Nirmitee Health

NIRMITEE / HEALTH

AI AGENT OPPORTUNITY & ROI KIT

A practical decision framework for Pharma & MedTech leaders choosing which workflows to automate - and which not to.

24 WORKFLOWS | 100-POINT SCORING | ROI CALCULATOR | 90-DAY PILOT

2026 EDITION | DECISION KIT + EDITABLE WORKBOOK

Designed for Digital, Medical Affairs, Clinical, Regulatory, Quality, Commercial and Technology leaders.

01 / EXECUTIVE BRIEF

The question is not “Can we build an agent?”

It is “Which workflow deserves automation, what value is realistic, and what evidence must exist before we scale?”

Life-sciences organizations face an expanding catalogue of AI possibilities: literature surveillance, regulatory intelligence, study operations, complaint triage, commercial account analytics and internal knowledge workflows. But a large list of ideas is not an investment strategy.

This kit converts ideas into a disciplined portfolio decision. It pairs workflow-level opportunity scoring with hard risk/data gates, a transparent capacity-based ROI model, and a 90-day pilot plan. It is deliberately built for teams that need to persuade both business sponsors and cross-functional reviewers.

The three decisions this kit helps you make

DECISION	WHAT TO ESTABLISH	WORKBOOK OUTPUT
1. Where to start	Which repeatable workflows have measurable value, accessible data, and realistic integration paths?	Use-Case Library + Opportunity Score
2. Whether it pays	What capacity is released after review time, adoption, implementation and recurring costs?	ROI Calculator + Settings
3. Whether to scale	Can the workflow meet accuracy, safety, traceability, adoption and change-control thresholds?	Pilot Charter + KPI Tracker + 90-Day Roadmap

> **A high opportunity score is not permission to deploy. Risk, data rights, validation and human accountability remain separate approval gates.**

Who should use it

Chief Digital/Information Officers; AI/Automation leaders; Heads of Medical Affairs, Clinical Operations, Regulatory, Quality, Safety and Commercial Excellence; RWE and data-platform teams; transformation PMOs and procurement leads.

What is included

Guide: 24 workflow patterns, prioritization rubric, ROI math, controlled architecture, validation gates, pilot charter, procurement questions and 90-day plan. **Workbook:** 10 editable tabs with formulas, scoring, templates and traceable source notes.

02 / CURRENT CONTEXT

Why prioritize with evidence in 2026

The regulatory conversation is moving toward context-specific, lifecycle-aware AI. That makes workflow selection and proof of value inseparable.

DATE / SOURCE	VERIFIED DEVELOPMENT	PRACTICAL IMPLICATION
14 Jan 2026 FDA + EMA	Jointly published ten guiding principles of good AI practice in drug development, including human-centric design, risk-based assessment, defined context of use and lifecycle management.	For AI used in medicine development, document intended use, people, data, evaluation and change management before scaling.
4 Jun 2026 EMA / HMA	Published the 2025 AI Observatory report describing activity in knowledge mining, productivity, automation and data handling/analysis.	Knowledge work and workflow automation are real areas of exploration, but use-case-specific validation is still needed.
Jan 2025 FDA	Issued draft guidance on AI credibility when models support drug/biologic regulatory decision-making. The guidance is draft and not binding.	Treat regulatory-evidence use cases differently from internal productivity workflows; assess applicability.
2023-2026 NIST	AI Risk Management Framework and Generative AI Profile offer voluntary governance and evaluation resources; NIST notes AI RMF revision work in 2026.	Use a structured risk lens, but do not mistake voluntary frameworks for certification.

These sources do not imply that every internal agent is a regulated medical device or that every productivity workflow is subject to drug-submission guidance. Applicability depends on the intended use, system role, jurisdiction, data and consequences of failure. Source URLs appear on the final page.

Our commercial thesis

> **The next valuable AI conversation is not “Show me a chatbot.” It is “Show me the workflow, the baseline, the human control, the systems it touches, and the measurable outcome.”**

03 / AUTOMATION DESIGN

Not every problem needs an AI agent

Use the least complex automation that can meet the workflow requirement reliably.

