Municipal AI procurement should avoid recreating those constraints. Recommended procurement safeguards include:

- open APIs,
- export rights,
- portable logs,
- interoperability requirements,
- model abstraction layers, and
- termination transition clauses.

Workforce and Organizational Resistance

Public-sector adoption challenges are rarely purely technical. Municipal leaders should expect concerns regarding job displacement, transparency, bargaining impacts, operational oversight, and trust. Strong deployment programs typically include staff briefings, labor engagement, pilot transparency, documented guardrails, and explicit statements regarding workforce intent.

> ***In many municipalities, the strongest operational use of AI will be reducing burnout and administrative overload rather than reducing headcount.***

SECTION 6 — ROI and Operational Impact

Municipal ROI Should Be Framed Carefully

Municipal governments do not optimize for profit. They optimize for service delivery, responsiveness, accountability, compliance, and operational sustainability. The strongest municipal AI ROI metrics therefore include:

- cycle-time reduction,
- backlog reduction,
- staff capacity recovery,
- resident satisfaction,
- compliance consistency, and
- audit readiness.

Operational Cost Modeling

Illustrative cost models can help cities estimate potential value. However, municipalities should avoid treating generalized AI ROI estimates as guaranteed outcomes. Actual savings depend on staffing models, workflow maturity, exception rates, integration costs, governance overhead, and adoption success.

Illustrative examples may include:

- reduced manual intake effort,
- reduced rework,
- reduced duplicate processing,
- improved scheduling efficiency, and
- reduced backlog accumulation.

The More Important Value Proposition

For many cities, the primary value is not immediate labor elimination. The more important value may be:

- avoiding operational degradation,
- supporting understaffed departments,
- stabilizing service levels, and
- improving scalability without proportional staffing growth.

That Distinction Matters

Cities that frame agentic AI as primarily a headcount-reduction tool will face workforce resistance, adoption challenges, and political difficulty.

Cities that frame it as a service-stabilization and capacity strategy — allowing existing staff to do more of the work that requires their expertise — are more likely to achieve sustainable deployment.

SECTION 7 — Municipal Readiness Model

Most municipalities today remain between Stages 1 and 2. That is normal. The transition to governed orchestration should be gradual and measurable.

Stage 1

AI Exploration

- experimentation with available tools,
- policy development and early governance,
- pilot evaluation and scoping, and
- staff education and awareness building.

Stage 2

AI-Assisted Operations

- summarization and search assistance,
- drafting support for staff,
- AI copilots in existing workflows, and
- workflow augmentation without autonomous action.

Stage 3

Semi-Autonomous Workflow Orchestration

- automated routing and triage,
- structured task execution with policy constraints,
- exception escalation to human reviewers, and
- partial workflow autonomy in low-risk areas.

Stage 4

Governed Autonomous Operations

- policy-constrained orchestration across workflows,
- operational monitoring and drift detection,
- audit integration and decision log management, and
- mature governance controls with regular red-team review.

Readiness Assessment

Before deploying any agentic system, cities should conduct an honest assessment of their current readiness stage.

Attempting to deploy Stage 3 or 4 capabilities without Stage 1 and 2 foundations in place is a common source of failed or ungoverned municipal AI deployments.

See Appendix A for the Municipal AI Readiness Checklist.

SECTION 8 — A Practical Municipal Deployment Framework

Start Small

Municipalities consistently underperform when they overscope initial deployments, attempt enterprise-wide transformation immediately, or pursue highly sensitive use cases first. The strongest early deployments are operationally narrow, measurable, low-risk, and highly visible.

Recommended First Deployment Characteristics

Ideal first deployments typically involve high-volume workflows, reversible outcomes, low regulatory risk, existing APIs, and measurable cycle times.

For many municipalities, service-request intake remains one of the strongest first candidates.

