TRUSTED
DATA AND TECHNOLOGY
INSTITUTE

AI for Public Leaders 101

Course guide

A course for the people who lead public agencies

6 modules · 18 short video segments · about 50 minutes of video
No quizzes: each module ends with a Leader's checklist

Trusted Data and Technology Institute (TDTI)
A nonprofit public benefit corporation
info@trusteddti.com

Course guide v1 · October 16, 2026 · © 2026 Trusted Data and Technology Institute. All rights reserved.

About the course

AI for Public Leaders 101 is a self-paced video course from the Trusted Data and Technology Institute (TDTI) for the people who lead public agencies. It covers the decisions a leader owns: where to use AI, how to govern it, how to buy it, how to lead staff through it and explain it to the public, and how to start with a 90-day plan. It does not teach hands-on use of AI tools.

Who it is for

Department heads and agency executives, technology and security leaders, and elected and appointed officials, in any public agency: city, county and state offices, schools, courts, health and social services, libraries and more. No technical background is needed. The course says "the people you serve" rather than assuming one kind of agency, and every example names its setting.

At a glance

Format	Narrated video with slides and captions. Six modules of three short segments each, watched at your own pace.
Length	About 50 minutes of video. Modules run 7:49 to 9:13; segments run about 2 to 4 minutes.
Assessment	None. Each module ends with a Leader's checklist of three or four actions to take in your agency, and one line to remember.
Before you start	Generative AI Fundamentals 101 is recommended, not required. Module 1 gives a short leader-level refresher.
Contact	Online form on the TDTI website, or **info@trusteddti.com**.

What you will be able to do

By the end of the course you will be able to:

- Explain three kinds of AI, what AI does well, and where it fails, and know your agency stays accountable
- Choose where to use AI by starting from a problem and screening it for value and risk to people
- Govern AI with a use policy, risk tiers, a named owner for every use, and an AI inventory
- Buy AI responsibly: ask vendors where the data goes, and put the answers in the contract
- Lead staff honestly, tell the public when AI helped, handle mistakes, and check who could be harmed
- Run a funded 90-day plan, judge whether the saving was worth the investment, and report it

Bring this course to your agency

Step	What happens
1	Contact TDTI through the online form on the TDTI website, or by email to **info@trusteddti.com**.
2	TDTI works with you to make the course available to your leadership team.
3	Leaders watch at their own pace, and use each module's Leader's checklist with their team.

A short introduction video (1:01) explains who the course is for and how to bring it to your agency.

How TDTI can help

TDTI can work with your team on any step of your 90-day plan:

- Choosing your first AI project
- Writing your AI use policy
- Running and implementing the pilot
- Assessing the return on investment

Course map

Each module has three segments and ends with a Leader's checklist. There are no quizzes.

#	Title	Segments	Video	Ends with
1	**What leaders need to know about AI**	Three kinds of AI, in plain terms, What AI does well, and where it fails, Where AI is already in your agency	9:10	Leader's checklist
2	**Choosing where to use AI**	Start from a problem, not a tool, A value-and-risk screen, Pilots: scope, measures, and the stop decision	8:03	Leader's checklist
3	**Governing AI in your agency**	What an AI use policy covers, Risk tiers, oversight and accountability, Your AI inventory, and frameworks to lean on	8:07	Leader's checklist
4	**Buying AI responsibly**	AI arrives inside software you already buy, Questions to ask vendors, Contract terms and checking after go-live	8:09	Leader's checklist
5	**Leading people and public trust**	Your workforce: roles, training, and labor partners, Telling the public, and handling a mistake, Equity and bias: who could be harmed	7:49	Leader's checklist
6	**Your 90-day plan**	Your first 90 days, Measure and report to your board or council, Course recap	9:13	Leader's checklist
	Course introduction	Who the course is for, what it covers, and how to bring it to your agency	1:01	
	Total		**50:29**	

Video lengths are measured from the current course videos and include the title cards between segments. The total does not include the introduction video.

MODULE 1

What leaders need to know about AI

9:10 of video · three segments · ends with a Leader's checklist

Leaders get a plain-language picture of three kinds of AI, what AI does well and where it fails, with a public example, and where AI is probably already at work in their agency.