PATTERN	BEST FIT	LIMIT / DECISION TEST
Rules / edit checks	Deterministic validation, field checks, simple routing, known conditions.	If a stable rule solves it, avoid generative AI.
RPA / workflow automation	Repeatable, predictable multi-system tasks with fixed steps.	Use when interpretation is limited and interfaces are stable.
AI copilot	Summarization, drafting, retrieval and recommendations for a human operator.	Often the best first production design in regulated work.
Agent with tools	Multi-step retrieval, classification and orchestration across systems under explicit permissions.	Use only when dynamic planning materially improves results.
Autonomous action	Bounded low-impact tasks with reliable tests, rollback and monitoring.	Require strong evidence and approved authority boundaries.

A five-question selection test

- **Repeatability:** Is the task performed frequently enough to justify design and maintenance?
- **Judgment:** Does it require interpretation across unstructured information, or can rules solve it?
- **Observability:** Can you measure baseline time, quality, exception rate and downstream effects?
- **Control:** Can the AI stop at a human gate before high-impact actions?
- **Data:** Are source systems, rights, provenance and integration paths clear?

> **Default first move: build a source-grounded, draft-only copilot. Earn autonomy through demonstrated accuracy, controls and operational evidence.**

04 / WORKFLOW LIBRARY

24 candidate workflows / Part 1

Use these as starting points for discovery, not as promises of achievable performance.

ID	WORKFLOW	OUTPUT	HUMAN GATE	KPI
MA01	**Literature alert triage**	Ranked evidence digest	Medical reviewer	Source accuracy; review time
MA02	**Medical information draft**	Source-linked response draft	Medical information approval	Unsupported claims; draft time
MA03	**MSL insight synthesis**	Structured themes and signals	Medical Affairs validation	Coding time; theme precision
MA04	**Evidence-gap mapping**	Claims-to-evidence register	Evidence strategy sign-off	Gap recall; agreement
MA05	**Congress intelligence**	Traceable scientific digest	Scientific reviewer	Turnaround; coverage
CL01	**Site feasibility triage**	Cited site shortlist	Clinical Operations	Cycle time; site suitability
CL02	**Eligibility pre-screen**	Potential candidate flags	CRC / PI	Recall; false negatives
CL03	**eCRF discrepancy detection**	Candidate data queries	Data manager	Query precision; latency
CL04	**RWE source fitness**	Data quality/fitness memo	Epidemiologist	Assessment time; missingness
CL05	**ePRO follow-up orchestration**	Escalation suggestion	Site/clinical staff	Completion; correct routing
RA01	**Regulatory intelligence watch**	Applicability digest	Regulatory Affairs	Detection time; accuracy
RA02	**Submission evidence cross-check**	Claim-source exceptions	Regulatory reviewer	Defect detection; review time

In the editable workbook, each of these workflows also includes its trigger, likely source systems, primary failure risk and candidate KPI. Select only 2-3 for the first scoring workshop.

04 / WORKFLOW LIBRARY

24 candidate workflows / Part 2

Every workflow needs a defined trigger, input, output, human gate and measurable failure mode.

ID	WORKFLOW	OUTPUT	HUMAN GATE	KPI
RA03	**Change impact assessment**	Affected market/product map	RA / QA	Missed impacts; review time
QA01	**Complaint intake triage**	Classification suggestion	Quality / Safety	Escalation recall
QA02	**CAPA trend detection**	Pattern candidates	Quality manager	Signal precision
QA03	**SOP training mapping**	Affected roles / courses	Quality / L&D	Training coverage
CO01	**Account whitespace**	Prioritized account opportunities	Sales + Compliance	Opportunity acceptance
CO02	**Tender monitoring**	Tender relevance + fit	Bid owner	Qualified tenders
CO03	**HCP pre-call briefing**	Approved-source briefing	Policy-trained user	Content compliance
CO04	**Portfolio cross-sell**	Evidence-linked bundle ideas	Sales rep	Attach rate; adoption
IT01	**IT knowledge retrieval**	Cited resolution suggestion	Service agent	Resolution time
IT02	**Data mapping assistant**	Proposed field mappings	Data architect	Mapping accuracy
IT03	**Contract obligation extraction**	Obligation register draft	Legal / Procurement	Extraction recall
IT04	**SOP Q&A with citations**	Versioned source answer	Document owner	Answer accuracy

In the editable workbook, each of these workflows also includes its trigger, likely source systems, primary failure risk and candidate KPI. Select only 2-3 for the first scoring workshop.

05 / PORTFOLIO SELECTION

The 100-point opportunity model

Seven weighted criteria. Scores are relative decision aids, not validated predictions of value.

