CAPMINDS

A TECHNICAL PLAYBOOK

NIH, Federal, and Healthcare Innovation Grant Winners

How to Convert Funding Into a Real Healthcare Product

Move from an approved research aim to a defined product, credible evidence, a deployable architecture, and a purchasing decision. For healthcare CTOs, principal investigators, product leaders, and hospital IT teams responsible for what happens after the award.

INSIDE THE PLAYBOOK

- Funding boundaries and product scope
- Clinical evidence and regulatory decisions
- Architecture, integration, and secure delivery
- Pilot operations, procurement, and commercialization

01 / OPERATING MODEL

Start with the award, then define the product

Choose the correct funding lane

An NIH research award may support a scientific aim without funding a market-ready product. SBIR/STTR adds program-specific eligibility and execution conditions. Other federal awards, contracts, cooperative agreements, and private innovation grants have their own terms. Record the actual instrument before borrowing rules from another program.

Create one authoritative baseline

Collect the Notice of Award, approved aims, budget and justification, award amendments, human-subjects conditions, reporting dates, and relevant funding notice.

The product owner translates these into delivery work; the PI and grants office confirm alignment. A backlog item is not authorized merely because it supports the company strategy.

Use five evidence gates

- G0: scope and spending authority.
- G1: user need, intended use, and data access.
- G2: tested architecture and risk controls.
- G3: approved pilot with measurable endpoints.
- G4: operating ownership and purchasing path.

Each gate needs an accountable approver and linked evidence, not a percentage-complete slide.

Read this playbook by decision

Pages 3-6 establish scope and evidence. Pages 7-12 address engineering dependencies. Pages 13-16 connect budget, pilot operations, adoption, and commercialization. The dates you enter are project commitments; they are not universal regulatory or delivery timelines.

Funding lanes: apply the instrument actually awarded

Instrument	Execution boundary	First control
NIH research grant	Approved research aims; commercial work is not automatically covered.	Map aims, costs, changes, and approvals.
SBIR / STTR	R&D; toward commercialization; eligibility and work-share rules apply.	Confirm phase, work split, and supplier role.
Federal contract / ARPA-H OT	Contract or negotiated agreement governs; OT is not a grant.	Track deliverables, milestones, payment, IP terms.
Private / hospital innovation award	Sponsor-specific agreement; federal rules do not automatically transfer.	Verify data access, IP, permitted spend, and exit.

Apply this chapter

Proceed when the team can identify the governing terms and the next accountable decision.

COMPANION EXCEL TAB / Award baseline

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [1], [2], [25], [26], [27]

02 / FUNDING CONTROLS

Make every work package traceable

Treat the award budget as a set of permitted activities with evidence obligations, not an unrestricted engineering balance.

Build a scope-to-cost ledger

For each work package, record the approved aim, deliverable, cost category, staff effort, supplier role, and acceptance evidence.

Separate research work, implementation work, and commercial work in the ledger so the grants team can determine which funding source supports each activity.

Detect scope drift before a sprint

A new patient population, a changed scientific objective, a new subaward, or a materially different intervention may need review.

A technically small change can have a large award impact. Capture the request, rationale, budget effect, affected approvals, and decision before committing expenditure.

Example: expand an RPM prototype

The approved study evaluates feasibility of collecting home measurements. The roadmap adds autonomous medication recommendations.

That introduces a different claim, risk profile, and validation burden. Keep the new function outside the funded baseline until the appropriate owners resolve scope and regulatory implications.

Document the basis for spending

Where applicable, federal cost principles require costs to be necessary, reasonable, allocable, and documented. Retain time records, invoices, and the rationale linking costs to award work.

Do not assume that unused budget permits a new purpose or that an extension automatically supplies more money.

Worked example: synthetic RPM feasibility project

Work package	Funded-scope hypothesis	Acceptance evidence
Collect home BP readings	Supports approved feasibility aim; grants owner confirms.	Source IDs, units, time, and patient mapping tested.
Nurse review queue	Supports the specified workflow evaluation.	One task per accepted event; disposition logged.
Autonomous medication advice	Excluded from this illustrative baseline.	Separate scope, regulatory, and evidence decision.

Apply this chapter

Commit funds only after required approvals and cost treatment are resolved.