What you will learn

- Three kinds of AI, in plain terms
- What AI does well, and where it fails
- Where AI is already in your agency

Segments

#	Segment	Length
1.1	**Three kinds of AI, in plain terms** Predictive, generative and agentic AI in plain terms, and why each needs a different level of control. **Key points:** Predictive AI scores, generative AI writes, agents act · Each needs a different level of control	2:58
1.2	**What AI does well, and where it fails** Where AI helps and where it fails, with New York City's MyCity chatbot as a public example. **Key points:** AI is fast, fluent, and sometimes wrong · Your agency stays accountable	3:05
1.3	**Where AI is already in your agency** AI you may already pay for, tools staff may already use, and three questions to find them. **Key points:** You can't govern AI you can't see · Ask IT, departments and vendors	3:00

Leader's checklist

- ☐ Know the three kinds: predictive, generative, agentic
- ☐ Ask where AI is already turned on
- ☐ Find out which tools staff use today
- ☐ Name one person to own the answers

Remember: AI writes what sounds right, not what is right. Your agency stays accountable.

MODULE 2

Choosing where to use AI

8:03 of video · three segments · ends with a Leader's checklist

Leaders learn to start from a problem rather than a tool, screen each idea for value and for risk to people, and run a small pilot with measures and stop conditions set in advance.

What you will learn

- Start from a problem, not a tool
- Screen each idea for value and risk
- Run a pilot with a clear stop decision

Segments

#	Segment	Length
2.1	**Start from a problem, not a tool** Problems worth solving, writing the problem down first, and tool-first versus problem-first. **Key points:** Start with a problem you can measure · Ask whether AI is even needed	2:31
2.2	**A value-and-risk screen** Two questions for every idea, low-risk internal drafting versus high-impact decisions about people. **Key points:** Screen every idea for value and risk · Treat decisions about people as high impact	2:41
2.3	**Pilots: scope, measures, and the stop decision** Keeping a pilot small, deciding success before you start, and deciding when to stop. **Key points:** Pilot small, with a fixed end date · Write measures and stop conditions first	2:44

Leader's checklist

- ☐ List three problems before you look at tools
- ☐ Screen each for value and risk to people
- ☐ Pick one low-risk pilot with an end date
- ☐ Write measures and stop conditions first

Remember: Pick the problem first, and keep AI away from decisions about people until you can govern it.

MODULE 3

Governing AI in your agency

8:07 of video · three segments · ends with a Leader's checklist

Leaders learn what an AI use policy covers, how to match controls to risk with a named owner for every use, and how to keep an AI inventory using published frameworks rather than a blank page.

What you will learn

- What an AI use policy covers
- Risk tiers, oversight and accountability
- An AI inventory, and frameworks to lean on

Segments

#	Segment	Length
3.1	**What an AI use policy covers** Why a policy now, five things it should cover, and templates you can start from. **Key points:** Write a policy now, even an interim one · Cover tools, data, review, disclosure, records	2:44
3.2	**Risk tiers, oversight and accountability** Matching controls to risk, extra safeguards for high-impact AI, and an owner for every use. **Key points:** Match controls to the risk tier · High impact: test, monitor, appeal	2:29
3.3	**Your AI inventory, and frameworks to lean on** Keeping an AI inventory, the NIST AI Risk Management Framework, and what trustworthy AI looks like. **Key points:** List every AI use, and review it yearly · Lean on NIST and GAO, not a blank page	2:47

Leader's checklist

- ☐ Adopt an AI use policy, starting from a template
- ☐ Set risk tiers, with extra safeguards for high impact
- ☐ Name an owner for every AI use
- ☐ Start an AI inventory, and review it yearly

Remember: If no one owns it, no one governs it.

MODULE 4

Buying AI responsibly

8:09 of video · three segments · ends with a Leader's checklist

Leaders learn that AI often arrives inside software they already buy, which questions to ask vendors, and which promises to put in the contract and keep checking after go-live.