CRITERION	WEIGHT	HOW TO SCORE
Volume / frequency	20%	How often does the workflow occur?
Manual effort	15%	How much time and rework exist today?
Business / clinical value	20%	What outcome would materially improve?
Data readiness	15%	Are sources accessible, permissioned and reliable?
Integration ease	10%	Can systems be connected without major redesign?
Change readiness	10%	Is there a committed owner and adoption path?
Risk exposure (inverse)	10%	How consequential is an erroneous output?

Scoring method

Rate each criterion from 1 to 5. For the first six, 5 means stronger opportunity or readiness. For Risk exposure, 5 means more consequential failure, so the model reverses it as (6 - risk). Weighted score = 20 x SUM(weight x adjusted rating). The maximum is 100.

Interpretation

75-100: pilot candidate if no hard gate. **60-74:** validate data or process first. **Below 60:** backlog. Any risk score of 4-5 or data-readiness score of 1-2: control review regardless of score.

> **The workbook lets you edit weights. It flags control reviews automatically. Always document the evidence behind each rating.**

06 / EXAMPLE DECISIONS

A shortlist is not an approval list

These examples are hypothetical scoring exercises, not measured industry benchmarks.

WORKFLOW	SCORE	AUTOMATIC GATE	WHY
Literature alert triage	**84**	**PILOT CANDIDATE**	High volume, accessible corpus, bounded output and medical review.
Medical information draft	**78**	**CONTROL REVIEW**	High potential value, but scientific claim and safety exposure require stronger controls.
MSL insight synthesis	**73**	**VALIDATE FIRST**	Permissions, structured input and reviewer agreement need evidence.
RWE source fitness	**Varies**	**CONTROL REVIEW**	Data rights/readiness and inferential consequences dominate initial scoring.

Create a balanced pilot portfolio

- **One quick win:** low-to-moderate risk, high frequency, stable source data and a visible human review gate.
- **One strategic workflow:** high value but more integration work, selected only if the organization can fund validation.
- **One hold:** a promising use case whose data rights, patient impact or regulatory consequences are not yet resolved.

Evidence required to move from “candidate” to “approved pilot”

A named sponsor, 30-50 representative baseline cases (or a justified sample plan), permitted data access, a documented human gate, measurable quality thresholds, critical failure scenarios, a technical integration map, a rollback route and a signed pilot charter.

> **A good portfolio deliberately contains “not now” decisions. Stopping an unsuitable agent early is a successful outcome of prioritization.**

07 / RISK AND AUTONOMY

Four levels of agent authority

Give the system the minimum authority needed for the business case.

LEVEL	PERMITTED BEHAVIOR	EXAMPLE	REQUIRED CONTROL
L0 / Observe	Read, summarize and flag only.	Literature digest	Source checks and human interpretation.
L1 / Draft	Prepare a record or action, without submission.	Draft MI response	Reviewer edits and approves.
L2 / Propose action	Prepare system updates that a human explicitly approves.	Suggested eCRF query	Role-based approval and audit trail.
L3 / Bounded execute	Execute low-impact reversible actions within policy.	Internal task routing	Permissions, monitoring, rollback, incident path.

Five hard gates

- **Intended use:** clearly define who uses the output, for what decision, and what AI is prohibited from doing.
- **Data permission:** verify rights, sensitive data handling, source provenance, retention and vendor terms.
- **Safety and quality:** test high-consequence errors and define escalation, override and suspension thresholds.
- **Traceability:** retain appropriate source IDs, prompt/model/version, actions, approvals and changes.
- **Change control:** plan for model updates, prompt revisions, retrieval changes and new system integrations.

The appropriate control set depends on context. The FDA-EMA January 2026 principles address AI in drug development; NIST AI RMF is voluntary. Neither should be treated as a one-size-fits-all certification scheme.

> **Do not use the pilot to make unreviewed patient eligibility decisions, safety escalation decisions, regulatory submissions or promotional claims.**

08 / ARCHITECTURE BLUEPRINT

Build around the workflow, not the model

A useful agent is a controlled service embedded into existing systems and human decisions.

STAGE	CONTROLLED FLOW	EVIDENCE / ARTIFACT
1. Trigger	New document, inquiry, case, data refresh or scheduled review.	Event ID, timestamp, workflow owner.
2. Retrieve	Access only approved sources through permissioned connectors.	Source IDs, versions, retrieval filters.
3. Interpret	Classify, compare, summarize or recommend with bounded prompts.	Prompt/model/tool versions; confidence limitations.
4. Validate	Apply deterministic rules, source checks and exception tests.	Test results, citations, policy checks.
5. Approve	Human accepts, edits, rejects or escalates material outputs.	Named reviewer, decision, timestamp.
6. Act	Update the approved system of record only within granted authority.	Write operation, API response, audit trail.
7. Monitor	Track quality, rework, adoption, failures and model/data changes.	KPI log, incident record, change log.