COMPANION EXCEL TAB / Scope and costs

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [1], [2], [3]

03 / PRODUCT DEFINITION

Choose one workflow worth changing

Observe the current path

Map the trigger, user, input, decision, handoff, and recovery path. For an abnormal home reading, determine who receives it, how patient identity is confirmed, who calls the patient, and what happens after hours. Interview frontline staff and the person accountable for unresolved work.

Specify a bounded MVP

Define one intended user, one care setting, one population, one primary decision, and the minimum data required.

Keep excluded uses explicit. A useful first release may surface validated readings and document follow-up; risk prediction and automated advice can remain separate future capabilities.

Separate four value tests

- Clinical: does the workflow support the intended care outcome?
- Operational: can staff execute it within capacity?
- Technical: can data arrive with usable meaning and provenance?
- Commercial: does an identified buyer have a purchasing route?

Passing a usability session answers only part of the product question.

Measure the work you introduce

Count review tasks per patient-day, time to disposition, missing data, duplicate alerts, and unresolved items at shift change.

An attractive interface can still increase clinical workload. Set a baseline with the pilot site before choosing targets; do not borrow unverified improvement percentages.

Worked example: define usefulness before features

Dimension	Illustrative project choice	Evidence required
Population and user	Adult outpatient cohort; nurse review team.	Site-approved eligibility and staffing model.
Minimum product	Validated BP display, review queue, disposition record.	Observed end-to-end workflow and usability tests.
Excluded capability	Diagnosis, autonomous dose changes, 24/7 emergency response.	Product claims and training reflect exclusions.

Apply this chapter

Approve the MVP only when a real workflow owner accepts its scope.

COMPANION EXCEL TAB / Product definition

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [5], [6]

04 / REGULATORY STRATEGY

Write the intended use before the claim

Assess each function separately

List data capture, display, calculation, prediction, recommendation, and control functions. For each, identify users, clinical purpose, input, output, and consequences of failure.

A product can contain both non-device and device functions; describing it as a platform does not settle classification.

Test CDS assumptions carefully

Do not assume that clinician-facing software automatically qualifies as non-device CDS.

Evaluate the applicable FDA criteria, including whether the clinician can independently review the basis of a recommendation. Record the rationale and unresolved questions instead of assigning a blanket exemption.

Align claims with evidence

Separate intended use from marketing language. A claim to organize readings is different from a claim to diagnose deterioration.

Freeze proposed claims before designing validation, then route changes through product, clinical, and regulatory review. Avoid claims broader than the tested population or setting.

Obtain feedback before expensive work

For uncertain device pathways, plan specialist review and appropriate FDA engagement before locking the study and architecture.

Track classification, submission route if applicable, required evidence, and dependencies. Research use or pilot status does not by itself settle regulatory obligations.

Worked example: claim-to-function review

Proposed function	Question to resolve	Artifact to retain
Display validated readings	Does the intended use remain a display function?	Function-specific regulatory rationale.
Flag out-of-range readings	What decision is supported, by whom, and on what basis?	Rule, rationale, failure analysis, intended user.
Recommend treatment	Does this introduce device oversight or change the pathway?	Regulatory assessment before implementation.

Apply this chapter

Do not approve clinical claims while the supporting regulatory rationale is unresolved.

COMPANION EXCEL TAB / Regulatory review

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

What will it take to turn your funded idea into a working product?

Explore how CapMinds helps translate clinical requirements into a focused MVP, practical architecture, and a clear development plan.

Explore Product Development ›

Sources: [7], [8], [9]

05 / EVIDENCE DESIGN

Design the pilot to answer a decision

A successful demonstration proves that software can run. A useful pilot establishes whether the intended workflow deserves the next investment.

Separate verification from validation

Verification asks whether the implementation meets its requirements: unit conversion, timestamp handling, authorization, calculations, and failure behavior. Validation asks whether the product meets intended user needs in the intended setting. Clinical performance claims may require additional clinical evidence.

Classify the activity before enrollment

Work with the institution to determine whether the activity is research, quality improvement, or another category and which reviews apply. For device investigations, assess applicable FDA requirements separately. Calling a project a pilot does not remove human-subjects or device-study obligations.

