The Note-Taker at Scale: Retention and Liability in Medical-Scribe Roll-Ups

A Valuation and Transaction Framework for Clinician Retention, Documentation Quality, Revenue Integrity, Data Governance and Liability in Ambient-AI Consolidation

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Medical-scribe roll-ups are becoming combinations of human documentation services, ambient artificial-intelligence software, clinical workflow integration, coding support and revenue-cycle operations. The category can reduce documentation burden and strengthen the clinical record. It can also concentrate liability. A buyer may inherit recordings, protected health information, model dependencies, clinician adoption risk, variable note quality, coding exposure, implementation obligations and customer contracts written for a different operating model.

This paper develops a transaction framework for testing whether an ambient-scribe business has become durable clinical infrastructure. The framework begins with intended use and the complete encounter-to-record workflow. It then connects purpose-specific data rights, model and software supply chains, representative evaluation, human oversight, privacy, security, regulatory classification, interoperability, implementation burden, clinician cohorts, enterprise contracts, revenue quality, gross margin, intellectual property, organisation and governance. The transaction analysis separates recurring software from people-intensive services, values retained clinician workflows, identifies liabilities and converts findings into price, consideration structure, contractual protection and a 180-day integration programme.

The control boundaries draw on current material from the World Health Organization, US Food and Drug Administration, Office of the National Coordinator for Health Information Technology, Centers for Medicare & Medicaid Services, US Department of Health and Human Services, Federal Trade Commission, European Union, UK health authorities, NIST, IMDRF, accounting and valuation standard setters, public transaction disclosures and peer-reviewed clinical studies.[1][2][3][4][5][6][7][8][9][10][11][12][13][14][15][16][17][18][19][20] Public transactions also show that clinical documentation assets can be acquired as strategic platforms, service combinations or capabilities embedded within broader healthcare technology businesses.[21][22][23][24]

A wholly hypothetical acquisition case illustrates a target with 9,000 licensed clinicians, 6,800 monthly active clinicians, 5,700 clinicians retained in productive use after twelve months and USD 54 million of contracted annual recurring revenue. Every adoption rate, price, cost, retention measure, margin, liability amount, multiple, synergy and valuation in the case is a modelling assumption created solely to demonstrate the method. None is a forecast, market benchmark, clinical claim or valuation opinion.

The paper concludes that scale is valuable when the combined business can reproduce safe and useful documentation across representative workflows, retain clinicians after novelty and subsidy fade, preserve lawful data use, maintain human accountability, support compliant coding and claims, and earn contribution margin after the full cost of deployment and review. Six figures and eight tables translate the

framework into an operating architecture, rights ledger, evaluation scorecard, clinician-retention funnel, liability chain, cohort economics, valuation bridge and ownership programme.

Clinical, coding, billing, privacy, employment, competition, tax, accounting and corporate-law obligations vary by product, customer and jurisdiction. Qualified clinical, regulatory, legal, coding, accounting, tax, insurance and valuation specialists should determine the requirements applicable to a specific transaction.

JEL Classification: G24, G34, I11, L86, O33

Keywords: ambient clinical documentation, AI medical scribe, healthcare software, mergers and acquisitions, clinician retention, documentation quality, liability, cohort economics

JEL Classification: G24, G34, I11, L86, O33

Keywords: ambient clinical documentation, AI medical scribe, healthcare software, mergers and acquisitions, clinician retention, documentation quality, liability, cohort economics

1. Define the acquisition decision and value perimeter

The buyer should begin with a precise acquisition decision: which capabilities, contracts, workflows, people, rights and liabilities are being purchased, and which economic benefits depend on post-closing change. A medical-scribe target may contain several businesses under one brand. Human scribes can attend encounters or work remotely. Transcription operations can convert recordings into draft text. Ambient software can produce structured notes. Coding modules can suggest evidence or classifications. Revenue-cycle teams can connect documentation to claims. Each component has different revenue quality, labour intensity, liability and scalability.

The value perimeter should distinguish current stand-alone value, contingent product value and buyer-specific synergy. Current value is supported by live customers, active clinicians, governed workflows, retained revenue and contribution after all delivery cost. Contingent value depends on products, integrations, evidence or migrations that remain unfinished. Buyer-specific synergy arises from distribution, infrastructure, workflow combinations or cost changes that the seller cannot realise independently. The same benefit should occupy one location in the valuation model.

Durable clinical infrastructure is a demanding test. The product must fit the encounter, produce a useful draft, preserve clinician review and attestation, integrate into the health record, support downstream users, remain available, protect data and operate under a controlled change process. Clinicians must continue using it after onboarding, initial enthusiasm and commercial subsidy. Customers must renew after considering security, integration, support, performance, liability and total cost.

The acquisition thesis should state the evidence that would disprove it. Rejection gates can include absent rights for a core processing purpose, clinically material errors without a reliable detection path, inability to reproduce reported performance, weak clinician retention, enterprise renewal dependent on discounts, negative mature cohort contribution, unresolved security exposure or liability that cannot be bounded through remediation and transaction structure.

Table 1. Medical-scribe roll-up value perimeter

Value component	Evidence unit	Primary valuation lens	Principal risk
ambient documentation software	active retained clinician and enterprise contract	cohort cash flow and supported market evidence	contracted seats exceed productive use
human scribe and review services	staffed workflow and customer	service cash flow	labour, quality and wage inflation
coding and revenue-cycle workflow	compliant completed task	incremental cash flow	unsupported coding or attribution

Value component	Evidence unit	Primary valuation lens	Principal risk
data and model capability	lawful right and productive use	incremental income or replacement evidence	restricted rights and double counting
integration estate	live interface and supported customer	avoided deployment cost and retention effect	EHR dependency and maintenance burden
buyer synergy	named initiative, owner and evidence gate	buyer-specific discounted cash flow	paid to seller before delivery

Proposed transaction architecture; rights, obligations and economics require contract-level verification.

2. Map the complete workflow before evaluating the model

Model accuracy has meaning only inside a workflow. Diligence should begin with direct observation, screen evidence, system logs and user interviews that reconstruct the current and proposed process. The map starts before the model receives data and ends after the final user action, downstream reconciliation and exception handling. It should show identities, systems, queues, handoffs, review points, wait states and failure recovery.

For ambient documentation, the workflow can include patient notification or consent, microphone selection, session start, speaker separation, audio transmission, transcription, structured extraction, note generation, clinician review, edit, attestation, coding, record write-back and deletion or retention. Each stage can fail independently. A strong transcript can still produce a poor note. A good note can still increase work when EHR navigation, template selection, attribution or signing is awkward.

The map should record time and quality for both the old and new process. Time saved during dictation may be offset by editing, searching, workflow interruptions or after-hours corrections. A product may transfer work from clinicians to medical assistants, coders, support teams or implementation staff. The investor should quantify total system work and identify who absorbs it.

Workflow fit also varies by specialty and encounter. A short primary-care visit, a paediatric consultation with several speakers, a mental-health encounter, an emergency department visit and a surgical consent discussion have different vocabulary, turn-taking, privacy and documentation structures. The diligence sample must represent the commercial expansion plan rather than the easiest current use case.