The minimum implementation evidence pack

Use-case card + data and permissions sheet + risk/oversight matrix + test protocol + release approval + monitoring plan + change log. The workbook provides starting templates for the use-case card, data/system map and KPI evidence.

Interoperability is part of ROI

Every additional manual handoff can erase apparent time savings. Before modeling value, list the systems of record, API/export capability, identifiers, refresh frequency, permissions and write-back constraints. A model that produces a good draft but requires double entry may have a weak business case.

> **Preferred pattern: permissioned retrieval -> source-grounded draft -> deterministic checks -> human approval -> auditable write-back.**

09 / FINANCIAL DISCIPLINE

ROI without inflated “AI savings”

Model released capacity after adoption, human review and recurring costs. Do not assume capacity is equivalent to cash.

ILLUSTRATIVE INPUT	VALUE	WHY IT MATTERS
Active users x cases / month	15 x 100	1,500 monthly cases.
Baseline / eligible / adoption	25 min / 75% / 80%	Only adopted eligible cases count.
Gross reduction / QA time	60% / 4 min	Net 11 minutes saved per eligible case.
Loaded labor cost	75 per hour	Illustrative currency units.
Annual recurring cost	60,000	Platform + maintenance + enablement.
One-time build	45,000	Integration, validation and rollout.

Worked result - hypothetical, not a benchmark

METRIC	ILLUSTRATIVE RESULT
Eligible adopted cases / year	**10,800**
Net hours released / year	**1,980**
Annual capacity value	**148,500**
Annual net capacity value	**88,500**
Simple payback	**6.1 months**
Three-year ROI	**98%**

Calculation: eligible adopted cases = 15 x 100 x 12 x 75% x 80% = 10,800. Net minutes saved/case = (25 x 60%) - 4 = 11. Hours = 10,800 x 11 / 60 = 1,980. Annual capacity value = 1,980 x 75 = 148,500. Annual net = 148,500 - 60,000 = 88,500. Payback = 45,000 / (88,500/12) = 6.1 months.

If those 1,980 hours are not redeployed, used to avoid future hiring, or reflected in reduced expenditure, the financial benefit may be substantially lower than the capacity estimate. Revenue, quality and patient outcomes require separate, defensible models.

> **The editable calculator includes 3-year ROI, discounted net value, simple payback and a break-even time-reduction threshold.**

10 / PILOT ECONOMICS

How to pressure-test the business case

A robust investment case must survive conservative assumptions and reveal what would falsify it.

TEST	QUESTION	ACTION
Baseline reliability	Were time and error rates measured, or estimated by stakeholders?	Use a representative time-and-motion sample.
Eligible population	What share of cases is safe and suitable for AI assistance?	Exclude exceptions and high-impact cases unless controlled.
Human review burden	How much time is spent verifying and correcting outputs?	Include incremental review in every calculation.
Adoption	Will people actually use the new workflow?	Pilot with representative users; measure eligible-case adoption.
Total cost	Are validation, data integration, security, support and vendor changes included?	Include one-time and recurring costs.
Benefit realization	Is the value capacity, avoided hires, reduced overtime, revenue or error reduction?	Separate each benefit and avoid double counting.

Minimum sensitivity review

Calculate the result under three scenarios: **conservative** (lower adoption, smaller time reduction, higher review burden), **base** (measured pilot assumptions), and **stretch** (only if supported by data). The workbook models a base case; duplicate the tab for scenario comparison.

Define “stop” conditions before launch

- Critical unsupported scientific claim or missed safety escalation.
- Unapproved personal or confidential data exposure.
- Failure to reconstruct a material AI output and reviewer action.
- No meaningful net time reduction after human review.
- Sustained user rejection or unacceptable integration burden.

> **A negative pilot is useful when it prevents an expensive rollout. The decision should be based on measured outcomes, not demo quality.**

11 / EDITABLE TEMPLATE

The one-page pilot charter

Copy these fields into a project intake or use the editable Pilot Charter tab.

