INDUSTRY + OPERATIONS

MANUFACTURING AI PLAYBOOK

Select one operational problem. Run a bounded pilot. Retain the evidence. Decide from measured value and risk.

For plant leaders, operations owners, production supervisors, quality, safety, IT/OT, cybersecurity, workforce and executive sponsors.

SELECT

One recurring, lower-risk workflow with a named owner.

TEST

Approved data, qualified review, baseline and stop condition.

DECIDE

Continue, revise, pause or stop from a retained record.

MISSISSIPPI ARTIFICIAL INTELLIGENCE NETWORK (MAIN)

MISSISSIPPI'S STATEWIDE AI INITIATIVE | MAINMS.ORG

ORIENTATION

CONTROL THE WORK BEFORE THE TOOL

A bounded pilot begins with an owner, an allowed data boundary, a reviewer, a baseline and a stop condition.

ORIENTATION CONTRACT

WHO	Plant and operations leaders; production, quality, maintenance, supply chain, safety, training, IT/OT and cybersecurity owners.
WHEN	A recurring delay, error, rework burden or information task has a measurable current process.
OUTPUT	A signed pilot charter and retained decision record.

LEAVE WITH

One decision-ready record naming the workflow, owner, baseline, approved system and data boundary, reviewer, evidence, stop condition and next review.

BRING TO THE FIRST SESSION

- The current process and accountable owner
- Approved source records and baseline definitions
- Qualified reviewers for operational and control boundaries
- A specific stop condition and decision date

DOES NOT AUTHORIZE

Unsafe deployment, autonomous process control, final safety or quality decisions, or unapproved data use.

FIRST WORKING SESSION - 60 MINUTES

01 NAME

State the recurring problem and accountable owner.

02 BRING

Current process, approved records and baseline definitions.

03 BOUND

Tool, data, access, reviewer and stop condition.

04 RECORD

Pilot decision, rationale, residual risk and next review.

SESSION OUTPUT **A named workflow, current-process record, approved boundary, qualified reviewer, stop condition and scheduled decision gate.**

PUBLIC AUTHORITY NIST MEP AI implementation guidance • OSHA Safety and Health Programs

Built on Mississippi’s Statewide AI Framework

The Mississippi Statewide AI Framework (August 2026) is the foundation for this playbook. It supplies shared priorities and a learning progression; this playbook applies them to manufacturing operations and workforce learning. Both are voluntary. Applicable law, adopted standards and institutional authority remain controlling.

Read the August 2026 Statewide AI Framework (PDF)

Learning progression. Workforce: early-career practice through mid-to-senior leadership. Workers may enter here without completing earlier education levels.

Five priorities in practice

Train participants and provide usable access (1); examine workforce effects and verify claims (2); approve data and controls (3); practice transferable operational skills (4); compare baseline, total effort and risk before expansion (5).

Priority numbers refer to literacy and access (1), ethics and critical thinking (2), privacy and governance (3), workforce readiness (4), and strategic alignment and measurable outcomes (5).

Selected Framework domains, practice and evidence

Framework domain	Practice and evidence to retain
6. AI-Enabled Workflow and System Design	Map the pilot from approved inputs to supervisor acceptance; retain the charter.
3. Critical Evaluation of AI Outputs	Check quantities, omissions and procedure references; retain corrections and total review time.
8. Human Agency, Judgment, and Restraint	Keep safety, quality and release decisions with qualified people; retain the decision rationale.
11. Cyber Security and Safety	Bound accounts, data and connections; retain tested stop and incident routes.

Keep the shared human-review boundary

Apply the Framework's human-review triggers (p. 35) before action: assessment or credentials; financial allocation or compensation; health or physical safety; discipline or employment; legal or compliance content; and sensitive or protected information. Its additional agent triggers cover autonomous system-of-record changes; persistent memory; tool access, credentials or external transactions; handoffs in safety-critical, regulated or student-facing processes; missing stop conditions or logged escalation to a named human owner; and institutional responses to synthetic media about identifiable people. Product or pilot approval does not replace action-specific review. Review the system and access before enabling them; review each proposed system-of-record change before execution. Advance permission cannot waive these triggers.