Figure 1. Healthcare AI voice and workflow map

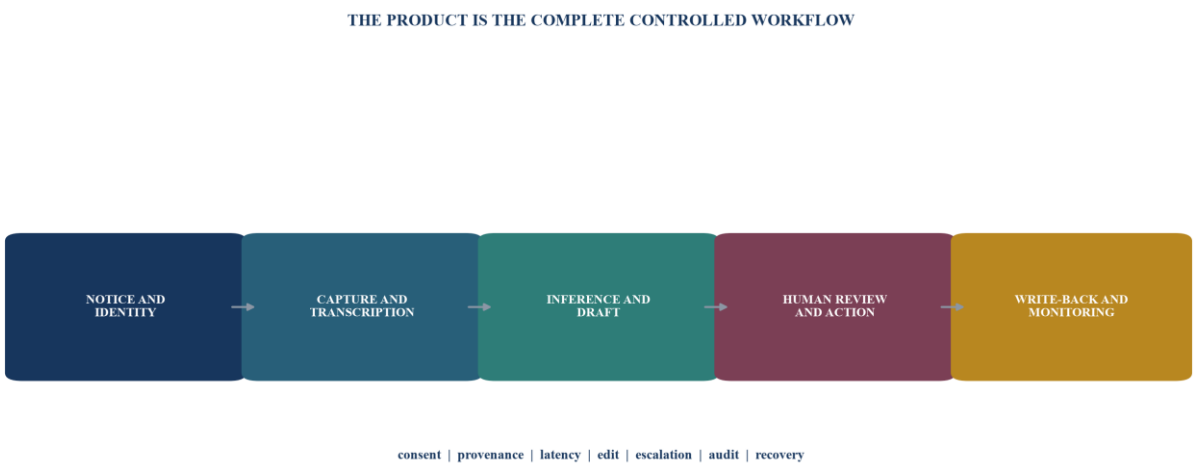

The map identifies the full control chain; actual workflows vary by product, care setting and jurisdiction.

3. Build the evidence room around production reality

The diligence room should be organised around claims that matter to customers and investors. For each claim, the company should provide the definition, measurement population, period, comparator, source system, exclusions, statistical method, owner and reproducible output. Evidence should distinguish internal tests, customer pilots, production deployments, peer-reviewed studies and marketing estimates.

Core materials include product specifications, intended-use statements, model cards, validation plans, test datasets, release histories, incident logs, clinical-safety files, privacy assessments, security reports, architecture diagrams, data-flow maps, customer contracts, business-associate agreements, subprocessor lists, implementation plans, support tickets, uptime records, EHR integration evidence, cohort revenue, invoices, cloud bills and customer-reference permissions. The request list should connect each document to an investment question.

Production logs deserve special attention. They can show encounter attempts, successful captures, abandonment, latency, output generation, review time, edit distance, acceptance, write-back, error, override and support events. The investigator should reconcile event definitions across versions. A metric may improve because a failed or difficult encounter disappeared from the denominator. A product release may change logging before it changes performance.

The evidence room should preserve negative evidence. Failed pilots, low-adoption cohorts, safety concerns, data-rights restrictions, unprofitable integrations and customer churn can reveal the limits of repeatability. Management should explain whether the issue arose from product capability, workflow design, customer readiness, pricing, implementation, governance or competition. The investor can then test whether the current plan addresses the cause.

Table 2. Evidence inventory for growth-investor diligence

Evidence domain	Minimum evidence	Reconciliation test	Decision use
product	specifications, versions, release notes	claims match shipped functions	scope and roadmap risk
model	datasets, evaluations, failures, monitoring	test population matches use	performance confidence
workflow	logs, observations, user references	total work and exceptions reconcile	adoption and value
rights	contracts, notices, consents, licences	every use has an identified basis	data and IP risk
deployment	statements of work, tickets, integrations	effort matches reported cost	margin and scalability
commercial	contracts, invoices, cohorts, churn	bookings bridge to cash and retention	revenue quality

Required depth should reflect the product's intended use, customer footprint and transaction materiality.

4. Create a data-rights ledger for every processing purpose

A healthcare AI company can possess data without holding every right required by its product and growth plan. Diligence should separate the right to receive data from the right to process it for a specific purpose. Recording, transcription, note generation, storage, quality assurance, support, model evaluation, product improvement, model training, benchmarking and customer export can require different contractual, privacy, consent and regulatory analyses.

The ledger should identify data category, subject, source, controller or covered entity, processor or business associate, subprocessor, location, purpose, legal or contractual basis, retention, deletion, export,

training use, audit right and incident responsibility. Audio deserves separate treatment from transcript and generated text. Audio can contain voices of patients, relatives, clinicians and bystanders. Derived features, embeddings, labels and model feedback may retain sensitivity or contractual restrictions.

HHS states that covered entities using business associates require written arrangements specifying permitted functions and safeguards for protected health information.[7] HHS cloud guidance also makes a business-associate agreement relevant when a cloud service creates, receives, maintains or transmits electronic protected health information on behalf of a covered entity or business associate.[8] Products outside HIPAA's covered-entity and business-associate perimeter may still face other privacy and breach obligations. The FTC's Health Breach Notification Rule is one US example for covered personal-health-record vendors and related entities.[19]

Rights should be tested through termination and change of control. The investor should determine whether a customer can revoke improvement rights, require deletion, prohibit model training, restrict cross-border processing or demand migration assistance. The company should demonstrate how data lineage supports those obligations. Contract language and system capability must agree.

Figure 2. Data-rights ledger across the healthcare AI lifecycle

Each processing purpose requires its own evidence; the figure does not determine a legal basis.

5. Understand the model and software supply chain

The product may combine speech recognition, speaker diarisation, language detection, specialty vocabularies, retrieval, large language models, rules, templates, coding engines and workflow orchestration. Diligence should identify which components are proprietary, licensed, open source, customer-specific or provided by a third party. The architecture should show where each component runs, which data it receives and which output it controls.

Dependency analysis should cover foundation-model terms, speech-model terms, hosting concentration, rate limits, regional availability, price changes, deprecation, safety filters and rights to generated output. A company can own the interface and workflow logic while relying on a small number of upstream providers for core inference. This may still be an attractive model, but the investor should price substitution cost and negotiating leverage.

The model register should record purpose, version, input, output, training or adaptation method, evaluation population, known limitations, approved use, monitoring and release authority. Deterministic

rules, statistical models and generative systems should remain distinct because their errors and test methods differ. A single product version may contain several model versions and prompt or retrieval configurations.

Change control is central. FDA's guidance on predetermined change control plans and lifecycle management shows why planned modifications, validation and monitoring matter for AI-enabled device functions.[4][5] The precise regulatory application depends on intended use and jurisdiction. Even when the software is outside device regulation, customers still need evidence that material changes are evaluated before release and that regressions can be detected, contained and reversed.

The investor should test the emergency substitution path. Can the company change an upstream model without losing function, rights, latency, geography or unit economics? Does it have regression tests and customer-notification rules? Architecture optionality is valuable only when it has been exercised or evidenced.

6. Evaluate performance at the workflow decision unit

Aggregate transcription accuracy is an incomplete measure. The investor should test the unit that affects workflow and care: a medication, allergy, symptom, diagnosis, duration, negation, speaker attribution, task, code, order element, urgency classification or patient instruction. A note can read fluently while omitting a clinically material fact or assigning it to the wrong speaker.

