



ORGANIZATIONAL AI ADOPTION + MODERNIZATION

California Public Agency

AI Readiness & Governance *Starter Kit.*

A practical field guide for moving from unmanaged experimentation to governed, measurable AI use.

SPECIAL DISTRICTS

CITIES + COUNTIES

PUBLIC AUTHORITIES

DISCOVER / GOVERN / PRIORITIZE / PILOT

Evaluate. Deploy. Monitor. Improve.

USE THIS TOMORROW

Start with the 90-minute leadership workshop on page 16. Leave with an owner, interim guardrails, a prioritized AI use case, and a 30-day decision.

CONTENTS + HOW TO USE

Use the page that matches the decision in front of you.

READ

This is a working kit, not a compliance certificate. Complete pages with your business owner, technical owner, and applicable review functions.

ASSESS	Readiness check	03	ORIENT	Standards + applicability	04
ASSIGN	Governance roles + AI lifecycle	05	DISCOVER	AI inventory	06
CONTROL	Data handling	07	FRAME	Use-case intake	08
DECIDE	Prioritization + risk	09	REVIEW	Vendor review	10
APPROVE	Production approval gate	11	EVALUATE	Pilot evaluation	12
PROTECT	Public-sector safeguards	13	OPERATE	Incident, change + retirement	14
PLAN	90-day roadmap	15	CONVENE	90-minute workshop	16
CHARTER	One-page pilot charter	17	TRIAGE	Backlog to appropriate action	18
COMMIT	Next priorities + references	19			

CHOOSE THE SHORTEST USEFUL PATH

Need immediate guardrails? Use pages 3, 6, 7, and 16. Evaluating a pilot? Use pages 8-12 and 17. Preparing for production? Complete pages 10-14 before approval.

Applicable law, regulation, contract, grant condition, insurance requirement, and adopted agency policy take precedence. Confirm applicability with qualified legal, privacy, security, records, procurement, accessibility, and labor professionals.

5-10 MINUTE EXECUTIVE CHECK

Where are we today?

COMPLETE

Check each statement only if the practice is currently documented or consistently performed. Total the checks, then identify the three largest gaps.

- ☐ 01 Strategy - We know which organizational problems AI should address.
- ☐ 02 Ownership - A named leader is accountable for AI.
- ☐ 03 Inventory - We know which AI tools and embedded features are in use.
- ☐ 04 Policy - Employees know what is and is not permitted.
- ☐ 05 Data - Rules define what information may enter each AI environment.
- ☐ 06 Security - AI tools receive security and vendor review.
- ☐ 07 Procurement - Contracts disclose AI, data handling, and material changes.
- ☐ 08 Use cases - A repeatable intake and prioritization method exists.
- ☐ 09 Risk - Uses are classified by impact and risk.
- ☐ 10 Evaluation - Evidence shows whether a system works before production.
- ☐ 11 Human oversight - Responsibility and review are explicit.
- ☐ 12 Monitoring - Performance, incidents, cost, and vendor/model changes are tracked.
- ☐ 13 Training - Training is tailored by role and risk.
- ☐ 14 Incidents - Staff can report, contain, investigate, and recover.
- ☐ 15 Outcomes - Time, quality, cost, service, or risk outcomes are measured.

Directional score: 0-4 unmanaged | 5-8 experimenting | 9-12 governed | 13-15 operationalizing

Critical gates override the score. Stop or restrict a use if sensitive-data handling is unknown, consequential output lacks accountable human review, or the organization cannot test, log, disable, investigate, and recover.

WHAT APPLIES?

Start with obligations. Use frameworks to operate.

DISCUSS For the proposed AI use, mark each layer required, conditionally applicable, recommended, or not applicable. Record who confirmed the determination.

Layer	What it means	How to use it
Binding obligations	Laws, regulations, contracts, grants, court orders, insurance, and adopted agency policy.	Identify before design; these requirements take precedence.
NIST AI RMF 1.0	Voluntary cross-sector AI risk-management framework: GOVERN, MAP, MEASURE, MANAGE.	Use as the governance backbone.
NIST GenAI Profile	Companion guidance for GenAI-specific risks and actions.	Use for generative AI systems and services.
California CDT / SAM / SIMM	Mandatory for covered state entities in scope; not automatically binding on every local agency.	Use applicable rules when covered and selected practices as a local-agency reference model.
NIST CSF 2.0	Voluntary cybersecurity-risk framework: GOVERN, IDENTIFY, PROTECT, DETECT, RESPOND, RECOVER.	Use to organize supporting security controls.
Sector overlays	Education, health, employment, law enforcement, federal work, critical infrastructure, and accessibility.	Apply when the data, population, decision, funding, or service brings the overlay into scope.