Predefine the evaluation

Specify the population, eligibility, site, comparator or baseline, primary endpoint, collection window, missing-data handling, and analysis plan. Determine sample size and uncertainty with qualified study support. A convenient sample of enthusiastic users rarely supports a broad clinical claim.

Use a claim-to-evidence table

Example: a routing claim needs delivery and disposition evidence; a diagnostic claim needs an appropriate clinical reference standard and performance analysis. Record dataset version, inclusion rules, adjudication, analysis code, and confidence intervals where relevant. Keep exploratory findings separate from predefined endpoints.

Worked example: requirements linked to tests

Requirement	Test fixture	Pass condition
REQ-01: no duplicate task	Submit source event 1001 twice.	One accepted reading and one review task.
REQ-02: quarantine identity conflict	Map one external ID to conflicting patients.	No clinical routing; exception is visible.
REQ-03: preserve temporal meaning	Deliver an older reading after a newer one.	Clinical order uses effective time; receipt time retained.

Apply this chapter

Start the pilot after required reviews and a signed evaluation plan.

COMPANION EXCEL TAB / Evidence and tests

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [9], [10], [11]

06 / ARCHITECTURE

Separate clinical state from transport

Build around observable state transitions so a delayed message, failed API call, or repeated event cannot silently change care.

Use explicit boundaries

A practical design separates ingestion, validation, identity resolution, normalized storage, workflow orchestration, and user-facing APIs. Keep source payloads or retrievable provenance where permitted. A normalized observation should retain source ID, patient mapping, code, unit, effective time, receipt time, and transformation version.

Persist before acknowledging

For asynchronous ingestion, acknowledge after durable acceptance. Carry a correlation ID through the pipeline.

Use an outbox or equivalent consistency pattern when a database state change must publish an event. Design consumers for duplicate delivery; do not assume exactly-once behavior across systems.

Model failure as a product state

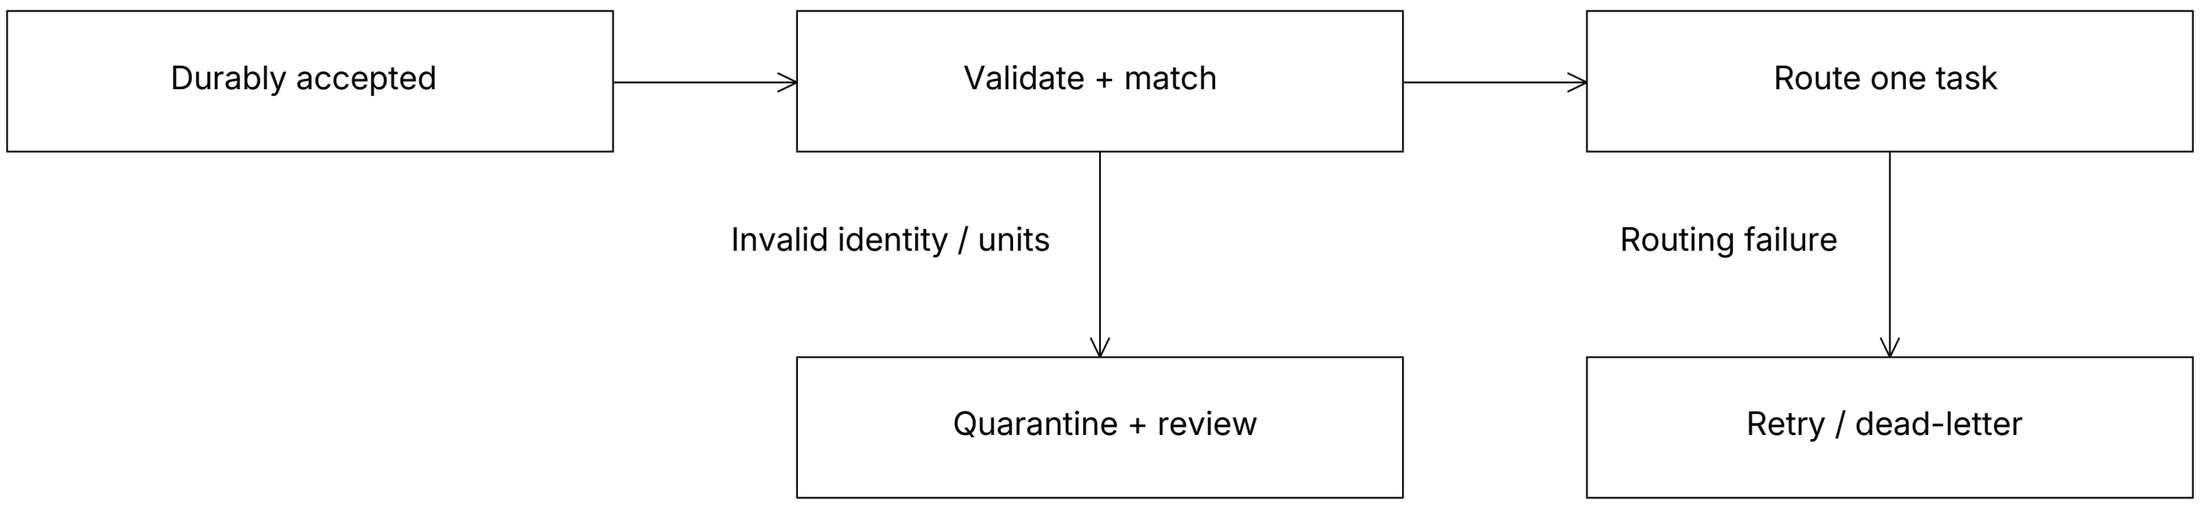
Distinguish received, validated, quarantined, routed, acknowledged, and resolved. Quarantine ambiguous patient matches and invalid units. A dead-letter queue needs an owner, replay controls, and an audit trail. Replaying an event must not generate a second clinical task for the same actionable condition.