Use evidence within its limits

Use the crosswalk to select practice and retain observable evidence. It is MAIN's implementation mapping, not a validated proficiency test, credential or certification. Local rubrics and pilot measures alone do not establish statewide proficiency or causal impact. Recheck alignment when the Framework changes; record the edition, local owner and next review.

Source: August 2026 Framework, pp. 5-7 (priorities), 14-30 (domains and progression), and 35-37 (review, measurement and maintenance).

RISK MAP

MATCH CONTROLS TO CONSEQUENCES

Risk depends on the tool, data, user, workflow, integration, audience and consequence of error - not on the word AI alone.

		EXAMPLES	CONTROL
01	BEGIN HERE **Draft support using approved, nonconfidential information**	Agendas, shift summaries, training outlines and plain-language explanations	Human review before use
02	CONTROLLED USE **Work that can influence an operating process**	Draft work instructions, maintenance summaries, inventory analysis and internal briefings	Qualified review plus retained evidence
03	DO NOT DELEGATE **Consequential decisions or direct control**	Safety approval, equipment control, product disposition, emergency response or sole-basis employment decisions	Authorized person decides; automation stops

THE BOUNDARY-CHANGE TEST

Before a workflow moves to another band, record exactly what changed and who accepts the new consequence.

DATA	ACCESS	CONSEQUENCE	REVIEW
Did sensitivity, retention or sharing change?	Did the tool gain connectors, autonomy or OT reach?	Can an error affect people, product, process or rights?	Can a qualified person verify and reverse the output?

RECLASSIFY **If any answer changes, stop the current use and obtain a new boundary decision before the next run.**

PUBLIC AUTHORITY NIST AI Risk Management Framework • NIST Generative AI Profile

SELECTION GATE

CHOOSE ONE WORKFLOW THAT CAN TEACH YOU SOMETHING

A pilot is a controlled learning exercise. Favor a stable, recurring task with evidence that a qualified reviewer can check.

FOUR QUESTIONS BEFORE APPROVAL

01

PROBLEM

What recurring delay, error, rework burden or learning need are we solving?

02

OWNER

Who owns the current process and accepts or rejects every result?

03

EVIDENCE

What current record proves the baseline and lets us verify each output?

04

BOUNDARY

What data, access, action or consequence forces escalation or stop?

USE-CASE SCREEN

	A STRONG FIRST PILOT	ESCALATE OR REJECT
Process	Stable, recurring and owned	Unmapped, changing or unowned
Information	Approved, minimal and traceable	Sensitive, controlled or unclear
Output	Draft or analysis with qualified review	Autonomous action or final disposition
Measurement	Baseline and review burden are observable	Success is only a plausible answer

GOOD FIRST CANDIDATES

SHIFT HANDOFF

Format approved notes; supervisor verifies against source records.

TRAINING OUTLINE

Draft from a current approved procedure; qualified trainer verifies every step.

WORK-ORDER TRENDS

Organize de-identified records; maintenance experts diagnose and decide.

PUBLIC AUTHORITY NIST MEP AI implementation guidance • NIST AI RMF Playbook

OPERATING SEQUENCE

MAKE THE PILOT A MEASURABLE LINE

Four weeks. One bounded workflow. Evidence retained at every gate. Scope does not expand automatically.

01 DEFINE + APPROVE

Name owner; select workflow; approve tool and allowed data; record baseline.

OUTPUT: signed pilot charter

02 TRAIN + TEST

Train a small group; use approved samples; require review; log errors and corrections.

OUTPUT: test log + corrections

03 REPEAT + CHALLENGE

Repeat the same task; compare quality and total time; test edge cases; check new risk.

OUTPUT: measured evidence

04 DECIDE

Review with stakeholders; continue, revise, pause or stop; record owner and next review.