A Realistic 90-Day Pilot Structure

Phase	Key Activities	Output
Phase 1 Discovery & Governance	Identify workflow. Establish baseline metrics. Document governance controls. Define escalation procedures and human approval gates.	Scoped use case document. Governance package v1.0 archived to ECM.
Phase 2 Integration & Testing	Configure APIs. Establish decision logging. Validate permissions. Conduct adversarial testing. Review outputs with domain staff.	Agent deployed to staging. Test report with findings addressed.
Phase 3 Controlled Production	Deploy with human oversight of all outputs. Monitor outcomes. Review incidents. Collect staff feedback. Measure against baseline.	Production deployment with human review gate active.
Phase 4 Review & Scale Decision	Compare baseline metrics. Assess governance effectiveness. Review staff experience. Determine whether expansion is justified.	Decision document: expand, adjust, or pause. Lessons documented.

The 90-day structure is not a deadline — it is a discipline. If any phase reveals governance gaps, integration problems, or staff readiness issues, the correct response is to resolve them before advancing, not to press forward on schedule.

SECTION 9 — Procurement and Vendor Evaluation

Municipal Procurement Reality

AI adoption in government must align with procurement cycles, budget approvals, public transparency requirements, cybersecurity reviews, and council oversight. Municipal leaders should expect deployment timelines to be shaped as much by governance and procurement processes as by technical implementation.

Recommended Vendor Evaluation Criteria

Cities should evaluate vendors on:

Evaluation Dimension	What to Assess
Auditability	Can all decision logs be exported in machine-readable form? Can permissions be independently audited?
API Openness	What API standards are supported? Can the city migrate workflows if the vendor relationship ends?
Interoperability	Does the system integrate with the city's existing platforms without requiring rip-and-replace?
Governance Transparency	What human-review controls exist? How are goal hierarchies and tool permissions documented?
Logging Access	How are prompts and logs retained? What is the retention period? Who has access?
Incident Reporting	What adversarial testing is performed? How are incidents reported to the city?
Data Residency	Where does city data reside? What jurisdiction governs it?
Security Maturity	What certifications are held? What is the vendor's security incident history?

Recommended Procurement Questions

Municipalities should ask vendors directly:

17. Can all decision logs be exported in machine-readable form?
18. Can permissions be independently audited?
19. What human-review controls exist?
20. How are hallucinations mitigated?
21. What adversarial testing is performed?
22. How are prompts and logs retained?
23. Can the city migrate workflows if the vendor relationship ends?
24. What API standards are supported?

Contract Language

For recommended procurement contract language covering decision log access, red-team evidence, incident notification, and data residency requirements, see the companion publication:

Governing Agentic AI in Public and Regulated Environments (Evening Star AI)

eveningstar.ai

SECTION 10 — Early Public-Sector AI Examples Worth Monitoring

While fully mature municipal agentic deployments remain early, governments globally are already operationalizing AI-enabled workflows. Examples worth studying include:

- Singapore GovTech initiatives — among the most advanced government AI deployment frameworks globally, with emphasis on explainability and auditability.
- Estonia digital government infrastructure — long-standing model for API-first, interoperable government services that provides a governance foundation for AI integration.
- New York City AI Action Plan — one of the first major US city-level AI governance frameworks with public disclosure requirements and use case inventories.
- Federal AI inventory programs — OMB-directed AI use case inventories providing visibility into how federal agencies are classifying and governing AI deployments.
- State-level AI governance initiatives — growing number of state governments developing AI policy frameworks, risk classification standards, and procurement guidance.
- Municipal digital service modernization programs — city-level digital transformation initiatives that are increasingly incorporating AI-assisted intake, routing, and service delivery.

The Lesson from Early Adopters

The lesson is not that governments have already solved autonomous governance.

The lesson is that experimentation is accelerating, governance expectations are maturing, and operational AI adoption is becoming institutional rather than experimental.

Cities that begin building governance capability now will be better positioned to adopt more capable systems as they mature.

CONCLUSION

Municipal governments are entering a period in which operational scalability will increasingly depend on intelligent workflow orchestration.

