Local Health, Global IP: Structuring GCC Bio-AI Platforms

A transaction framework for converting governed regional health data and scientific capability into defensible global intellectual property and commercial pathways

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Gulf health systems are assembling longitudinal clinical records, genomic programmes, research infrastructure and artificial-intelligence capability at a scale that can support precision medicine and biopharmaceutical discovery. The investment question is how to organise those inputs so that locally governed data creates durable regional health benefit and globally defensible commercial rights. A platform can fail even when its science is promising if the parties have not separated data access, model development, patentable inventions, clinical evidence, regulatory submissions, product rights and overseas licensing.

This paper develops a Local Health, Global IP Transaction Framework for governments, health systems, universities, biotechnology companies, pharmaceutical partners and investors. It maps the value chain from lawful data access through curation, model development, target discovery, experimental validation, clinical translation and commercial licensing. It assigns ownership, control, access, publication, improvement, regulatory and exit rights to the entity best able to perform and finance each function. The framework treats health and genomic data as governed inputs rather than automatically transferable property. It distinguishes the data from curated datasets, software, model weights, validation evidence, patents, know-how, regulatory dossiers and product rights.

A wholly hypothetical case examines a GCC bio-AI platform requiring USD 120 million across secure infrastructure, data engineering, discovery programmes and translational validation. The case assumes three health-system contributors, two technology partners and one strategic pharmaceutical collaborator. Funding is released against data-governance readiness, reproducible model performance, experimentally validated targets, accepted development candidates and signed licensing economics. The proposed structure combines USD 45 million of sponsor equity, USD 25 million of strategic capital, USD 20 million of grant and programme funding, USD 15 million of milestone-linked venture debt and USD 15 million of partner-funded research. Every case figure is a hypothetical management assumption; it is not observed company data, a market forecast, a valuation opinion or an offer of financing.

Current primary and authoritative evidence supports the need for explicit governance. UAE health-ICT and personal-data rules protect confidentiality, integrity, authorised access and cross-border processing. Abu Dhabi and Dubai standards address health-information exchange, interoperability and genomic-data governance. Saudi Arabia's biotechnology strategy links genomics, data access and precision medicine. WHO, FDA and EMA materials emphasise intended use, data governance, model credibility, lifecycle management and human accountability. WIPO guidance shows that healthcare AI is protected through a layered bundle of patents, copyright, trade secrets, contractual rights and licences. The paper concludes that regional data access should remain purpose-bound and

auditable while global value is commercialised through clearly defined IP, evidence and market rights.

JEL Classification: G24, G31, G34, I11, L65, O31, O32, O34

Keywords: bio-AI, health data, genomics, intellectual property, GCC, platform transactions, technology transfer, licensing, joint ventures, precision medicine

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, an offer, or a solicitation. It contains no client information.

1. DEFINE THE TRANSACTION DECISION

The transaction decision is which legal and economic structure can combine health data, scientific capability, AI technology and development capital without weakening patient protection or commercial incentives. The platform may be a government-backed infrastructure company, a health-system consortium, a university spin-out, a corporate joint venture, a contractual alliance or a series of programme-specific vehicles.

The choice should follow the intended output. A population-health analytics platform has different rights and liabilities from a drug-discovery company, a diagnostic developer or a licensing vehicle. The parties should define the health outcome, scientific question, model context of use, validation pathway, product class, target geography and expected commercial transaction before allocating ownership.

The structure should also answer who can stop a programme. Data stewards, scientific leaders, regulators, funders and product sponsors can each hold distinct approval rights. A financing case becomes credible when those approvals, evidence gates and economic consequences are visible before capital is committed.

Table 1. Bio-AI platform layers and transaction treatment

Layer	Primary contribution	Financeable evidence	Typical transaction treatment
Health and genomic data	Lawful access, provenance and stewardship	Consent or legal basis, quality, permitted use and audit trail	Purpose-bound access agreement
Curated dataset	Cleaning, linkage, annotation and standardisation	Reproducible data pipeline and documented transformations	Licensed dataset or platform asset
Software and models	Code, architecture, weights and workflows	Version control, performance and portability	Background licence plus improvement rights
Scientific discovery	Targets, biomarkers and hypotheses	Reproduced experimental evidence	Foreground IP and programme rights
Clinical and regulatory evidence	Protocols, trials, dossiers and surveillance	Accepted study outputs and regulator-ready traceability	Sponsor-controlled evidence package
Product and market rights	Development, manufacture and commercialisation	Candidate acceptance, approvals and channel capability	Territorial or field licence, option or acquisition

Each layer requires a distinct rights, evidence and financing analysis.