CHARTER FIELD	REQUIRED ANSWER
Business owner	Name / function / decision authority
Intended use	What exact decision or task is supported?
In scope	Trigger, users, eligible cases, systems
Out of scope	Prohibited actions, populations, decisions
Data rights	Source owner, permission, retention, classification
AI behavior	Read, draft, recommend or execute?
Human oversight	Who reviews what, and when?
Success metric	Measured baseline, target, sample
Critical failures	What must never happen?
Validation	Representative cases, reviewers, acceptance
Traceability	Sources, versions, outputs, approvals
Fallback	Manual path, suspension and rollback
Costs	Build, integration, recurring, monitoring
Decision	Go / extend / stop authority and date

The objective is a falsifiable experiment: a narrow workflow, a clear baseline, a named owner and a pre-agreed threshold for stopping or scaling.

12 / EVIDENCE SCORECARD

Measure accuracy and adoption together

Choose primary and guardrail metrics before anyone sees the pilot results.

DIMENSION	PILOT METRIC	EXAMPLE ACCEPTANCE RULE
Scientific quality	Source citation accuracy	At least 98% in a defined blinded sample.
Scientific quality	Unsupported claim rate	At most 1%; adjudicate severity.
Safety / Quality	Critical failures	Zero unresolved critical events.
Productivity	Net time per eligible case	At least 30% improvement after review.
User value	Reviewer acceptance rate	At least 90% accepted without material rework.
Oversight	Escalation correctness	At least 95% of test exceptions correctly routed.
Traceability	Execution trace completeness	100% of required records reconstructable.
Adoption	Eligible-case usage	At least 75% during steady-state pilot.
Privacy	Unapproved data access	Zero confirmed events.
Operations	Availability	At least 99% during agreed operating hours.

Important: thresholds above are illustrative. A clinical, safety, regulatory or GxP workflow may require different validation, sampling, performance and governance. Set thresholds with qualified functional reviewers and document why they are adequate.

The go/no-go hierarchy

1. No unresolved critical failures. 2. Data and audit requirements met. 3. Scientific/operational quality threshold met. 4. Net value demonstrated. 5. Owner accepts ongoing monitoring.

> **Productivity cannot compensate for a failed critical safety, quality, privacy or traceability gate.**

13 / EXECUTION PLAN

A 90-day route from idea to decision

The goal is not a production rollout by day 90. It is a defensible decision about whether and how to proceed.

WEEKS	ACTIVITY	EVIDENCE
01-02	Select use case; capture baseline	Use-case card; baseline sample
03-04	Risk, data-rights and integration review	Approved risk tier; data/system map
05-06	Controlled prototype and architecture	Versioned draft-only prototype
07-08	Blinded tests, edge cases and defect review	Test report; issue register; retest
09-10	Limited supervised pilot	Release approval; reviewer training; KPI log
11-12	Monitor, refine, verify economics	Change log; actual costs; measured capacity
13	Cross-functional go / extend / stop	Signed decision and scale conditions

Recommended governance cadence

- **Weekly:** business owner + product + reviewer inspect quality, errors, adoption and exceptions.
- **At major gates:** Quality/Regulatory/Safety/Privacy/IT sign off as applicable to intended use.
- **At day 90:** Finance validates costs and benefit realization; sponsor decides next investment.

Exit outcomes

GO: quality and guardrails passed; positive measured business case; monitoring owner assigned. **EXTEND:** promising value but insufficient sample or unresolved noncritical issues. **STOP:** critical failure, weak value, unavailable data or disproportionate controls.

> **A pilot without a signed scale/stop decision is a demo, not an operating-model experiment.**

14 / THREE WORKFLOW BLUEPRINTS

What a bounded agent actually looks like

These are design examples. Each requires context-specific review and validation.

A. Medical literature surveillance

Trigger: New peer-reviewed publication or guideline update

Workflow: Retrieve approved sources -> deduplicate -> identify relevant change -> generate source-linked digest -> route to medical reviewer.

Control: No automated medical recommendation or external scientific communication. Human confirms claims and source fidelity.

Measure: Coverage, citation correctness, reviewer time, false negatives.

B. eCRF discrepancy suggestion

Trigger: New or modified study data record

Workflow: Read protocol and form metadata -> run deterministic edit checks -> flag potential inconsistencies -> draft query for data manager.

Control: Data manager approves queries. AI never alters source data or silently closes queries.

Measure: Query precision, turnaround, review effort, audit completeness.

C. MedTech account whitespace

Trigger: New distributor or hospital usage report

Workflow: Normalize usage -> match catalog and UoM -> identify gaps -> estimate addressable volume -> draft rep next-best-action.