Agentic AI is not a replacement for public judgment. It is not a substitute for democratic accountability. It is not a shortcut around governance.

However, governed AI systems may allow municipalities to process work faster, reduce administrative friction, improve responsiveness, support overstretched staff, and create more resilient operating models.

The Cities Most Likely to Succeed

They will not necessarily be the ones that move fastest. They will be the ones that:

- start with realistic use cases,
- implement strong governance,
- maintain human accountability,
- engage staff early, and
- scale operationally rather than rhetorically.

The future of municipal AI adoption will likely belong not to fully autonomous government, but to governed, transparent, accountable augmentation of municipal operations. That distinction matters.

Get Support from Evening Star AI

Implementation consulting, municipal readiness reviews, and use case scoping:
raygauger.com | rcgauger@gmail.com

Governance framework, Governance Stack tooling, research, and publications:
eveningstar.ai | eveningstarai@protonmail.com

Research Contributor and Fellow opportunities available — contact eveningstarai@protonmail.com

APPENDIX A — Municipal AI Readiness Checklist

Governance

- ☐ AI governance owner identified
- ☐ Decision logging requirements defined
- ☐ Human escalation procedures documented
- ☐ Risk-tiering framework established
- ☐ Incident response process identified

Technology

- ☐ API inventory completed
- ☐ Identity and access management reviewed
- ☐ Logging infrastructure available
- ☐ Data classification policies reviewed
- ☐ Security review procedures established

Operations

- ☐ Baseline metrics documented
- ☐ Workflow maps completed
- ☐ Staff stakeholders identified
- ☐ Pilot success criteria defined
- ☐ Change-management plan prepared

Procurement

- ☐ Vendor evaluation criteria documented
- ☐ Data residency requirements defined
- ☐ Export and portability requirements established
- ☐ Cybersecurity review process defined
- ☐ Contractual governance language prepared

APPENDIX B — Suggested Reference Categories

Municipal leaders evaluating AI adoption should review materials from the following organizations and frameworks:

Organization / Framework	Relevance
NIST AI Risk Management Framework (AI RMF)	Risk classification, governance pillars, and responsible AI deployment standards
GAO AI Accountability Framework	Federal accountability expectations; applicable to state and local government by analogy
OMB AI Governance Guidance	Federal AI use case inventory requirements; minimum practice standards
ICMA	Municipal management research; AI adoption surveys; local government operational guidance
NASCIO	State and local CIO perspectives; technology governance; cybersecurity integration
National League of Cities	Workforce research; resident service trends; city-level policy analysis
Public Technology Institute	Municipal technology procurement and governance resources
Brookings Institution	AI policy research; public sector technology analysis
RAND Corporation	Government technology risk analysis; independent policy research
Deloitte Government Insights	Public sector AI adoption analysis; operational cost modeling
Stanford HAI	AI governance research; accountability frameworks; public sector applications
OECD AI Principles	International AI governance standards; public sector applications

ABOUT EVENING STAR AI

Evening Star AI focuses on AI governance, regulated-environment deployment strategy, operational AI implementation, and policy-aligned AI adoption frameworks.

Our work is grounded in practical systems delivery — evidence-focused, anti-theater, and designed for organizations that must be accountable for what their AI systems do.

Companion governance resources and implementation materials are available through Evening Star AI publications and advisory engagements.

Service Area	Description	Contact
Implementation Consulting	Municipal readiness reviews, use case scoping, governance stack deployment, and pilot program support	raygauger.com rcgauger@gmail.com
Research & Publications	Governance framework development, red-teaming resources, Governance Stack tooling, and policy-aligned AI adoption guides	eveningstar.ai eveningstarai@protonmail.com
Fellowship Program	Research Contributor and Fellow opportunities to co-develop governance tooling, contribute operational examples, and review publications	eveningstarai@protonmail.com

Evening Star AI | eveningstar.ai | eveningstarai@protonmail.com

Implementation Consulting: raygauger.com | rcgauger@gmail.com