Choose boundaries by failure mode

A modular monolith can be appropriate when one team owns the workflow and scale is modest.

Separate services when isolation, independent scaling, or organizational ownership justifies the operational cost. Document latency budgets, recovery objectives, dependency failures, and the safe degraded experience.

Failure-state map / persist each transition and retain the source event ID



Replay after correction retains lineage. Retry uses backoff; duplicate events never create a second task.

Apply this chapter

Demonstrate a failed dependency and recovery path before expanding deployment.

COMPANION EXCEL TAB / Architecture risks

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

07 / INTEROPERABILITY

Prove the EHR contract with real tests

FHIR support is a startingpoint. The integration contract is the intersection of the specification, the EHR implementation, and the customer configuration.

Build an operation-level matrix

For every workflow, record the resource, read/search/write operation, search parameters, profile, terminology, scopes, and environment. Confirm production onboarding and customer authorization. A successful Patient read does not establish permission to write an Observation or create a clinical task.

Exercise authorization and limits

Use the SMART configuration and supported authorization flow. Test expired tokens, denied scopes, revoked access, pagination, rate limiting, and retry behavior. Store secrets outside code and logs. Distinguish an empty clinical result from a failed request so users do not mistake an outage for no history.

Resolve meaning before throughput

Test code systems, units, effective versus recorded time, timezone offsets, missing values, and corrected results. Keep source identifiers and provenance. Patient matching must use an approved rule with an exception path; demographic similarity alone should not silently select a record.

Make write-back a separate release gate

Define reconciliation for partially completed writes and discover how the EHR handles duplicates and corrections. Test the resulting clinician workflow in the destination system. Measure response-time distributions and error rates in the agreed environment; never publish a universal EHR latency or capacity promise.

Illustrative integration acceptance matrix

Operation	Contract to verify	Failure test
Patient search / identity mapping	Customer-approved match rule; permitted identifiers.	Ambiguous match routes to manual resolution.
Observation read/search	Supported code, date, category, pagination, and scope.	Expired token and empty result remain distinguishable.
Observation write (if required)	Create/update support; profile; authorization; correction rules.	Retry after timeout reconciles before creating again.

Apply this chapter

Sign off the exact operations and failure cases required by the pilot.

COMPANION EXCEL TAB / EHR integration

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [15], [16]

08 / DATA RIGHTS AND PRIVACY

Know which data you may actually use

Access to a dataset, permission to delivercare, permission to conduct research, and permission to train a model are different questions.