OUTPUT: decision record

EVIDENCE RETAINED AT EVERY GATE

QUALITY

Accuracy, corrections and reviewer acceptance

TOTAL TIME

Draft, checking, correction and handoff time

REPEATABILITY

Variation across users, records and edge cases

RISK

Data, access, safety, quality, workforce and legal effects

CONTROLLED DECISION

CONTINUE / REVISE / PAUSE / STOP

PUBLIC AUTHORITY NIST AI RMF Playbook • NIST MEP AI implementation guidance

GOVERNANCE

PUT A NAMED PERSON AT EVERY CONTROL POINT

A vendor, model or prompt does not own the decision. The pilot owner coordinates qualified review and retains the record.

ACCOUNTABLE ROLE	OPERATING RESPONSIBILITY	DECISION RIGHT
PILOT OWNER	Defines the problem, scope, baseline, evidence and decision meeting.	**Can pause the pilot at any time.**
PROCESS REVIEWER	Checks every output against source records, procedure and operational context.	**Accepts, corrects or rejects each output.**
CONTROL REVIEWERS	IT/OT, cybersecurity, privacy, safety, quality, HR, legal or procurement review the boundaries that apply.	**Approve controls within their authority.**
EXECUTIVE SPONSOR	Ensures resources, resolves conflicts and owns expansion beyond the pilot.	**Cannot waive legal, safety or security requirements.**

STOP AUTHORITY Any named owner or qualified reviewer may pause work for sensitive-data exposure, unsafe output, unauthorized access or action, missing evidence, or review burden above the approved limit.

PUBLIC AUTHORITY NIST AI Risk Management Framework • OSHA Safety and Health Programs

ILLUSTRATIVE EXAMPLE - REPLACE WITH LOCAL EVIDENCE

SHIFT-HANDOFF SUMMARY

A production team tests draft assistance for a consistent handoff. The tool does not change settings, issue instructions, make safety decisions or replace supervisor review.

01 **SOURCE**
Approved shift notes and current handoff form

02 **DRAFT**
AI formats a concise summary inside the approved environment

03 **VERIFY**
Outgoing supervisor checks every event, quantity, hold and owner

04 **RETAIN**
Source references, corrections, reviewer, final summary and decision

CONTEXTUAL PHOTOGRAPH

EVIDENCE LEDGER

PROBLEM + OWNER	Recurring handoff burden; operations owner and supervisor named.
BASELINE + TARGET	Current preparation time, missing fields, corrections and reviewer effort measured before a local target is set.
DATA + SYSTEM	Approved, minimized shift information only; credentials, unrestricted personnel data and uncontrolled OT access excluded.
HUMAN REVIEW	Supervisor checks events, quantities, safety items, maintenance references, owners and unresolved conditions.
STOP	Pause for protected-data exposure, unsafe or fabricated instruction, missing evidence, unauthorized action or excessive review burden.
DECISION	Continue only when measured value is repeatable and residual risk remains inside the approved boundary.

COMPLETE WHEN **Every output is traceable; total employee time is counted; failures and corrections are logged; residual risk is explained; and the decision has an effective period and next review.**

PUBLIC AUTHORITY NIST MEP AI implementation guidance • OSHA Safety and Health Programs

STARTER PATTERNS

FIVE MANUFACTURING AI WORKFLOWS

Replace bracketed placeholders, use only approved information, require the system to identify assumptions and missing data, and have a qualified person verify the result.

01 LOWER RISK — Draft a shift-handoff summary

USE WHEN

A supervisor needs a consistent format for production status, open issues and assigned follow-up.

ASK

Create a concise template for [LINE]. Include status, downtime, holds, safety events, shortages, open work orders, owners and next-shift priorities. Mark missing information as not provided.

VERIFY

Quantities, equipment status, holds, incidents, owners and source references.

BOUNDARY Draft only. No production action, safety decision or final instruction.

02 CONTROLLED USE — Create a training outline

USE WHEN

A qualified trainer needs a draft lesson structure based only on a current approved procedure.

ASK

Using the approved procedure, draft a [LENGTH]-minute outline for [AUDIENCE] with objectives, prerequisites, demonstration points, practice, knowledge checks and trainer sign-off.