What you will learn

- AI arrives inside software you already buy
- Questions to ask vendors
- Contract terms, and checking after go-live

Segments

#	Segment	Length
4.1	**AI arrives inside software you already buy** AI in updates and renewals, which features need a full review, and AI in every purchase. **Key points:** AI arrives in updates and renewals · Ask 'does this use AI?' every time	2:42
4.2	**Questions to ask vendors** Where your data goes, how the vendor knows it works, security and records, and whether you can leave. **Key points:** Ask about data, testing, security and exit · Get the answers in writing	2:40
4.3	**Contract terms and checking after go-live** Contract terms, testing with your own data, and checking after go-live when the model changes. **Key points:** Get the promises in the contract · Test with your own data · Keep checking after launch	2:41

Leader's checklist

- ☐ Ask 'does this use AI?' in every purchase
- ☐ Get no-training and data-return terms in writing
- ☐ Test with your own data before launch
- ☐ Schedule checks after go-live

Remember: Ask where the data goes before you ask what the tool can do.

MODULE 5

Leading people and public trust

7:49 of video · three segments · ends with a Leader's checklist

Leaders learn to talk with staff honestly about AI, train every role and involve labor partners, tell the public when AI helped, handle a mistake, and check who could be harmed.

What you will learn

- Your workforce: roles, training, labor partners
- Telling the public, and handling a mistake
- Equity and bias: who could be harmed

Segments

#	Segment	Length
5.1	**Your workforce: roles, training, and labor partners** Talking about jobs honestly, training by role, and bringing labor partners in early. **Key points:** Be honest about what changes · Train every role	2:41
5.2	**Telling the public, and handling a mistake** Saying when AI helped, publishing your policy and inventory, and responding when AI gets it wrong. **Key points:** Say when AI helped · Publish your policy and inventory · Plan your response to mistakes now	2:22
5.3	**Equity and bias: who could be harmed** Where bias comes from, asking who could be harmed, and checking results by group. **Key points:** Ask who could be harmed, every time · Check results by group · Give people a way to appeal	2:39

Leader's checklist

- ☐ Tell staff what AI will and won't change
- ☐ Train every role, starting with managers
- ☐ Publish your policy and AI inventory
- ☐ Check high-impact uses for unequal results

Remember: Tell people when AI helped, and tell them fast when it went wrong.

MODULE 6

Your 90-day plan

9:13 of video · three segments · ends with a Leader's checklist

Leaders build a 90-day plan: set the ground rules, fund and run one pilot, decide on the evidence, measure whether the saving was worth the investment, report to their board or council, and review the whole course.

What you will do

- Plan your first 90 days
- Measure and report the results
- Recap the course

Segments

#	Segment	Length
6.1	**Your first 90 days** Days 1 to 30: the ground rules. Days 31 to 60: prepare and fund one pilot. Days 61 to 90: run, check, decide. **Key points:** Rules first, then one funded pilot · Decide on evidence, including cost	3:28
6.2	**Measure and report to your board or council** What to measure, whether the saving is worth the investment, reporting in plain language, and keeping governance alive. **Key points:** Measure time, quality, feedback and cost · Show whether the value beat the cost · Report problems as well as wins	3:10
6.3	**Course recap** The six modules in one look, questions for your next leadership meeting, your next step, and how TDTI can help.	2:28

Leader's checklist

- ☐ Name your AI owner
- ☐ Adopt an interim AI use policy
- ☐ Start your AI inventory
- ☐ Pick one low-risk pilot

Remember: Start with one pilot, measure it, and report it.

After the course

Module 6 ends with questions to take to your next leadership meeting. If your team can't answer them yet, that is your starting point.

Questions for your next leadership meeting

- Where is AI in use today?
- Who owns each use?
- What is our policy?
- What is our first pilot?

Your 90-day plan

Days	What to do
1 to 30	Set the ground rules: name an AI owner, adopt an interim AI use policy, and start your AI inventory.
31 to 60	Prepare one pilot: choose a high-value, low-risk task, set measures and stop conditions, budget and fund the pilot, and train the pilot team.
61 to 90	Run the pilot with people reviewing, compare the results with your measures and with what the pilot actually cost, and decide: expand, change or stop.

Is the saving worth the investment?

- Value: hours saved, times what an hour of staff time costs, plus value money can't count, such as faster service and fewer errors.
- Full cost: licenses, training, setup, and the staff time spent reviewing the AI's work.
- Hours freed are not dollars saved unless the budget changes. Say what the time went to.
- Compare a full year of running costs with a full year of value, not just the pilot.