Map purposes and recipients

Document what is collected, why, where it is processed, who can access it, and how long it is retained. Identify the covered entity, business associate, subcontractor, and other roles as applicable. Determine HIPAA applicability from the actual relationship and activity, not from the app category alone.

Design practical access boundaries

Use tenant-aware authorization on every request and background job. Test object-level access, export permissions, support access, and service accounts. Keep PHI out of unnecessary logs and analytics. Document how support staff obtain time-limited access and how that access is reviewed.

Control secondary use

Verify consent, authorizations, contracts, and applicable policy before reusing data for training, evaluation, or commercial analytics. Separate production care data from development datasets. Do not treat a vendor contract or a BAA as blanket authorization for every secondary purpose.

Plan lifecycle operations

Specify retention, deletion, exports, legal holds, and backup treatment with the customer. Cloud service eligibility and a signed BAA do not configure the product securely. Record responsibility for risk analysis, encryption, incident handling, access review, and subcontractor oversight.

Illustrative normalized observation contract

Field group	Example / type	Invariant
Identity	tenant_id, source_id, event_id, revision	Unique by tenant + source + event + revision.
Clinical content	patient_ref; code/system; value; UCUM unit	Validated mapping; preserve original unit and value.
Time and provenance	effective_at, received_at, transform_version	Timezone-aware timestamps; reconstruct transformation.

Apply this chapter

Use real data only after the permitted purpose and responsibilities are documented.

COMPANION EXCEL TAB / Data responsibilities

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [12], [13]

09 / SECURE DELIVERY

Make security evidence part of each release

Threat-model the clinical workflow

Trace identity, API, ingestion, storage, support, and integration boundaries.

Prioritize cross-tenant access, forged measurements, exposed tokens, tampered rules, and workflow denial of service. Connect each threat to a testable control and a named operational owner.

Secure the build chain

Require reviewed changes, protected branches, reproducible release identifiers, dependency inventories, secret scanning, and vulnerability triage.

Preserve the relationship between deployed version, tested build, configuration, and infrastructure definition. A scan report without disposition and retest evidence is incomplete.

Test recovery, not just prevention

Exercise restore, credential revocation, tenant isolation, rollback, and a compromised integration credential.

Define who can suspend a function and how clinicians continue safely. Set vulnerability response targets based on risk and contractual commitments, then track evidence that the process works.

Apply device requirements where relevant

For applicable medical devices, use current FDA cybersecurity guidance and assess section 524B obligations.

Plan appropriate security documentation, software component visibility, vulnerability management, and update processes. Do not present NIST alignment or a penetration test as a certification of overall compliance.

Release evidence: show the control working

Risk	Executable test	Release evidence
Cross-tenant access	Tenant A token requests tenant B observation.	Denied response; no data; attributable audit record.
Secret exposure	Scan repository and built artifacts with seeded test secret.	Detection + blocked release + rotation procedure.
Dependency vulnerability	Introduce known vulnerable fixture in test pipeline.	Triage record, remediation/reasoned exception, retest.

Apply this chapter

Release only with accepted residual risks and an executable response path.

COMPANION EXCEL TAB / Release security

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Will your healthcare software be ready for a hospital’s technical review?

See how CapMinds approaches secure development, EHR integration, and deployment requirements before they become expensive obstacles.

Explore Our Engineering Approach >

Sources: [14], [17]

10 / AI AND DECISION LOGIC

Keep the model inside a controlled workflow

If the product uses AI, its operating boundary must be as explicit a sits evaluation dataset. If it does not, apply the same discipline to rules and calculations.

Version the whole decision system

Record model or rules version, prompt where applicable, retrieval source, preprocessing, terminology mapping, threshold, and output schema. A stable model with changed input normalization is a changed system. Preserve enough lineage to reconstruct the basis of an output without unnecessarily duplicating sensitive data.

Constrain clinical actions

Separate drafting from signing and recommending from executing. Use structured outputs, permission checks, explicit clinician confirmation where appropriate, and a safe abstention path. A low-confidence result should enter a defined review state rather than disappear or produce fabricated certainty.

Evaluate failure slices

Test missing fields, stale observations, conflicting histories, out-of-distribution inputs, and subgroup performance. For generative outputs, assess unsupported statements and whether cited evidence supports the text. Use a held-out evaluation set with documented leakage controls; avoid treating polished examples as validation.