CALIFORNIA STATE VS. LOCAL

CDT requires covered state entities to maintain a GenAI inventory and complete applicable risk/procurement steps. A city, county, air district, or other special district is not automatically subject to every CDT requirement. Label each practice: required, conditionally applicable, or recommended.

VERSION NOTE

NIST is revising AI RMF 1.0. California policies and forms continue to evolve. Record the source version and review date rather than making a timeless claim of 'compliance.'


DECISION RIGHTS

Governance is ownership made visible.

ASSIGN Write a name beside every role for each AI use case. One person may hold multiple roles, but each accountability must be explicit.

Role	Accountability
Executive sponsor	Sets outcomes, risk tolerance, resources, and resolves cross-department barriers.
AI accountable lead	Maintains inventory, intake, decisions, roadmap, reporting, and reassessment.
Business owner	Owns the workflow, baseline, outcome, adoption, and operational consequences.
Technical owner	Owns architecture, integration, access, logs, reliability, cost, and rollback.
Review functions	Security, privacy, legal, records, procurement, accessibility, HR/labor, and communications as applicable.
Operational reviewer / accountable decision-maker	Checks defined AI outputs or proposed actions and remains accountable for the decision or work product.

AI GOVERNANCE LIFECYCLE



USE THE LOOP FOR EVERY AI USE

Retirement means revoking access, preserving required records, deleting data under policy and contract, removing integrations, notifying users, and verifying the replacement or fallback. See page 14.

CONTROL WHAT YOU CAN SEE

You cannot govern AI you do not know exists.

COMPLETE

List purchased tools, consumer tools, APIs, copilots, embedded features, custom systems, contractor use, pilots, and retired systems. Mark suspected shadow use as unverified.

Tool / system	Dept. + owner	Purpose + users	Data accessed	Vendor / model	Approval + risk	Human review	Retention / record	Status

FIRST PASS

Ask procurement for contracts, IT for SSO/SaaS/browser-extension data, departments for tools, developers for APIs/models, and staff for anonymous self-reporting. The objective is visibility, not punishment.

FOUNDATIONAL CONTROL

Classify the information before selecting the tool.

DECIDE

For each AI use, identify the highest data class involved. Permit the use only in an environment approved for that class and purpose.

Class	Examples	Consumer / public AI	Approved enterprise / API
PUBLIC	Published reports, public web pages, released datasets	Only if agency policy permits; still assess accuracy, records, copyright, and accessibility.	Generally permitted within the approved use.
INTERNAL	Procedures, drafts, internal operational information	Avoid unless explicitly approved.	Permitted according to classification, contract, and need.
CONFIDENTIAL	Personnel, procurement-sensitive, privileged, security, or nonpublic business information	Prohibited by default.	Only in an approved architecture and workflow.
RESTRICTED	Student records, PHI, law-enforcement data, credentials, legally protected information	Prohibited.	Only with specialized authority, controls, agreements, and review.

SECRETS

Never paste passwords, tokens, private keys, certificates, or connection strings into prompts. Approved integrations should retrieve secrets from a managed vault and expose only the minimum necessary capability.

DOCUMENT THE DATA FLOW

Source -> collection -> prompt/context -> model/vendor -> subprocessors -> storage/logs -> output -> downstream action -> retention/deletion

- ☐ Data owner

☐ Allowed purpose

☐ Residency

☐ Deletion

☐ Breach notice
- ☐ System of record

☐ Minimum fields

☐ Retention

☐ Export

☐ Record status

START WITH FRICTION, NOT TOOLS

What problem is important enough to solve?

COMPLETE

Interview the workflow owner. Answer each prompt with current-state evidence before naming an AI product or solution.

CURRENT WORKFLOW

What happens today? Where does work enter, move, wait, fail, and finish?

PROBLEM + OUTCOME

What pain are we solving? What measurable result would matter?

VOLUME + LABOR

How often? How many items? How much staff time?

ERRORS + REWORK

Where does quality break down? What is the consequence?

DATA

What information is used, created, retained, or disclosed?

DECISION IMPACT

What happens to a person, right, service, permit, payment, job, or public statement if wrong?

AI ROLE

Assist / retrieve / summarize / generate / recommend / decide / act

HUMAN OVERSIGHT

Who reviews, approves, corrects, and remains accountable?

CONSTRAINTS

Systems, procurement, records, accessibility, labor, budget, time, vendor, security.

BASELINE

Current time, cost, quality, backlog, satisfaction, risk, or service level.

VALUE X FEASIBILITY X RISK

Choose the right first pilot.