The evaluation plan should define intended use, population, setting, reference standard, adjudication, metrics, thresholds and failure response. It should include known positives, known negatives, difficult edge cases and out-of-distribution conditions. Samples should cover specialties, encounter types, accents, languages, noise, telehealth, interruptions, several speakers and relevant patient groups. The commercial roadmap should determine the test matrix.

Performance dimensions include factual completeness, factual correctness, unsupported content, negation, temporality, speaker attribution, structure, action extraction, latency and abstention. Workflow dimensions include review time, edit distance, acceptance, override, abandonment, support use and downstream correction. Safety dimensions focus on clinically material errors and whether controls detect them before action.

The company should preserve versioned results. A global score can improve while a material subgroup regresses. Thresholds should therefore be set by use case and harm, not by marketing convenience. Evaluation data should be sufficiently independent from development data, and customer-specific tuning should be separated from core-product performance.

Table 3. Model-evaluation design by decision unit

Decision unit	Example failure	Evaluation measure	Required review
medication	name, dose or negation error	exact and clinically weighted error	clinician or qualified reviewer
speaker	symptom assigned to wrong person	attribution precision and recall	encounter-level adjudication
summary fact	omitted or unsupported statement	factual completeness and support	blinded reference review
task	missed urgency or wrong routing	sensitivity, specificity and escalation	operational and clinical owner
workflow action	unauthorised or incorrect write-back	controlled-action success and recovery	safety and system owner

Metrics and thresholds require product-specific clinical and statistical design.

Figure 3. Model-evaluation scorecard from speech to safe workflow action

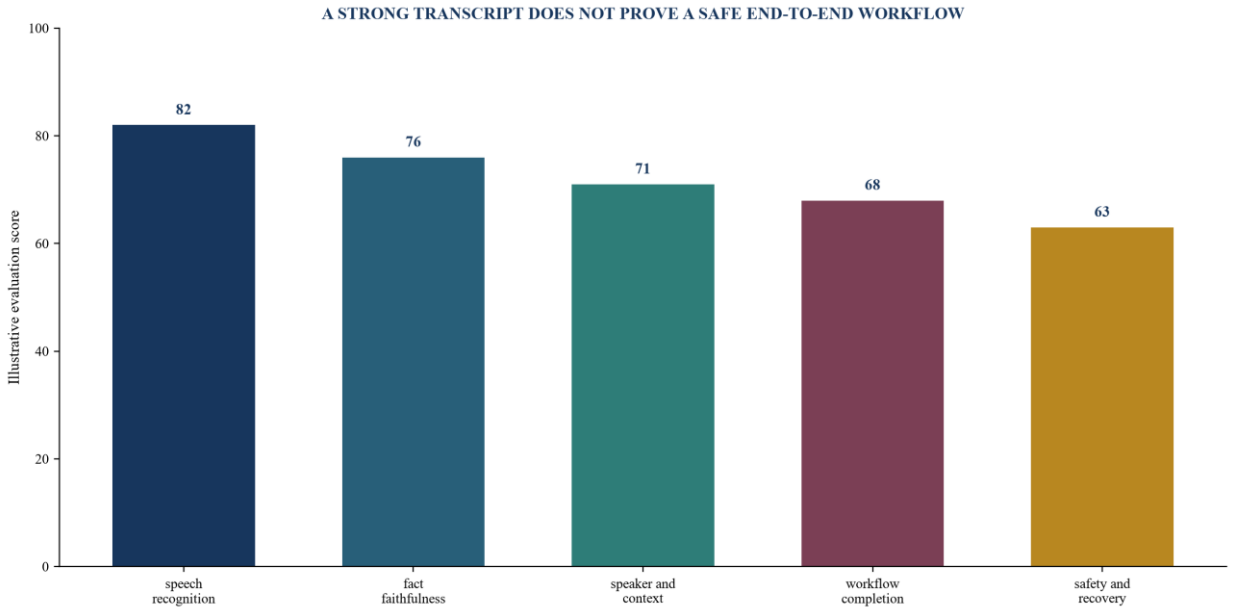

Scores are illustrative dimensions; investors should use product-specific metrics and thresholds.

7. Test clinical safety and human oversight as operating systems

Human oversight should be specified as an operating control. “Human in the loop” does not explain who reviews, what they see, how long they have, which errors they can detect, what happens under workload pressure or how the system responds when review fails. Diligence should map the accountable role, review interface, escalation route and evidence retained for each material output.

The review design should expose source context and uncertainty where useful. A user may need the relevant transcript span, speaker, timestamp, confidence, supporting record and reason for a suggested action. A long fluent draft can make errors harder to find. Interface design, alert burden and time pressure therefore affect the practical reliability of human oversight.

WHO guidance on large multi-modal models highlights risks including inaccurate outputs, bias, automation bias, privacy and cybersecurity.[2] NHS DCB0129 requires clinical-risk management for health-IT manufacturers, while DCB0160 addresses deployment and use by care organisations.[12] These standards illustrate the shared responsibility between product and adopter. The company should show a hazard log, clinical-safety case or equivalent risk file appropriate to its markets and use.

Incident management should connect product telemetry, customer reports, clinical review, engineering investigation and executive decision rights. Severity definitions should capture potential harm and near misses, not only confirmed injury. Material incidents should trigger containment, customer communication, root-cause analysis, corrective action, regression testing and monitoring.

The investor should test the degraded mode. If transcription, inference, EHR connectivity or identity resolution fails, can the user continue care safely? The workflow needs clear status, no silent write-back, retry rules and a manual route. A safe failure path can matter more than average model performance.

8. Examine subgroup performance, language and accessibility

Voice products operate across accents, languages, code-switching, speech impairments, age groups, microphone quality and clinical environments. An aggregate result can conceal weak performance in commercially important or vulnerable populations. Diligence should align subgroup testing with the company's deployed footprint and expansion plan.

The company should explain how it defines groups, obtains lawful evaluation data, establishes reference truth and avoids unstable conclusions from small samples. Performance differences should be assessed at clinically and operationally meaningful units. A modest word-error difference may become material when it affects medication names, negation, numbers or urgent symptoms.

Language support should distinguish user-interface translation, speech recognition, clinical vocabulary, note generation and workflow execution. A product may recognise conversational speech in a language while lacking validated specialty terms, templates or downstream EHR fields. Mixed-language encounters require separate evidence. Marketing labels such as multilingual should be decomposed into tested functions and settings.

Accessibility includes hearing, speech, cognitive, visual and motor needs as well as digital literacy. Patient-facing voice agents should handle repetition, interruption, uncertainty, interpreter involvement and transfer to a human. They should avoid creating a barrier for people who cannot or do not wish to use voice automation.

WHO's broader AI governance principles place inclusiveness and equity alongside autonomy, safety, transparency, responsibility and sustainability.[1] The investor should translate those principles into versioned subgroup evaluation, customer deployment criteria, monitoring and remediation. The commercial model should include the cost of supporting the populations it promises to serve.

9. Audit privacy, consent, security and resilience

The data-flow map should follow audio, transcript, generated content, identifiers, metadata, logs, support access, backups and model feedback across devices, networks, clouds, subprocessors and customer systems. Each transfer needs an owner, purpose, protection, retention rule and incident route. Security certificates and penetration tests are useful evidence, but they do not replace an architecture-specific threat model.

Voice creates particular risks. Always-on or ambient capture can include unintended speakers and sensitive discussion. Session start and stop controls, visible status, consent or notice, microphone permissions, local buffering and deletion behaviour require testing. The product should prevent one patient's content from entering another record and should surface identity uncertainty before write-back.