VERIFY

Every step against the controlled document, PPE, site rules and current revision.

BOUNDARY The approved procedure remains controlling; the trainer verifies and signs off.

ALWAYS RETAIN Prompt or instruction; approved inputs; tool/version; output; corrections; reviewer; accepted record; exception; and decision.

PUBLIC AUTHORITY NIST MEP AI implementation guidance • NIST AI RMF Playbook

STARTER PATTERNS

WORKFLOWS FOR MAINTENANCE, SUPPLY CHAIN AND QUALITY

AI may organize approved information and draft a review aid. Qualified people retain diagnostic, commercial, safety, engineering and disposition authority.

03 CONTROLLED USE — Summarize work-order trends

USE WHEN	ASK	VERIFY
Maintenance leaders need help organizing approved, de-identified work-order data.	Group by equipment class, failure description, downtime band and repeat occurrence. Identify patterns to investigate; do not diagnose root cause.	Classifications, counts, periods, missing records and apparent patterns.

BOUNDARY Maintenance experts retain diagnosis and work authorization.

04 CONTROLLED USE — Prepare a material-shortage response

USE WHEN	ASK	VERIFY
Supply-chain and operations teams need a structured internal briefing.	Use confirmed inventory, delivery dates, affected orders, options, owners and open questions. Separate facts from assumptions.	Inventory, purchase orders, lead times, contracts and decision authority.

BOUNDARY Qualified owners make customer and commercial commitments.

05 CONTROLLED USE — Prepare a quality problem-solving meeting

USE WHEN	ASK	VERIFY
A quality team needs a neutral structure for reviewing verified observations.	Organize facts, unknowns, containment, process stages, evidence to collect and questions. Do not determine root cause or disposition.	Lot, part, date, measurement, defect, containment and source records.

BOUNDARY Quality authorities retain containment, disposition and corrective action.

ALWAYS RETAIN Prompt or instruction; approved inputs; tool/version; output; corrections; reviewer; accepted record; exception; and decision.

PUBLIC AUTHORITY NIST MEP AI implementation guidance • NIST AI RMF Playbook

CONTROL BOUNDARY

TREAT THE AI TOOL AS ANOTHER SYSTEM

Keep the information below out of unapproved tools. Any exception requires specific authorization for the system and data use, plus applicable legal, security and contractual controls.

BLOCK UNLESS SPECIFICALLY AUTHORIZED

01 Trade secrets and proprietary formulas

02 Customer drawings and technical specifications

03 Controlled, classified or export-controlled information

04 Source code, credentials, network diagrams or security details

05 Employee, applicant, customer or supplier personal information

06 Nonpublic quality, incident, contract, pricing or financial records

FIVE CONTROL GATES

01 **IDENTITY**
Organization-approved accounts, least privilege and MFA

02 **CONNECTIONS**
Disable unnecessary connectors, plug-ins and integrations

03 **RETENTION**
Confirm storage, deletion and model-training settings

04 **INVENTORY**
Record tool, owner, use, version, access and review date

05 **INCIDENT**
Route AI issues through the existing incident process

BEFORE PROCUREMENT OR PILOT USE

DATA FLOW
What enters, where it is processed, what is retained and who can retrieve it

MODEL USE
Whether inputs or outputs train a shared or future model

ACCESS
Accounts, roles, connectors, APIs, logging and administrative rights

EXIT
Export, deletion, revocation, incident support and vendor offboarding

PUBLIC AUTHORITY NIST Generative AI Profile • CISA AI resources

PROCUREMENT AND CONTROL

REVIEW THE TOOL IN PROPORTION TO ITS REACH

The depth of review should match access, autonomy, integration, data sensitivity and potential consequences. A public chatbot and an OT-connected system do not receive the same review.