Control post-release change

Monitor input shifts, operational errors, overrides, and relevant performance signals. Predefine escalation thresholds and rollback. Route new models, thresholds, populations, and claims through impact assessment and any required review. Do not assume that a previous authorization covers every subsequent modification.

Decision logic: keep the tested system reproducible

Change	Impact question	Required record
Unit conversion changes	Do alert thresholds receive equivalent values?	Regression dataset and rule outputs.
Threshold or model changes	Does risk, clinical performance, or intended use change?	Impact assessment and approval; applicable regulatory review.
Missing source data	Can the product distinguish unavailable from normal?	Explicit unavailable state and tested review path.

Apply this chapter

Approve the complete decision system and its failure behavior, not just a model score.

COMPANION EXCEL TAB / Decision logic

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [7], [8], [9]; implementation controls are recommended design practices

11 / PARTNERS AND OWNERSHIP

Buy deliverables you can operate

Supplier selection should expose delivery boundaries, IP rights, funding constraints, and operational dependencies before the contract is signed.

Classify the engagement first

Ask the grants office to determine the appropriate relationship and applicable requirements. For NIH SBIR/STTR, check work-performance conditions, location requirements, and prior approvals. Do not assume that a vendor invoice makes an arrangement permissible or that a new subaward can begin without review.

Resolve ownership in writing

Identify pre-existing IP, newly developed IP, third-party components, licensing limits, and rights to modify and maintain the software. Coordinate subject-invention reporting with institutional or company IP owners. Open-source use needs license review; paid development does not automatically settle every ownership question.

Require a technical handover

Specify source repository access, build instructions, deployment definitions, data schemas, API contracts, tests, dependency licenses, credentials handoff, and runbooks. Acceptance should demonstrate a clean build and deployment under the recipient-controlled environment, not only a vendor-hosted demo.

Tie payment to observable outputs

Use milestone acceptance tied to requirements and evidence. Include remediation, change control, subcontractor visibility, data return, and exit assistance. Evaluate the supplier with a bounded technical exercise using synthetic data and agreed criteria before assigning a critical integration or clinical workflow.

Supplier scorecard: suggested weighting, not a procurement rule

Criterion	Weight	Evidence to inspect
Integration and failure handling	35%	Customer-relevant operation and recovery demonstration.
Secure delivery and traceability	35%	Build, test, dependency, release, and handover records.
Ownership and maintainability	30%	Recipient-controlled clean deployment; IP/license terms.

Apply this chapter

Sign only when award eligibility, ownership, and acceptance are explicit.

COMPANION EXCEL TAB / Supplier evaluation

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [2], [18], [19]

12 / BUDGET AND SCHEDULE

Fund the bottlenecks before the backlog

The critical path often runs through data access, institutional review, EHR onboarding, and evidence collection rather than feature implementation.

Estimate complete work packages

Include engineering, clinical input, integration onboarding, validation, security, hosting, support, and required reporting. Separate one-time setup from recurring cost. Identify which costs are supported by the award, another funding source, or customer payment; keep the allowability decision with the appropriate owner.

Use a transparent planning model

Illustrative only: $300,000 of deployable funding at a net $50,000 monthly burn implies six months of runway before timing effects and obligations. Two months of blocked access can consume $100,000 without generating pilot evidence. Use actual cash restrictions and committed liabilities in the real forecast.

Schedule external dependencies

Give every dependency a requested date, external owner, lead-time assumption, and fallback. A production credential, executed agreement, or approved protocol is a gate, not an engineering task marked almost complete. Track blocked days so the team can distinguish capacity problems from access problems.

Reforecast at each evidence gate

Track actual expenditure, committed expenditure, estimated remaining cost, and available funding by source. Do not disguise overruns with unapproved scope cuts that invalidate the study. Reduce nonessential features, renegotiate milestones where permitted, or secure additional funding before expanding commitments.

Illustrative delivery budget: $300,000 total; not an award allocation

Cost package	Planning amount	Assumption
Engineering + integration	$150,000 + $45,000	Bounded MVP; one EHR; actual onboarding priced separately.
Clinical evaluation + security / QA	$45,000 + $30,000	Study design and security scope still require quotes.
Hosting / support + planning headroom	$15,000 + $15,000	Headroom is not a chargeable grant cost by itself.