Security diligence should cover authentication, role access, least privilege, encryption, key management, secrets, tenant isolation, software supply chain, vulnerability management, secure development, logging, anomaly detection, backup, recovery and customer notification. Generative systems add prompt injection, data leakage, malicious content, retrieval poisoning and model-service dependency. NIST's AI RMF and Generative AI Profile provide a structured risk-management reference.[13]

For regulated device functions, FDA's 2026 cybersecurity guidance addresses design, quality-system and submission considerations.[5] Application again depends on the product and intended use. All companies should demonstrate resilience through tested recovery objectives, dependency failures and incident exercises.

Contractual promises should reconcile with operations. If the agreement promises regional processing, rapid deletion, subprocessor notice or a recovery objective, the company should show system evidence. A gap between sales language and production capability can become a renewal, liability and valuation risk.

10. Classify regulatory exposure by function and jurisdiction

Regulatory classification follows intended use and function, not the company's preferred category. A documentation tool that records and drafts a note can occupy a different position from software that recommends diagnosis, prioritises urgent symptoms or executes an order. Product modules and claims should therefore be classified separately.

In the United States, FDA's January 2026 clinical-decision-support guidance explains criteria for certain non-device CDS functions and gives examples distinguishing non-device from device software.[3] FDA's other digital-health guidance addresses AI-enabled device software, change control and cybersecurity.[4][5] ONC's HTI-1 rule adds algorithm-transparency requirements for predictive decision support supplied through certified health IT.[6] The diligence team should map whether the company is a certified health-IT developer, integrates with one or supplies functionality that triggers contractual flow-down.

In the European Union, the AI Act creates obligations based on system role and risk classification, including high-risk categories connected with regulated products and specified uses.[9] The European Health Data Space adds staged requirements for electronic health data and EHR systems while preserving the application of GDPR, medical-device and AI rules.[10] The effective dates and product impact require current jurisdiction-specific analysis.

In the United Kingdom, MHRA guidance and its software and AI as a medical device programme frame product classification and lifecycle expectations.[11] NHS deployments may also require clinical-risk, data-protection, technical-security, interoperability and usability evidence through DCB standards and assessment criteria.[12]

The regulatory matrix should record module, claim, user, action, jurisdiction, classification view, evidence, registration or certification, post-market duties and change consequence. Management's conclusion should be supported by qualified advice where material. A label such as administrative software does not settle the analysis.

11. Underwrite integration architecture and interoperability

Integration is part of the product. Enterprise value can depend on reliable identity, context launch, encounter matching, schedule access, documentation write-back, task routing, coding, consent, audit and analytics across customer systems. A vendor with strong models and fragile integration can face long sales cycles, high services cost and variable outcomes.

The architecture review should identify each interface, protocol, authentication method, data object, write permission, polling or event mechanism, latency, retry, error queue, observability and support owner. Standards such as HL7 FHIR can improve portability, but customer implementation, versions, profiles and proprietary constraints still matter.[20] A claimed integration should be evidenced in production rather than inferred from an API specification.

The investor should distinguish reusable connectors from customer-specific work. Reusable configuration can support scale. Bespoke transformations, template mapping, workflow redesign, security review and testing can create recurring labour even when the technical interface is standard. Implementation data should show hours by role, elapsed time, dependency waits, defects and post-go-live support.

Write-back deserves a separate control path. The system should confirm patient, encounter, document type, author, version and final approval. It should prevent duplicate or stale writes and preserve an audit trail. Failed writes need a visible queue and accountable recovery. Clinical content should not disappear into a generic integration log.

The diligence model should value interoperability through evidence: shorter time to live, lower implementation cost, fewer defects, higher adoption, lower support load and easier expansion. Connector count alone can overstate readiness when integrations are old, lightly used or maintained manually.

12. Convert implementation burden into a total cost to live

Customers buy a working change, not an isolated model. The total cost to live includes security and legal review, clinical governance, procurement, integration, configuration, workflow design, training,

communications, support, measurement and dual running. These costs can sit with the vendor, customer, implementation partner or EHR provider. The investor should capture the complete system cost and identify who pays.

Implementation should be decomposed into repeatable stages: qualification, contracting, data and security approval, technical build, configuration, clinical-safety review, pilot, training, go-live, monitoring and scale. For each stage, the company should provide median and distribution of elapsed time, internal hours, customer hours, blockers, rework and abandonment. Averages can conceal a long tail that consumes cash and leadership attention.

The cost curve should separate fixed platform investment, reusable product work and customer-specific labour. Early cohorts may be expensive while the company learns. Later cohorts should show evidence of declining effort for comparable complexity. Growth created by adding implementation headcount at the same rate as revenue has different economics from product-led repeatability.

Customer-side burden affects adoption and renewal even when it does not appear in vendor cost of revenue. Clinicians may spend time reviewing, correcting and learning. Informatics teams may maintain templates and governance. IT teams may operate interfaces. The customer reference should test whether realised benefit exceeds this burden.

Figure 4. Hypothetical integration cost curve by customer complexity

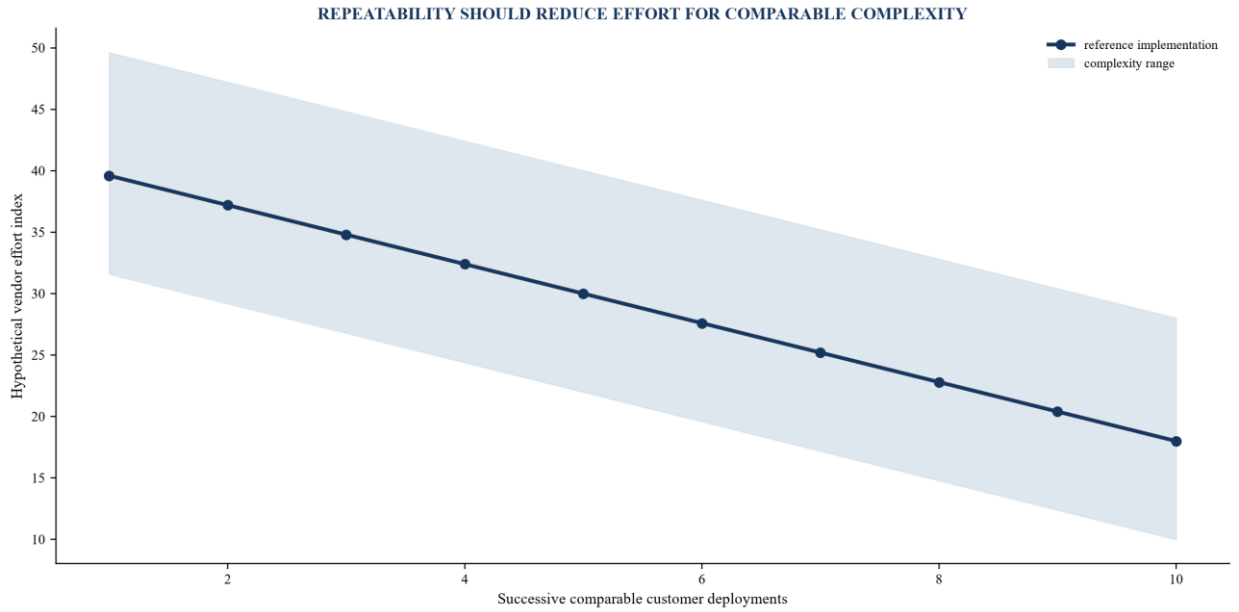

Values are hypothetical; live diligence requires customer-level hours, cash costs and elapsed-time evidence.