REVIEW AREA	REQUIRE BEFORE PILOT	ESCALATE OR STOP
OWNER + PURPOSE	Named owner; approved use; excluded uses	No accountable owner or vague enterprise-wide purpose
DATA	Allowed/prohibited data; retention; training; deletion	Unknown processing, reuse, location or deletion
IDENTITY + ACCESS	SSO/MFA; least privilege; admin and user roles	Shared accounts, excessive privilege or weak revocation
INTEGRATION	Connectors, APIs and permissions are necessary and bounded	Unneeded mail, file, code, network or OT reach
OUTPUT + REVIEW	Qualified reviewer, source traceability and correction path	Unverifiable output or automation bias
SECURITY	Logging, incident reporting, vulnerability handling and support	No usable logs, notification or response commitment
CONTRACT	Confidentiality, IP, audit, change notice and exit terms reviewed	Unclear ownership, unilateral change or lock-in
MEASUREMENT	Baseline, success criteria, burden and stop conditions recorded	Success defined only by adoption or plausible output

DECISION **APPROVE WITH BOUNDARY / RETURN FOR EVIDENCE / REJECT**

PUBLIC AUTHORITY NIST AI Risk Management Framework • CISA AI resources

EVIDENCE GATE

MEASURE THE RESULT AND THE BURDEN OF CONTROL

A pilot is not successful merely because the system produced an answer. Count verification, correction, exception handling and operator effort.

MEASURE	BASELINE	PILOT RESULT	DECISION QUESTION
QUALITY	Current error and revision level	Verified error and revision level	Did the reviewed output meet the defined standard?
TOTAL TIME	Current completion time	Draft + review + correction time	Did total time improve after review was included?
CONSISTENCY	Current variation	Observed variation	Was performance repeatable across users and examples?
RISK	Known process risks	New or changed risks	Did data, safety, cyber, quality or workforce concerns increase?
USEFULNESS	Current user experience	Participant feedback	Would trained users choose the controlled workflow again?

DECISION LOGIC

CONTINUE	REVISE	PAUSE	STOP
Measured value + controlled residual risk	Promising result; boundary or burden needs work	Evidence gap or correctable control failure	Value unproven or control is inadequate

IMMEDIATE STOP Sensitive-data exposure; unauthorized access or action; deceptive or unsafe output; repeated material error; loss of evidence; or a new event affecting safety, quality, security, privacy, rights or customer commitments.

PUBLIC AUTHORITY NIST AI RMF Playbook • OSHA Safety and Health Programs

FROM PILOT TO OPERATING CONTROL

CONNECT THE WORKFLOW TO POLICY, TRAINING AND INCIDENT RESPONSE

The designated operations decision-maker authorizes routine use after qualified control owners review the pilot. Expansion needs a documented boundary.

01

POLICY

Approved tools and uses; restricted data; human oversight; records; procurement; incident routing.

02

TRAINING

Role-specific practice on boundaries, source checking, disclosure, correction, escalation and stop authority.

03

OPERATIONS

Named owner; version and access inventory; review cadence; quality/safety integration; change control.

04

INCIDENT

Known route for exposure, unsafe output, unauthorized action, deception, outage, vendor change or lost evidence.

OPERATING-CONTROL TEST

Routine use requires measured value, accepted residual risk, named authority, trained roles, operating controls and a live incident route.

EXPANSION RECORD

SCOPE

Workflow, users, sites, systems, data and excluded actions

EVIDENCE

Baseline, pilot results, correction burden, incidents and residual risk

AUTHORITY

Operations decision-maker, qualified reviewers, conditions, effective date and stop owner

CHANGE

Triggers for re-review: model, vendor, data, access, integration, law, policy or observed performance

PUBLIC AUTHORITY NIST AI Risk Management Framework • CISA AI resources

COMMON QUESTIONS

USE THE METHOD, NOT A SHORTCUT

The answers below preserve the human decision boundary while routing readers to the practical next step.

WHAT IS A MANUFACTURING AI PLAYBOOK?

A planning and operating resource for selecting a use case, protecting data, testing a workflow, requiring human review, measuring results and deciding whether a controlled pilot should continue.

CAN AI WRITE MANUFACTURING SOPS?