DECIDE

Score each candidate 1-5. Favor value, readiness, and feasibility; treat risk and change burden as constraints. Select one candidate for discovery or a controlled pilot.

Candidate	Value	Readiness	Feasibility	Risk	Change burden	Baseline?	Decision

T1 / LOW

Drafting, formatting, brainstorming, or summarizing approved non-sensitive information.

T2 / CONTROLLED

Internal retrieval, document processing, coding assistance, or workflow automation with internal information.

T3 / HIGH IMPACT

Regulatory analysis, permitting or enforcement support, employment, sensitive data, or authoritative public output.

T4 / RESTRICTED

Autonomous consequential decisions, uncontrolled confidential data, or actions without appropriate oversight.

BEFORE PURCHASE OR RENEWAL

Ask what happens to data, behavior, and accountability.

REVIEW

Obtain written evidence from the vendor and contract. Record unresolved answers, required mitigations, owner, and due date before approval.

- ☐ 01

Which models and hosting services are used?
- ☐ 02
- Where is organizational data processed and stored?
- ☐ 03
- Is agency data used to train or improve any model?
- ☐ 04
- What is retained, for how long, and can retention be disabled?
- ☐ 05
- Which subprocessors receive data?
- ☐ 06
- What SSO, MFA, RBAC, and administrator controls exist?
- ☐ 07
- Are prompt, output, action, and administrator logs available?
- ☐ 08
- Can the tool act in other systems? What permissions does it receive?
- ☐ 09
- How are model, feature, price, and subprocessor changes disclosed?
- ☐ 10
- What evaluation, safety, accessibility, and security evidence is available?
- ☐ 11
- How are incidents reported and within what timeframe?
- ☐ 12
- Can agency data and logs be exported, corrected, preserved, and deleted?
- ☐ 13
- What survives contract termination?
- ☐ 14
- What uptime, support, recovery, and exit commitments exist?
- ☐ 15
- Who owns inputs, outputs, configurations, and derived artifacts?
- ☐ 16
- Does the vendor disclose newly embedded AI functionality?

CONTRACT FOR CHANGE

AI capabilities can appear after initial purchase. Require notice and review for material model, feature, data-use, subprocessor, autonomy, and retention changes.

NO-GO UNTIL ANSWERED

A demo is not a production system.

APPROVE

Complete with the business owner, technical owner, and applicable review functions. Check an item only when evidence exists; unresolved critical items block approval.

PURPOSE + OWNERSHIP

- ☐ Business outcome and baseline defined
- ☐ Business and technical owners named
- ☐ AI role and prohibited uses documented
- ☐ Human responsibility explicit

RISK + DATA

- ☐ Applicability and risk assessment complete
- ☐ Data flow and classification approved
- ☐ Privacy, security, records, accessibility reviewed
- ☐ Vendor and contract accepted

EVIDENCE

- ☐ Representative test set created
- ☐ Expected failures and misuse tested
- ☐ Accuracy/quality thresholds approved
- ☐ Bias and affected-group impacts assessed where relevant

OPERATIONS

- ☐ Least privilege and authentication enabled
- ☐ Logs, alerts, cost limits, and monitoring enabled
- ☐ Fallback, rollback, incident, and recovery tested
- ☐ Change triggers and re-evaluation date set

RECORD THE PRODUCTION DECISION

- ☐ NO-GO
- ☐ CONDITIONAL APPROVAL
- ☐ APPROVED FOR PRODUCTION
- ☐ RETURN FOR MORE TESTING

Approver: _____ Conditions / residual risk: _____ Next review: _____

PROVE USEFULNESS

Measure the workflow, not AI activity.

COMPLETE AFTER THE PILOT

For each outcome, enter the current baseline, pre-set target, actual result, supporting evidence, and Pass/Fail. Do not substitute prompt counts or logins for operational outcomes.

Workflow measure	Current baseline	Target	Actual pilot result	Evidence	Pass / Fail
Quality / correctness					
Time / throughput					
Cost / capacity					
Reliability / availability					
Security / privacy					
Fairness / accessibility					
User / public experience					

TEST DESIGN

- Use representative, difficult, known-failure, and prohibited cases.
- Compare against the current workflow and qualified review where appropriate.
- Set thresholds before final results are known; document errors and downstream impact.
- Retest after material model, prompt, data, workflow, vendor, permission, or integration changes.

SEND THE EVIDENCE TO THE PRODUCTION GATE

Use this completed table at the page 11 approval meeting. The decision is recorded there once, with conditions, residual risk, and next review date.

BEYOND TECHNICAL PERFORMANCE

Protect records, access, rights, and public trust.