Table 4. Integration responsibility and cost matrix

Workstream	Vendor responsibility	Customer responsibility	Evidence for repeatability
security and privacy	architecture, controls, responses	risk review and approval	standard pack and cycle time
interface	connector and monitoring	access, testing and change control	reusable code and defect history
workflow	product configuration	operating design and ownership	configuration range and adoption
clinical safety	product hazards and controls	deployment hazards and governance	approved files and incidents
training and rollout	materials, enablement, support	attendance, local champions, policy	time to competent use

Workstream	Vendor responsibility	Customer responsibility	Evidence for repeatability
value measurement	telemetry and definitions	baseline, comparator, outcome review	reconciled benefit evidence

The matrix should be completed for each material customer segment and EHR environment.

13. Rebuild customer economics from cohorts and usage

Contracted annual value does not prove durable economics. The investor should build cohorts by start period, segment, product, geography, EHR, deployment complexity and contract type. Each cohort should show contracted value, live value, usage, expansion, contraction, churn, implementation cost, support cost, cloud and model cost, gross margin, cash collection and customer concentration.

The bridge begins with eligible users or encounters, activated users, active use, successful workflow completion, reviewed output and sustained use. Seats can overstate value when usage is low. Encounter pricing can overstate predictability when volume or specialty mix changes. Enterprise pricing can conceal unprofitable heavy users. The commercial model should reconcile pricing unit to value and cost drivers.

Retention should be measured on several bases. Logo retention captures relationship continuity. Gross revenue retention excludes expansion. Net revenue retention includes expansion and contraction. Usage retention shows whether the product remains embedded. Margin retention tests whether renewed revenue remains economical after inference, support and integration. Definitions and currency treatment should be consistent.

Expansion requires causal evidence. A customer may add clinicians because the pilot worked, because a separate budget opened or because the vendor discounted an enterprise conversion. Expansion into coding, inbox or patient voice can have a different buyer, risk review and support model. The investor should separate same-product scale from cross-product sales.

Customer references should test counterfactual behaviour. Would the customer renew at current price? What measurable workflow changed? Which users benefit and which do not? What burden remains? Which product failure would cause removal? Reference enthusiasm should be reconciled to invoices, usage and renewal decisions.

14. Test revenue quality, pricing power and gross margin

Revenue quality depends on enforceable contracts, deployment status, acceptance, usage, renewal, collection and delivery cost. Bookings can include pilots, cancellation rights, future modules or contingent volumes. Diligence should bridge signed contracts to recognised revenue, invoices, cash and live utilisation by customer.

Pricing should reflect the customer's value unit while remaining observable and governable. Per-clinician pricing is simple but can detach from encounter volume. Per-encounter pricing aligns usage but creates volume variability. Enterprise pricing supports budget predictability while requiring controls for heavy usage. Outcome-linked pricing can be attractive where the measure, baseline and attribution are credible. Hybrid structures may fit different workflows.

Pricing power should be tested through renewal history, discounts, competitive replacements, procurement concessions and willingness to expand. A high list price with extensive implementation credits, free pilots or bundled modules may create a weaker net price. The investor should calculate effective price after incentives and service commitments.

Gross margin should include speech, foundation-model, storage, observability, EHR fees, support, clinical operations, implementation, customer success and third-party royalties where attributable to delivery. Capitalised development and allocation policy should be understood separately. Upstream inference price

declines can improve margin, while usage growth, larger context, richer outputs or redundancy can offset the benefit.

The company should demonstrate margin by cohort and product. Mature customers should reveal the steady-state support and compute profile. New products may have lower margin while evidence and workflow mature. The valuation model should reflect the time and investment required to reach target margin rather than applying a software benchmark by label.

Figure 5. Hypothetical customer cohort economics bridge

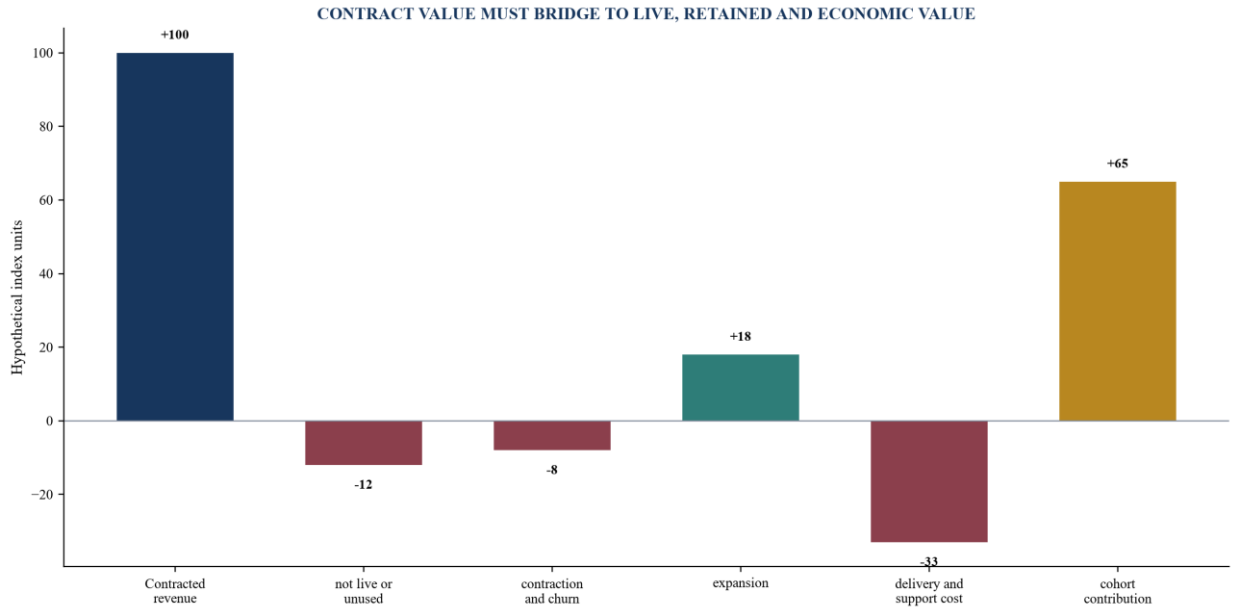

Values are hypothetical index units and do not represent a market benchmark or forecast.

15. Separate clinical evidence from marketing claims

Evidence strength depends on design. A vendor case study can establish feasibility and customer experience. A pre-post study can describe change while remaining exposed to selection, timing and concurrent initiatives. A matched observational study can improve comparison while retaining confounding. A randomised trial can strengthen causal inference for its tested product, users, period and outcomes. No single study proves every specialty, customer or product version.

The investor should build a claim-to-evidence matrix. Each clinical, workflow and economic claim receives an intended population, endpoint, study design, version, comparator, result, limitation and permitted commercial wording. Evidence for documentation time does not automatically prove reduced burnout, higher patient capacity, better note quality or improved revenue.