Control: Sales/Compliance approves action. Claims, pricing, contract/GPO and data permissions are checked.

Measure: Opportunity acceptance, qualified pipeline, rep time, false matches.

The same architectural pattern - retrieve, interpret, check, approve, act, monitor - can support different functions, but risk and evidence requirements must change with the context.

15 / BUYER DUE DILIGENCE

12 questions to ask AI vendors

A good demo is not a substitute for documented operating controls.

#	QUESTION
1	Which model(s), tools and subprocessors power this workflow?
2	Can we control or approve model, prompt and retrieval-source updates?
3	What customer data is stored, retained, logged or used for training?
4	Can we enforce least privilege, data classification and tenant isolation?
5	Can each material output be linked to source IDs and versions?
6	What is the test methodology for factual accuracy and failure modes?
7	How are unsupported claims, prompt injection and unsafe tool actions handled?
8	Can a human reject, edit, override and escalate every material action?
9	How are audit trails, change records and evidence exported?
10	What operational monitoring, incident handling and rollback exist?
11	Can the workflow fall back safely if the model or API is unavailable?
12	What happens to our data, evidence and workflows if we exit the vendor?

Red flags: claims of “100% accurate AI,” no documented model dependencies, unbounded autonomous write access, opaque data retention, no versioned testing, no human override, or a value case that counts all saved minutes as cash.

> **Ask the vendor to demonstrate a failure case and the resulting human escalation, not only a perfect success path.**

16 / DECISION SUMMARY

From workflow ideas to investable AI

The strongest AI opportunities are not always the most technically impressive.

Five takeaways

- **Start with a workflow and baseline.** The unit of value is a repeatable task with an owner, not an agent persona.
- **Use a portfolio, not a popularity contest.** Prioritize value, data, integration, change and failure impact together.
- **Count the full cost of assistance.** Human review, maintenance, monitoring and adoption belong in ROI.
- **Keep authority proportional to evidence.** Draft first; execute only after control and validation are demonstrated.
- **Make scale conditional.** The pilot must produce a quality record, a capacity/value record and a signed decision.

> **A better first AI investment is a workflow that creates measurable value and earns trust - not an autonomous demo that cannot be governed.**

A useful next step

Nirmitee Health supports Pharma and MedTech organizations in mapping high-value workflows into connected AI and digital products - from data architecture and interoperability to controlled automation and production engineering. Where clinical evidence or study operations are involved, the wider Clinera ecosystem may be relevant.

Invitation: Book a 30-minute AI Workflow Opportunity Review. Bring one workflow, its approximate monthly volume and the systems involved. The outcome is a practical discussion of value, architecture, human controls and pilot feasibility - not a product pitch.

Suggested CTA: Explore an AI Workflow Opportunity Review | nirmitee.health

17 / REFERENCES AND SCOPE

Source notes

Primary regulatory and standards sources. Research checked 28 October 2026.

FDA + EMA - Good AI Practice in Drug Development, Jan 2026
https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development

EMA / HMA - 2025 AI Observatory Report, 4 Jun 2026
https://www.ema.europa.eu/en/about-us/how-we-work/data-regulation-big-data-other-sources/artificial-intelligence

FDA - AI Credibility for Drug/Biologic Regulatory Decisions, draft Jan 2025
https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

NIST - AI Risk Management Framework, 2023; 2026 revision underway
https://www.nist.gov/itl/ai-risk-management-framework

NIST - Generative AI Profile, 26 Jul 2024
https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

EMA - AI in Medicinal Product Lifecycle Reflection Paper, Sep 2024
https://www.ema.europa.eu/en/about-us/how-we-work/data-regulation-big-data-other-sources/artificial-intelligence

What is original to this kit

The 24-workflow catalogue, 100-point prioritization rubric, sample scores, illustrative financial assumptions, control gates, pilot acceptance criteria and 90-day plan are Nirmitee Health planning frameworks. They are not sourced industry benchmarks or endorsed regulatory methods.

Important limitations

This kit is educational and is not legal, regulatory, clinical, financial or quality-system advice. It does not establish that a particular AI system complies with any law or standard. Determine applicable obligations by jurisdiction, intended use, system classification, GxP context and organizational policy. Financial outputs are scenarios and should be validated with measured baselines and realized benefit mechanisms.

© 2026 Nirmitee Health | Knowledge Library | Share with attribution.