Apply this chapter

Proceed when the next evidence milestone is funded and its dependencies are owned.

COMPANION EXCEL TAB / Budget tracker

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [1], [3]; financial example is illustrative

13 / CLINICAL DEPLOYMENT

Launch a supported workflow, not a URL

Run end-to-end acceptance

Test representative journeys from data creation through clinician disposition and audit export. Include wrong-patient prevention, out-of-order data, corrected results, duplicate events, access revocation, unavailable APIs, and recovery. Verify both system behavior and the staff action expected after each failure.

Use a controlled launch

Choose an initial cohort consistent with the approved plan. Confirm training, access, fallback procedures, incident contacts, and evidence collection before activation. Review early issues at a predefined cadence. Expand only when the measured workload, safety conditions, and technical behavior support the next cohort.

Assign operational ownership

Define first-line support, technical escalation, clinical escalation, after-hours coverage, and the authority to suspend use. A monitoring dashboard is not a clinical response service. Confirm that users understand who is responsible for reviewing work and what remains in the existing care pathway.

Define a stop-and-reconcile path

Specify stop conditions such as a patient-identity defect or a critical authorization failure. On suspension, identify outstanding tasks, preserve evidence, notify accountable staff, and reconcile affected records. Resume only after correction, retesting, and the designated approval; a rollback alone may not resolve clinical work already created.

Pilot stop-and-reconcile runbook

Trigger	Immediate action	Restart evidence
Wrong-patient association	Suspend affected routing; clinical owner reviews exposure.	Correct mapping; identify affected records; retest.
Unresolved queue exceeds capacity	Activate agreed staffing or pause enrollment.	Backlog reconciled; safe capacity demonstrated.
Integration outage	Display unavailable status; use approved fallback.	Reconcile missing/repeated events; approve recovery.

Apply this chapter

Launch only when clinical and technical owners sign the same readiness record.

COMPANION EXCEL TAB / Pilot readiness

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [5], [6], [9]

14 / COMMERCIAL ADOPTION

Define the purchase decision before pilot exit

Identify how the product is bought

Distinguish an enterprise software budget, a service contract, a reimbursed clinical service, and another funding route. For reimbursement-dependent models, verify coverage, coding, payment, and provider requirements separately. An available code does not guarantee payment for the product or workflow.

Calculate the customer economics

Estimate implementation effort, interface fees, hosting, monitoring, support, clinical review time, training, and maintenance. Measure avoided work or other benefits against the baseline. Keep financial assumptions separate from observed results; do not turn an unmeasured workflow benefit into a guaranteed ROI claim.

Prepare a procurement evidence room

Organize architecture and data-flow diagrams, security responses, test summaries, applicable regulatory rationale, service responsibilities, agreements, and pricing assumptions. Name owners for open procurement questions. Build this incrementally so pilot completion is not followed by months of avoidable documentation work.

Agree an exit decision

Before the pilot ends, define the evidence threshold, decision-maker, budget process, and choices: purchase, extend with a specific question, redesign, or stop. A letter of interest is not a purchase commitment. Record the condition that would change the buyer decision rather than counting enthusiasm as adoption.

Illustrative business case: measure before claiming savings

Input	Example calculation	Decision implication
Baseline review effort	500 reviews/week x 4 minutes = 33.3 hours.	Collect representative baseline timing first.
Hypothesized future effort	500 x 3 minutes = 25 hours; 8.3 hours released.	Hypothesis only; workload and outcomes must be measured.
Financial interpretation	8.3 x $60/hour = about $500/week capacity value.	Not cash savings unless actual spending changes.

Apply this chapter

Advance commercially when evidence, ownership, and a purchasing route converge.

COMPANION EXCEL TAB / Purchase case

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

What’s standing between your healthcare pilot and its first customer?

Explore how CapMinds supports the integrations, deployment, and workflow improvements your product needs for real customer adoption.