The 2025 rapid review of ambient AI scribes found only six eligible real-world studies from a much larger search set and reported sparse evidence, mixed standardisation and limited productivity effects.[15] A later systematic review likewise identified implementation benefits alongside accuracy variation, manual edits and evaluation gaps.[18] These reviews support disciplined evidence collection rather than a category-wide conclusion.

A pragmatic randomised trial across 238 outpatient physicians tested two ambient-scribe applications against usual care and reported application-specific outcomes.[16] Another emergency-department comparison found similar or lower note quality and greater clinician contribution for AI-generated notes versus human scribes in that setting.[17] These studies show why comparator, setting, user and metric matter.

The company should maintain post-deployment evidence. Version changes, new specialties, languages, EHRs and workflow actions can alter performance. Monitoring should include adoption, edits, errors, incidents, subgroups, support and customer outcomes. Commercial claims should follow the approved evidence set.

16. Examine intellectual property, data advantage and defensibility

Defensibility can arise from workflow design, distribution, integrations, evaluation systems, proprietary data rights, clinical trust, regulatory capability and operating evidence. A model alone may be replicable or sourced from an upstream provider. The investor should identify which assets remain valuable if a foundation model improves, prices fall or customers gain similar features from an incumbent platform.

The IP review should cover employee and contractor assignment, background IP, open-source software, model and dataset licences, customer feedback, patents, trademarks, domain names, code provenance and restrictions in commercial partnerships. Generated code and content should be assessed under applicable law and contract. Material components should have a clear owner and permitted commercial use.

A data advantage requires lawful, contractually permitted and technically usable data. Large volume is insufficient when rights prohibit training, labels are weak, populations are narrow or versions cannot be linked to outcomes. The company should explain how data improve a defined capability and how the improvement is measured. Customer-specific data may create switching cost without supporting a transferable model.

Evaluation infrastructure can be a durable asset. Curated test sets, error taxonomies, clinical adjudication, subgroup analysis, production monitoring and release gates can accelerate safe improvement. Their value depends on representativeness, rights, freshness and reproducibility.

The investor should run a substitution test. If the upstream speech or language model changed, what remains? If an EHR vendor launched a competing feature, what workflow, evidence, integration or distribution advantage would protect renewal? Defensibility should be described as a tested economic mechanism.

17. Assess organisation, governance and deployment capability

The organisation should match the risk and operating model of the product. Relevant capabilities can include machine learning, speech engineering, clinical informatics, product management, security, privacy, regulatory affairs, clinical safety, integrations, implementation, customer success and revenue-cycle expertise. The required mix depends on intended use and customer segment.

Key-person analysis should identify who owns model architecture, clinical judgement, regulated quality systems, major integrations, customer relationships and incident response. Knowledge should be evidenced in documentation, review processes and succession rather than concentrated in a founder or a few engineers. Compensation, retention and hiring plans should align with the expansion case.

Governance should connect product release, clinical safety, privacy, security and commercial claims. A release committee can require completed evaluation, hazard review, security testing, documentation, support readiness and customer communication. Exceptions should have a named approver, rationale, expiry and monitoring plan.

Sales incentives require review. Teams may be rewarded for contracted value before security approval, go-live or sustained use. Product claims can outrun evidence when competitive pressure is high. Compensation and approval gates should support qualified revenue and safe deployment.

Board reporting should combine product performance, incidents, rights, regulatory status, customer deployment, retention, margin and roadmap risk. A growth board needs early-warning indicators: stalled

implementations, low active use, rising edits, subgroup regressions, support concentration, inference-cost changes and renewal exposure. Governance becomes an asset when it produces timely decisions and credible evidence for customers.

18. Reconstruct clinician retention from licensed seat to durable use

Clinician retention is the core operating proof. Enterprise contracts can keep nominal seats in reported annual recurring revenue while clinicians stop using the product, use it intermittently or revert to manual documentation. The buyer should obtain clinician-level, privacy-preserving event data and reconcile licensed, enabled, activated, monthly active, productively active and retained users by customer, specialty, site, product version and cohort.

Productive use needs a stable definition. A clinician who launches the application once is not equivalent to one who completes a representative share of eligible encounters, reviews drafts, signs notes on time and remains within the approved workflow. The measure should capture eligible encounters, completed ambient sessions, usable drafts, edit burden, time to signature, exception rate and sustained use. Customers may define eligibility differently, so the buyer should retain both source definitions and a normalised analytical view.

Retention analysis should separate logo, revenue, seat and clinician behaviour. A customer can renew while reducing licensed seats. Revenue can expand through price while active use falls. A multi-year contract can delay the economic effect of dissatisfaction. Conversely, active clinician use can rise before contracted value changes. The acquisition model should connect these evidence states without treating one as a substitute for another.

Figure 6. Clinician adoption and retention funnel

Hypothetical operating funnel; definitions and values are created solely to demonstrate the method.

Table 6. Clinician-retention evidence

Measure	Required definition	Diligence use	Failure signal
activation	first completed eligible encounter	onboarding conversion	licence provisioning counted as use
monthly active	minimum completed workflow in month	sustained breadth	sporadic launch or test activity

Measure	Required definition	Diligence use	Failure signal
productive use	eligible encounter share plus acceptable review	workflow depth	usage without useful output
edit burden	clinically and operationally meaningful change	note quality and total work	hidden correction effort
time to signature	encounter end to attestation	downstream readiness	delayed closure or backlog
twelve-month retention	productive use after defined anniversary	durability	contract masks clinician attrition
expansion	added retained clinicians or modules	growth quality	price or shelfware drives expansion

Proposed cohort controls; thresholds should be set for the product, specialty and customer context.

19. Link note quality to liability and revenue integrity

Liability follows the complete chain from capture to downstream use. A missing fact can affect care. An invented statement can become part of the clinical record. Incorrect speaker attribution can alter meaning. A coding suggestion can contribute to an unsupported claim. A privacy or security failure can expose protected information. A delayed or unavailable note can disrupt handover, orders, referrals and billing. The buyer should map each hazard, controller, detection mechanism, escalation route, contractual allocation, insurance response and residual economic exposure.

Human review remains an operating control. The treating clinician's attestation does not remove the vendor's obligations for product design, representations, security, incident response or agreed service. The transaction team should avoid treating clinician sign-off as a universal transfer of liability. Allocation depends on facts, contract and law. The economic model should carry remediation, customer credits, defence cost, insurance retention, regulatory response and potential revenue loss where supported.

Revenue integrity requires traceability from encounter to claim. The diligence sample should compare source evidence, generated note, clinician edits, signed record, suggested codes, submitted claim, payer response and final cash. Any lift in coding or collections should be separated from volume, payer mix, fee-schedule change, staffing, backlog release and other interventions. Management claims need reproducible data and qualified clinical and coding review.

Table 7. Liability and control allocation

Exposure	Controlling evidence	Primary control	Transaction response
omitted or invented content	source audio, draft, edits and signed note	representative evaluation and clinician review	remediation covenant and claim-specific protection
privacy and recording	notice, consent, purpose and retention records	configured capture and deletion controls	rights condition and indemnity
coding and claims	record, suggestion, reviewer and remittance	qualified coding governance	price adjustment, escrow or covenant
model or vendor change	version, release and dependency register	change control and monitoring	operating covenant and notification right
security and availability	architecture, tests, incidents and recovery	security programme and resilience tests	remediation condition and insurance review
integration failure	interface logs, backlog and reconciliation	staged migration and rollback	holdback linked to verified go-live

Illustrative diligence map; legal conclusions require qualified jurisdiction-specific advice.