Your next step

- Share this course with your leadership team
- Enroll your staff in Generative AI Fundamentals 101
- Start your 90-day plan this week

Sources

The course content is checked against the published sources below. Each slide in the course files lists the sources it relies on, and the course files hold the exact passages quoted. Access dates are in October 2026.

Source	Used in
OpenAI, 'A practical guide to building agents' (PDF; no date shown; accessed 2026-10-16) https://cdn.openai.com/business-guides-and-resources/a-practical-guide-to-building-agents.pdf	M1
OpenAI Help Center, 'How your data is used to improve model performance' (accessed 2026-10-16) https://help.openai.com/en/articles/5722486	M1, M4
NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, July 2024 (accessed 2026-10-16) https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf	M1, M5
Colin Lecher and Katie Honan, 'Mamdani to Kill the NYC AI Chatbot We Caught Telling Businesses to Break the Law', The Markup / THE CITY, January 30, 2026, updated February 4, 2026 (accessed 2026-10-16) https://themarkup.org/artificial-intelligence/2026/01/30/mamdani-to-kill-the-nyc-ai-chatbot-we-caught-telling-businesses-to-break-the-law	M1, M2, M6
Colin Lecher, 'NYC's AI Chatbot Tells Businesses to Break the Law', The Markup (co-published with Documented and THE CITY), March 29, 2024 (accessed 2026-10-16) https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law	M1
NIST CAISI, 'Request for Information Regarding Security Considerations for Artificial Intelligence Agents', Federal Register Vol. 91, No. 5, January 8, 2026 (accessed 2026-10-16) https://www.govinfo.gov/content/pkg/FR-2026-01-08/html/2026-00206.htm	M1
Mata v. Avianca, Inc., No. 22-cv-1461 (PKC) (S.D.N.Y. June 22, 2023), Opinion and Order on Sanctions (accessed 2026-10-16) https://www.nhd.uscourts.gov/sites/default/files/pdf/Mata-v-Avianca-sanctions-order.PDF	M1
City and County of San Francisco, Generative AI Guidelines, revised November 9, 2023 (accessed 2026-10-16) https://www.sf.gov/sites/default/files/2023-11/Generative%20AI%20Guidelines%20-%20CCSF.pdf	M1, M3, M5
NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023 (accessed 2026-10-16) https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf	M2, M3, M4, M5
OMB Memorandum M-25-21, 'Accelerating Federal Use of AI through Innovation, Governance, and Public Trust', April 3, 2025 (accessed 2026-10-16; no replacement found in a web search that day) https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf	M2, M3, M5, M6
City of Boston, Interim Guidelines for Using Generative AI, Version 1.1, May 18, 2023 (accessed 2026-10-16) https://boston.gov/sites/default/files/file/2023/05/Guidelines-for-Using-Generative-AI-2023.pdf	M3, M5
Digital Government Hub, 'GovAI Coalition Templates and Resources' (published 2025; accessed 2026-10-16) https://digitalgovernmenthub.org/library/govai-coalition-templates-and-resources/	M3, M5

Source	Used in
U.S. GAO, 'Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities', GAO-21-519SP, published June 30, 2021 (accessed 2026-10-16) https://www.gao.gov/products/GAO-21-519SP	M3
Abigail Choi, 'Public Records Act Meets AI', Liebert Cassidy Whitmore, April 28, 2026 (accessed 2026-10-16) https://www.lcwlegal.com/news/public-records-act-meets-ai-what-recent-records-requests-reveal-about-chatgpt-use-in-government/	M3, M4
Microsoft Learn, 'Enterprise data protection in Microsoft Copilot and Microsoft Copilot Chat' (last updated 2026-08-18; accessed 2026-10-16) https://learn.microsoft.com/en-us/copilot/microsoft-365/enterprise-data-protection	M4
OMB Memorandum M-25-22, 'Driving Efficient Acquisition of Artificial Intelligence in Government', April 3, 2025 (accessed 2026-10-16) https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf	M4

The OMB memoranda (M-25-21 and M-25-22) apply to U.S. federal executive agencies. The course presents them as a model for state and local agencies, not as rules that bind them. The scenarios and the numbers in them are invented examples. Teaching models such as the three kinds of AI, the value-and-risk screen, the risk tiers and the 90-day plan were set by TDTI and are recorded in the course plan.