Explore Implementation Support ›

Sources: [20], [21]

15 / NEXT-STAGE EXECUTION

Turn the next 90 days into decision gates

Days 1-30: remove ambiguity

Establish the award baseline, intended use, workflow owner, permitted data access, and initial architecture. Start external approvals and onboarding immediately. Exit with a funded scope and a register of unresolved dependencies. If access is blocked, use synthetic fixtures to test contracts while resolving the real dependency.

Days 31-60: produce engineering evidence

Implement the smallest end-to-end workflow, prove failure handling, and complete traceability for critical requirements. Exercise authorization, identity mapping, unit handling, and recovery. Review risk and intended-use changes before moving into any real-world evaluation.

Days 61-90: reach the next justified gate

Where approvals and readiness permit, execute the defined evaluation and collect evidence. Otherwise complete the remaining prerequisites and reforecast openly. Decide whether to expand, redesign, or stop based on the agreed criteria. Preserve findings and unresolved issues for the next funding or purchasing decision.

Use available support deliberately

Eligible NIH recipients can investigate TABA and other SEED support for specific commercialization gaps. The April 2026 TABA notice sets request limits of $6,500 per Phase I project and $50,000 per Phase II project. Confirm applicability and request procedures; these funds are not automatic.

Evidence gates: release decisions, not calendar promises

Gate	Required record	Decision owner
G1: scope and access	Approved baseline; intended use; permitted data access.	PI + product + grants / clinical owners.
G2: tested system	Critical requirements pass; accepted risks; recovery evidence.	Engineering + clinical + security reviewers.
G3/G4: pilot and adoption	Required approvals; pilot evidence; buyer and operating owner.	Site sponsor + product + procurement.

Apply this chapter

Leave the planning session with a named owner and evidence requirement for every gate.

COMPANION EXCEL TAB / Execution milestones

Record the decision, its supporting evidence, the accountable owner, and the next action in the separate planning and execution workbook.

Sources: [4], [22]

REFERENCE DESK

Primary sources for execution

[01] NIH: Prior Approval Requirements

[02] NIH: SBIR/STTR Administrative Requirements

[03] Federal cost principles: 2 CFR 200.403

[04] NIH: TABA Policy Update, April 20, 2026

[05] AHRQ: Workflow Assessment Toolkit

[06] AHRQ: Health IT Evaluation Toolkit

[07] FDA: Device Software Functions

[08] FDA: Clinical Decision Support Software

[09] NIH SEED: Digital Health Regulatory Guide

[10] IMDRF: SaMD Clinical Evaluation

[11] OHRP: Quality Improvement Activities FAQs

[12] HHS: Resources for Health App Developers

[13] HHS: HIPAA and Cloud Computing

[14] NIST: Secure Software Development Framework

[15] HL7: SMART App Launch

[16] ONC: API Certification Conditions

[17] FDA: Medical Device Cybersecurity Guidance

[18] NIH: Intellectual Property Policy

[19] NIST: iEdison

[20] CMS: Technology and Innovation

[21] CMS: Coding

[22] NIH SEED: Product Development Support

[23] FDA: Quality Management System Regulation

[24] FDA: Device Software Submission Content

[25] NIH: SBIR/STTR Eligibility

[26] ARPA-H: Submission FAQs and Award Instruments

[27] NSF: SBIR/STTR Funding Overview

APPLY THE CURRENT REQUIREMENT

For regulated device development, read [23] and [24] alongside the software and clinical-evidence sources. QMSR became effective February 2, 2026. The February 2026 cybersecurity guidance supersedes the June 2025 guidance. The 2024 NIH digital-health guide remains useful for engineering structure, but current FDA guidance controls the regulatory detail.

Award-specific terms, approved protocols, and applicable requirements determine the execution path. The engineering patterns and worksheets in this playbook are practical planning tools; they are not agency-issued forms.

CAPMINDS

YOUR NEXT PRODUCT DECISION

Bring one funded idea.

Leave with a clearer next step.

You do not need a finished requirements document to start a useful conversation. Bring the clinical workflow, the current stage of the product, and the dependency that is holding the next milestone.

Focus the discussion

- What must the product demonstrate next?
- Which data, integration, or engineering dependency blocks it?
- What evidence would make the next investment decision easier?

Discuss your healthcare product

Start with one workflow, one architecture question, or one delivery milestone.

Book a product discussion ›