20. Rebuild cohort economics and recurring revenue

Reported annual recurring revenue should be reconciled to contracted value, go-live, active clinician use, invoices, collections, credits, renewals and contribution. Human scribe labour, implementation, integration maintenance, clinical quality assurance, coding review, model inference, customer support, security, insurance and customer-specific configuration should be assigned to the cohorts that consume them. Capitalised development and shared infrastructure should remain visible.

The hypothetical case begins with USD 54 million of contracted annual recurring revenue. USD 48 million has gone live, USD 45 million is invoiced on a run-rate basis and USD 43 million has been collected or has a normal collection history. The illustrative cohort incurs USD 4 million of compute and model-vendor cost, USD 7 million of human scribe and quality review, USD 5 million of integration and support, USD 3 million of clinical, privacy, security and insurance cost, and USD 2 million of customer-specific implementation amortisation. The resulting USD 22 million contribution is a modelling output, not a benchmark.

The buyer should test revenue classification and margin by product. A bundled contract can contain software, services and implementation. Software revenue may have high incremental margin after a stable deployment, while human review and customer-specific operations scale differently. Removing people before quality evidence supports the change can create liability and churn. Retaining every legacy process can prevent the combination from achieving its thesis. The ownership plan should therefore link workforce change to measured note quality, exception handling and customer acceptance.

Table 8. Hypothetical recurring-revenue and contribution bridge

Evidence state or cost	Amount	Interpretation
contracted annual recurring revenue	54	signed commercial perimeter
live annual recurring revenue	48	customers through defined go-live
collected or normal-collection recurring revenue	43	revenue with cash evidence
compute and model vendors	(4)	variable technology delivery
human scribe and quality review	(7)	people-intensive delivery and control
integration and support	(5)	live-customer operating burden
clinical, privacy, security and insurance	(3)	control infrastructure
implementation amortisation	(2)	customer-specific activation cost
illustrative contribution	22	before shared corporate cost, tax and financing

All amounts are hypothetical USD millions created solely to demonstrate the method.

21. Translate evidence into valuation and consideration

Valuation should begin with stand-alone cash flow from retained cohorts. A market multiple can serve as a reasonableness check after normalising software and services mix, growth, gross margin, clinician retention, enterprise renewal, concentration, implementation burden, liability and cash consumption. A high contracted growth rate should not override weak productive use or negative mature cohort contribution.

Contingent product modules should be valued through probability-weighted cash flow or funded milestones. Data and model capability should receive value only when the target has lawful, durable rights and the capability creates measurable incremental economics. Replacement cost is evidence of resources

consumed, not proof of economic value. Buyer synergy should be modelled by initiative, owner, cost, timing, tax and execution risk.

Consideration can follow evidence. Closing cash can reflect verified stand-alone value. Deferred consideration can depend on defined customer renewal, retained productive clinicians, compliant go-live, contribution or another auditable outcome. Clinical or regulatory outcomes require careful drafting because the seller may not control them after closing. Holdbacks, escrow, warranties, indemnities and covenants can address identified exposures subject to enforceability, insurance and governing law.

22. Design the roll-up and integration programme

The roll-up thesis should state which capabilities will be standardised, federated or retired. Data purpose and customer obligations can constrain consolidation. EHR integrations may use different architectures. Clinical models can perform differently across specialties and languages. Human scribe operations may contain local employment, credential, scheduling and quality systems. The programme should preserve continuity while establishing a common evidence and control architecture.

During days 0-30, the buyer should freeze product and model versions, secure key personnel, preserve logs and customer obligations, confirm critical subprocessors, reconcile incidents and establish a group clinical, privacy, security and revenue-integrity forum. During days 31-90, it should normalise cohort definitions, run representative note-quality tests, map rights and retention, reconcile revenue and cost, and select limited migration pilots. During days 91-180, it should execute approved pilots, compare results with acceptance criteria, remediate evidence gaps and scale only the workflows that pass.

The board dashboard should combine clinician activation and retention, eligible encounter penetration, note edit burden, time to signature, exceptions, incidents, customer renewal, live recurring revenue, contribution, implementation backlog, model and vendor changes, and realised integration value. Each measure needs a definition, source, owner, threshold and action.

23. Decision and conclusion

A medical-scribe roll-up becomes durable clinical infrastructure when clinicians and enterprises continue using it because the controlled workflow improves documentation and downstream operations. Contracted seats, demonstrations and category growth cannot establish that result by themselves. The evidence must connect authorised data, representative performance, clinician review, integration, productive use, renewal, contribution and cash.

The acquisition decision should preserve the distinctions that matter. Software and services have different economics. Clinician retention and enterprise renewal reveal different behaviours. Note quality and coding outcomes require separate controls. Contractual allocation and practical control determine different parts of liability. Stand-alone value, contingent products and buyer synergy should occupy separate valuation locations.

The framework converts those distinctions into action. The buyer maps the complete workflow, builds a purpose-specific rights ledger, reproduces evaluation, reconstructs clinician cohorts, traces revenue to cash, assigns full delivery cost, maps liability, prices the evidence and embeds unresolved risk into consideration and the ownership plan. A 180-day programme then moves from protection and measurement to controlled pilots and evidence-gated scale.

The final investment question is whether the combined business can repeatedly produce an accurate, useful and timely clinical record while keeping an accountable clinician in control, meeting customer and legal obligations, and earning durable contribution. Evidence that survives this chain supports value. A break identifies remediation, structure, price adjustment or a reason to stop.

Frequently asked questions

What proves that an ambient scribe has become durable clinical infrastructure?

Sustained productive clinician use, acceptable note quality and edit burden, timely attestation, enterprise renewal, controlled incidents and positive mature-cohort contribution provide connected evidence.

Why are contracted seats insufficient for valuation?

Contracts can include unused seats, delayed go-live, discounts or clinicians who stop using the workflow. Valuation should reconcile contracts to active retained use, renewal, contribution and cash.

How should a buyer evaluate note quality?

Use representative specialties, settings, languages and edge cases; compare source evidence, generated draft, clinician edits and signed note; measure clinically material omissions, additions, attribution, structure and downstream effects.

Who carries liability when the clinician signs the note?

Allocation depends on the facts, contract and applicable law. Clinician attestation is one control and does not remove vendor obligations for product design, claims, privacy, security, incidents or agreed service.

How should coding or collection improvement be underwritten?

Trace encounter evidence through note, coding review, claim, payer adjudication and collected cash; control for volume, payer mix, tariff changes, staffing and backlog release.

What costs should be included in cohort contribution?

Include compute, model vendors, human review, integration, support, clinical quality, coding governance, privacy, security, insurance and customer-specific implementation, alongside applicable shared costs.

How can transaction structure protect the buyer?

Use evidence-gated price, deferred consideration, conditions, representations, warranties, indemnities, escrow, holdbacks and operating covenants matched to specific controllable risks.

What belongs in the first 180 days?

Protect continuity, freeze versions, preserve evidence, retain key people, reconcile rights and economics, normalise cohort measures, test representative quality, pilot migrations and scale only after predefined acceptance gates.