DISCUSS

For the proposed AI use, identify which safeguards apply, the responsible reviewer, required evidence, and any approval condition.

PUBLIC RECORDS

Could prompts, outputs, evaluation results, decision-support records, or agent logs relate to public business? Define the authoritative record, search/export capability, retention, holds, exemptions, and deletion.

ACCESSIBILITY

Can people with disabilities access the public service and AI-generated content? Test the interface, documents, alternatives, and assistive-technology experience.

CIVIL RIGHTS

Could the system create or amplify unequal outcomes? Identify affected groups, proxies, appeal/correction routes, and accountable human review.

NOTICE + TRANSPARENCY

When should staff or the public know AI is used? State the system's role, limits, source, review status, and how to challenge an error.

WORKFORCE + LABOR

Does the use change duties, monitoring, performance evaluation, staffing, or required skills? Involve HR, labor, and affected staff early.

PUBLIC COMMUNICATION

Who approves authoritative public output? Require source verification, accessibility, records handling, correction, and incident escalation.

COUNSEL + RECORDS OFFICER

Prompts and outputs may become agency records depending on content, purpose, use, and retention. Do not retain everything forever or allow everything to disappear by default. Set a deliberate policy.

PREPARE BEFORE FAILURE

Stop, investigate, recover, and retire.

RETAIN

Assign owners and contact paths before production. Use this sequence for an AI error, disclosure, harmful output, misuse, unauthorized action, or material model/vendor change.

1**REPORT**

One clear channel for staff, vendors, and the public to report suspected AI errors, harmful output, disclosure, misuse, or unauthorized action.

2**CONTAIN**

Disable the feature, revoke credentials, isolate integrations, preserve relevant evidence, and switch to the fallback process.

3**ASSESS**

Determine data, people, systems, decisions, records, vendors, and legal or contractual notifications affected.

4**CORRECT**

Fix or roll back the system, correct downstream records/decisions, communicate appropriately, and validate recovery.

5**LEARN**

Document root cause, update tests and controls, retrain users, revise the risk decision, and set follow-up verification.

CHANGE TRIGGERS

Reassess when the model, vendor, training/data use, subprocessor, prompt, retrieval source, permissions, users, workflow, autonomy, integration, law, policy, or contract materially changes.

RETIRE SAFELY

Revoke access; export required records; preserve decisions and documentation; delete data under policy and contract; remove integrations; notify users; verify the replacement or fallback; close vendor access and cost.

FROM VISIBILITY TO EVIDENCE

A practical first 90 days.

PLAN

Choose an accountable owner and completion date for each item. Adjust sequencing to legal, procurement, labor, technology, and operational constraints.

DAYS 1-30

KNOW WHAT YOU HAVE

- ☐ Name sponsor and accountable lead
- ☐ Inventory approved, embedded, experimental, and suspected use
- ☐ Interview departments and map sensitive data
- ☐ Review current policy, contracts, records, and accessibility
- ☐ Identify 10-20 workflow pain points
- ☐ Issue interim acceptable-use guidance

DAYS 31-60

PUT GUARDRAILS AROUND IT

- ☐ Classify use cases and establish review paths
- ☐ Create vendor and procurement review
- ☐ Select 1-3 controlled pilots
- ☐ Define baselines, evaluation, logs, and human oversight
- ☐ Train pilot users and managers
- ☐ Confirm incident and fallback procedures

DAYS 61-90

PROVE SOMETHING

- ☐ Run pilots within approved boundaries
- ☐ Measure quality, time, cost, risk, and user experience
- ☐ Document failures and unexpected impacts
- ☐ Decide: stop, revise, productionize, or expand
- ☐ Publish lessons and update guardrails
- ☐ Create the 12-month adoption roadmap

NEED THE ROADMAP BUILT WITH YOUR TEAM?

MDang Solutions combines readiness, governance, workflow discovery, technical planning, and pilot design into one facilitated engagement.

FACILITATOR GUIDE

90-minute AI readiness working session.

CONVENE Use this agenda to leave with named ownership, interim guardrails, an initial inventory, a ranked AI opportunity, and a scheduled next decision.

Time	Activity	Prompt	Output
0-10	Frame the session	Which operational outcome matters most? What is out of scope today?	Shared objective
10-25	Assign ownership	Who sponsors, owns, reviews, and decides?	Named sponsor + lead
25-40	Expose use	What is purchased, embedded, tested, free, custom, or suspected?	Initial inventory
40-55	Set interim rules	Which tools and data are allowed? Where is human review mandatory?	Immediate guardrails
55-75	Rank friction	Which workflows consume time, create backlog, errors, risk, or poor service?	Ranked candidates
75-90	Commit	Which candidate earns discovery or a controlled pilot? What happens next?	Owner + date + evidence

INVITE

Executive sponsor; AI accountable lead; IT; security/privacy; records; procurement; legal; HR/labor; accessibility; one or two operational owners; and a facilitator. Keep the group small enough to decide.