It may help draft a template, but qualified personnel must verify every instruction against approved procedures, equipment documentation, safety requirements, quality systems and applicable standards.

WHAT SHOULD NOT GO INTO A PUBLIC TOOL?

Trade secrets, customer drawings, controlled technical data, source code, credentials, security details, employee information, proprietary formulas or other nonpublic records without specific approval.

CAN AI MAKE SAFETY OR QUALITY DECISIONS?

It should not independently approve safety procedures, engineering changes, process parameters, equipment operation, product release or quality disposition. Authorized people retain responsibility.

WHERE SHOULD A MANUFACTURER BEGIN?

With one bounded, lower-risk workflow using approved, nonconfidential information; a named owner; a small pilot group; human review; and measures of quality, total time, repeatability, burden and risk.

NEXT **Record scope, owner, approved data and tool, baseline, reviewer, evidence, stop conditions and next review. Use the pilot charter and decision record on the next page, or copy the web fields into the existing planning record.**

PUBLIC AUTHORITY NIST MEP AI implementation guidance • EEOC AI and the ADA resources

WORKING RECORDS

RECORD THE PILOT AND THE DECISION

Copy these fields into the existing planning record. Complete essential approvals before testing. Retain the decision and its evidence.

PILOT CHARTER: COMPLETE BEFORE TESTING

Problem and current process	[problem; current process and baseline source]
Scope and exclusions	[workflow; users; test period; excluded tasks and actions]
Accountable roles	[owner; qualified reviewer; approving roles; stop authority]
Tool, information and access	[approved tool/account; permitted information; access and retention limits]
Baseline and quality	[current effort and quality standard; comparison limits]
Verification and evidence	[authoritative sources; reviewer checks; acceptance evidence]
Local targets and stop limits	[quality; full staff effort; cost; risk; conditions that pause testing]
Decision and review dates	[decision meeting; owner; next review and change triggers]

PILOT DECISION: RETAIN WITH THE EVIDENCE

Disposition and effective date	[continue / revise / pause / stop; effective date]
Decision authority	[decision owner and qualified reviewers]
Evidence and uncertainty	[baseline comparison; errors/corrections; full effort; cost; incidents; limits]
Next boundary and review	[permitted next scope; corrections/approvals; owner; review date/triggers]

Pause for an unresolved essential approval or an output that a qualified reviewer cannot check. Expansion requires a new decision for the changed boundary. This record is MAIN implementation synthesis.

PUBLIC AUTHORITY + LOCAL EVIDENCE

GROUND THE METHOD IN AUTHORITATIVE SOURCES

MAIN developed the operating model, decision logic, examples and working tools in this publication. The public sources below support specific risk, safety, security, workforce and manufacturing boundaries.

01 **NIST AI Risk Management Framework**
AI risk governance and lifecycle foundation

02 **NIST AI RMF Playbook**
Suggested actions for operationalizing the framework

03 **NIST MEP: AI implementation strategies**
Manufacturing-specific implementation considerations

04 **NIST 2026 AI/ML roadmap for smart manufacturing**
Research and technical direction for smart manufacturing

05 **NIST Trustworthy AI in Critical Infrastructure concept note**
Emerging critical-infrastructure risk context; not a final profile

06 **NIST Generative AI Profile**
Generative-AI risks and actions extending the AI RMF

07 **CISA artificial intelligence resources**
Cybersecurity and secure-use context

08 **OSHA Recommended Practices for Safety and Health Programs**
Worker participation, hazard prevention and program improvement

09 **EEOC resources on AI and the ADA**
Employment and disability-related compliance context

10 **NIST Cybersecurity Framework 2.0**
Cybersecurity governance and control context for connected systems

REVIEW RECORD

OWNER
Named local owner controls applicability, links and currency.

CHECK
Review when law, standards, vendor, model, data, access, policy or observed performance changes.

STOP
Do not rely on an expired, inaccessible or superseded source without qualified review.

Contextual photo: Freek Wolsink / Pexels. Not a MAIN facility or AI deployment.

PUBLIC AUTHORITY