References

[1] World Health Organization, Ethics and governance of artificial intelligence for health, 2021, https://www.who.int/publications/i/item/9789240029200

[2] World Health Organization, Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models, 2025, https://www.who.int/publications/i/item/9789240084759

[3] US Food and Drug Administration, Clinical Decision Support Software, Final Guidance, January 2026, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software

[4] US Food and Drug Administration, Guidances with Digital Health Content, https://www.fda.gov/medical-devices/digital-health-center-excellence/guidances-digital-health-content

[5] US Food and Drug Administration, Clinical Decision Support Software FAQs, https://www.fda.gov/medical-devices/software-medical-device-samd/clinical-decision-support-software-frequently-asked-questions-faqs

[6] Office of the National Coordinator for Health Information Technology, HTI-1 Final Rule, https://healthit.gov/regulations/hti-rules/hti-1-final-rule/

[7] Office of the National Coordinator for Health Information Technology, Transparent and Trustworthy AI in Health Care, https://healthit.gov/blog/ai-ml/transparent-and-trustworthy-ai-in-health-care/

[8] Centers for Medicare & Medicaid Services, Evaluation and Management Services, https://www.cms.gov/training-education/medicare-learning-networkr-mln/compliance/medicare-provider-compliance-tips/evaluation-management-services

[9] Centers for Medicare & Medicaid Services, Evaluation and Management Services Guide, 2025, https://www.cms.gov/outreach-and-education/medicare-learning-network-mln/mlnproducts/downloads/eval-mgmt-serv-guide-icn006764.pdf
[10] Centers for Medicare & Medicaid Services, Medical Record Maintenance and Access Requirements, https://www.cms.gov/files/document/mln4840534-medical-record-maintenance-and-access-requirements.pdf
[11] US Department of Health and Human Services, HIPAA Cloud Computing Guidance, https://www.hhs.gov/hipaa/for-professionals/special-topics/health-information-technology/cloud-computing/index.html
[12] US Department of Health and Human Services, Business Associates, https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/index.html
[13] US Department of Health and Human Services, Sample Business Associate Agreement Provisions, https://www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions/index.html
[14] US Department of Health and Human Services, Summary of the HIPAA Security Rule, https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html
[15] US Federal Trade Commission, Complying with the Health Breach Notification Rule, https://www.ftc.gov/business-guidance/resources/complying-ftcs-health-breach-notification-rule-0
[16] Office of the National Coordinator for Health Information Technology, Information Blocking, https://healthit.gov/information-blocking
[17] European Union, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence, https://eur-lex.europa.eu/eli/reg/2024/1689/oj
[18] European Union, Regulation (EU) 2025/327 on the European Health Data Space, https://eur-lex.europa.eu/eli/reg/2025/327/oj/
[19] UK Medicines and Healthcare products Regulatory Agency, Software and artificial intelligence as a medical device, https://www.gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device
[20] NHS England, DCB0129 and DCB0160 clinical risk management standards, https://digital.nhs.uk/data-and-information/information-standards/governance/latest-activity/standards-and-collections/dcb0129-clinical-risk-management-its-application-in-the-manufacture-of-health-it-systems
[21] Microsoft, Microsoft completes acquisition of Nuance, 2022, https://news.microsoft.com/source/2022/03/04/microsoft-completes-acquisition-of-nuance-ushering-in-new-era-of-outcomes-based-ai/
[22] Microsoft, Acquisition of Nuance announcement, 2021, https://news.microsoft.com/source/2021/04/12/microsoft-accelerates-industry-cloud-strategy-for-healthcare-with-the-acquisition-of-nuance/
[23] Microsoft, Annual Report 2022, https://www.microsoft.com/investor/reports/ar22/
[24] IKS Health, IKS Health purchases AQuity Solutions, 2023, https://ikshealth.com/insights/press-releases/iks-health-purchases-aquity-solutions/
[25] HealthChannels, ScribeAmerica and Elite Medical Scribes merger announcement, 2019, https://www.healthchannels.com/wp-content/uploads/2019/07/Press_Release-Elite_Merger.pdf
[26] AQuity Solutions, Quality Overview 2023, https://ikshealth.com/wp-content/uploads/2024/07/AQuity-Quality-Overview-2023.pdf
[27] National Institute of Standards and Technology, AI Risk Management Framework, https://www.nist.gov/itl/ai-risk-management-framework
[28] National Institute of Standards and Technology, Cybersecurity Framework 2.0, https://www.nist.gov/cyberframework
[29] International Medical Device Regulators Forum, Software as a Medical Device: Clinical Evaluation, 2017, https://www.imdrf.org/documents/software-medical-device-samd-clinical-evaluation
[30] HL7 International, FHIR Overview, https://www.hl7.org/fhir/overview.html
[31] Kanaparthy NS et al., Real-World Evidence Synthesis of Digital Scribes Using Ambient Listening and Generative AI, JMIR AI, 2025, https://pubmed.ncbi.nlm.nih.gov/41071988/
[32] Shah SJ et al., Ambient AI Scribes in Clinical Practice: A Randomized Trial, 2025, https://pubmed.ncbi.nlm.nih.gov/41497288/
[33] Morey J et al., Ambient Artificial Intelligence Versus Human Scribes in the Emergency Department, 2025, https://pubmed.ncbi.nlm.nih.gov/41251650/
[34] Clinical Implementation of Artificial Intelligence Scribes in Health Care: A Systematic Review, 2025, https://pubmed.ncbi.nlm.nih.gov/40306686/
[35] ACI-BENCH: a Novel Ambient Clinical Intelligence Dataset for Benchmarking Automatic Visit Note Generation, https://arxiv.org/abs/2306.02022
[36] IFRS Foundation, IFRS 3 Business Combinations, https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/
[37] IFRS Foundation, IAS 38 Intangible Assets, https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/
[38] IFRS Foundation, IFRS 13 Fair Value Measurement, https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/
[39] International Valuation Standards Council, International Valuation Standards, https://www.ivsc.org/standards/
[40] US Department of Justice and Federal Trade Commission, Merger Guidelines, 2023, https://www.justice.gov/atr/2023-merger-guidelines
[41] UK Competition and Markets Authority, Merger Assessment Guidelines, https://www.gov.uk/government/publications/merger-assessment-guidelines
[42] European Commission, Guidelines on the assessment of horizontal mergers, https://competition-policy.ec.europa.eu/mergers/legislation_en
[43] UK Information Commissioner's Office, Guidance on AI and data protection, https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/
[44] European Data Protection Board, Guidelines and recommendations, https://www.edpb.europa.eu/our-work-tools/general-guidance/guidelines-recommendations-best-practices_en
[45] US Department of Health and Human Services, Guidance on HIPAA and online tracking technologies, https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/hipaa-online-tracking/index.html
[46] US Department of Health and Human Services, Is a software vendor a business associate?, https://www.hhs.gov/hipaa/for-professionals/faq/is-software-vendor-business-associate/index.html
[47] Centers for Medicare & Medicaid Services, Administrative Simplification: Claims Attachments Final Rule, https://www.cms.gov/newsroom/fact-sheets/administrative-simplification-adoption-standards-health-care-claims-attachments-transactions
[48] US Food and Drug Administration, Reports on Non-Device Software Functions, https://www.fda.gov/about-fda/cdrh-reports/reports-non-device-software-functions
[49] American Medical Association, Digital health research and resources, https://www.ama-assn.org/practice-management/digital-health
[50] Agency for Healthcare Research and Quality, Patient Safety Network, https://psnet.ahrq.gov/

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.