MATERIALS

- ☐ This kit + readiness check
- ☐ Known tools, contracts, and policies
- ☐ Backlog + one representative workflow

INTERIM GUARDRAILS TO APPROVE

- ☐ Approved tools/accounts only
- ☐ No restricted data without approval
- ☐ Human verifies decisions and public output
- ☐ No autonomous consequential action
- ☐ Report errors, misuse, and disclosure

FACILITATION SUPPORT

If internal history, competing priorities, or unclear ownership make the conversation difficult, an independent facilitator can preserve neutrality and convert discussion into decisions. MDang Solutions offers this as a focused starting engagement.

ONE-PAGE PILOT CHARTER

Define the experiment before building it.

COMPLETE

Approve this charter before pilot access or build work begins. Keep the pilot bounded, reversible, measurable, and time-limited.

PILOT NAME + DECISION DATE	BUSINESS OWNER / TECHNICAL OWNER
PROBLEM AND CURRENT BASELINE	TARGET OUTCOME AND PASS/FAIL THRESHOLD
USERS, DATA, SYSTEMS, VENDOR/MODEL	AI ROLE + REQUIRED HUMAN REVIEW
TOP RISKS, PROHIBITED USES, AND MITIGATIONS	TEST SET, EVIDENCE, AND MONITORING
FALLBACK, ROLLBACK, INCIDENT ROUTE	BUDGET, TIME, SUPPORT, AND EXIT

NEXT DECISION

Take completed pilot evidence to the production gate on page 11. Do not expand access or integrate with production systems by default.

BACKLOG TRIAGE

Move each problem to the appropriate next action.

DISCUSS

For each backlog item, decide whether AI is appropriate now, what foundation must be fixed first, and whether a bounded experiment is justified.

Operational problem	Is AI appropriate now?	What must be fixed first?	Possible controlled experiment
Missing procedures and succession knowledge	Possibly – as drafting and retrieval support, not as the source of truth.	Owner, approved sources, access, records rules, validation workflow.	Draft one procedure from approved evidence; SME verifies every statement.
Legacy application and migration backlog	Possibly – for explanation, tests, documentation, and bounded refactoring.	Repository controls, architecture target, test coverage, dependencies, priorities.	Generate tests and documentation for one low-risk module; compare to expert review.
Old Access purchasing database	Not first. This is primarily a process, controls, data, and system-of-record decision.	Map workflow, approvals, fiscal controls, records, data quality, replacement criteria.	Use AI only to summarize interviews or draft requirements from approved material.
Excessive project requests and parallel work	Possibly – for intake support and portfolio visibility; not autonomous prioritization.	Decision rights, service catalog, intake criteria, capacity data, portfolio governance.	Classify and deduplicate a historical sample; managers verify routing and errors.

THE PROCESS-FIRST RULE

AI is not the answer to every backlog. If ownership, workflow, data quality, test discipline, capacity, or the system of record is unclear, address that foundation before automating it.

TURN THE ASSESSMENT INTO ACTION

Your next three priorities.

COMMIT

Write the three actions your agency will complete next, name one accountable owner for each, and set a decision or completion date.

1

2

3

WANT A SECOND SET OF EYES?

Complete the readiness check and bring your top three workflow candidates. In a 30-minute introductory conversation, MDang Solutions will compare notes and identify the most useful next decision - whether or not a larger engagement is warranted.

Move AI from discussion to a practical plan.

mdangsolutions.com

boss@mdangsolutions.com | Fresno / Clovis, CA | Remote support beyond

SELECTED OFFICIAL REFERENCES

NIST AI RMF 1.0: nist.gov/itl/ai-risk-management-framework
NIST GenAI Profile (NIST AI 600-1): nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
NIST Cybersecurity Framework 2.0: nist.gov/cyberframework
California CDT Technology Letters: cdt.ca.gov/policy/technology-letters/
California SIMM 5300 / 5305-F and SIMM 150: cdt.ca.gov/policy/simm/
California Public Records Act: oag.ca.gov/consumers/general/pra
ADA Title II web and mobile accessibility: ada.gov/resources/2024-03-08-web-rule/
California AI in Education: cde.ca.gov/ci/pl/aiincaifornia.asp

Information only. Not legal, compliance, procurement, employment, privacy, accessibility, or cybersecurity advice.