2. MAP THE CONTRIBUTION CHAIN

The contribution chain begins with lawful access to clinical, genomic or research data. It continues through curation, feature engineering, model training, hypothesis generation, laboratory validation, candidate selection, clinical evidence and commercial development. Each stage changes the risk, cost and legal character of the output.

Raw health facts may not themselves be exclusive intellectual property. A curated compilation, software pipeline, trained model, new biomarker, experimental method, patentable invention or confidential development package can support a different bundle of rights. The transaction documents should identify each transformation and the contributor responsible for it.

This map also prevents value duplication. A party should not claim the full value of a therapeutic programme because it supplied one input. Contribution value should reflect scarcity, lawful substitutability, cost, performance, risk assumed and the incremental effect on the programme's probability and economics.

3. SEPARATE STEWARDSHIP FROM OWNERSHIP

Health data requires stewardship that can survive a change in shareholders, technology vendors or commercial partners. Stewardship includes the lawful basis for use, access control, security, quality, retention, audit, patient rights and regulator engagement. These duties should remain with an accountable entity that has authority over the data environment.

Commercial ownership applies more naturally to created assets such as software, model components, patents, laboratory methods, regulatory submissions and contractual product rights. The structure should avoid language that purports to transfer a health system's underlying patient data when the intended arrangement is controlled access for a defined purpose.

The platform can create a clean economic interface by granting compute-to-data access, controlled research workspaces or approved extracts. Commercial partners receive defined use rights and validated outputs. The steward retains policy control, monitors use and can suspend access when conditions are breached.

4. DESIGN THE GOVERNED DATA ZONE

The governed data zone is the technical and contractual boundary in which approved processing occurs. It should identify systems, locations, users, purposes, data classes, permitted tools, export rules, logging, model-output review and incident response. Data residency and cross-border processing should be analysed for each jurisdiction and dataset.

The UAE's official health-ICT summary emphasises confidentiality, validity, integrity, availability and authorised access. Its personal-data framework addresses consent, controller obligations and cross-border transfer. Abu Dhabi's exchange standards require common coding, minimum datasets and interoperability. Dubai's genomic-governance standards sit alongside policies for sharing, quality, classification, consent, confidentiality, security and AI.

A transaction should convert these obligations into operating controls. The data-access agreement, security schedule, model-development plan and audit protocol should describe the same permitted activity. Contradictions between documents create risk and can make evidence unusable.

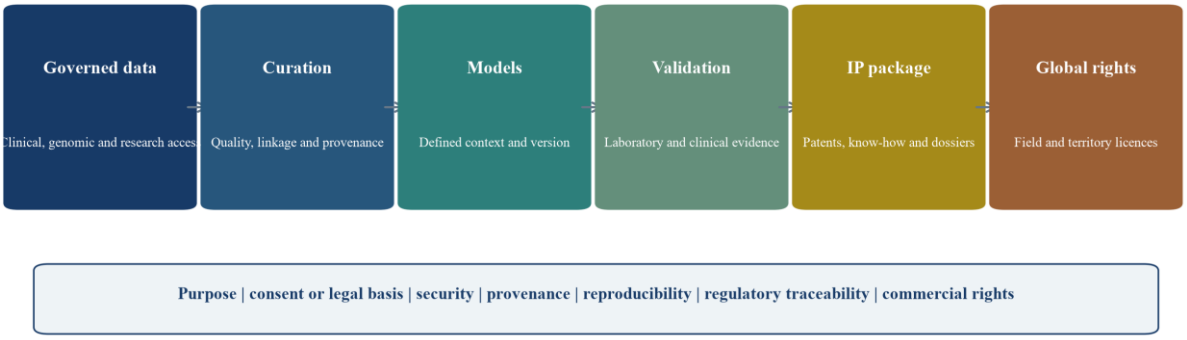

Figure 1. Governed regional data to global commercial rights

Sensitive data remains inside the governed zone while approved evidence and defined IP rights move through validation and licensing gates.

5. CHOOSE THE PLATFORM PERIMETER

The platform perimeter determines which assets and obligations sit inside the venture. A broad platform may include data infrastructure, laboratories, discovery programmes and product development. A narrow platform may provide secure analysis and licence outputs to separate programme companies.

Broad structures capture more upside and require more capital. They also combine operational, scientific, regulatory and product risk. Narrow structures can ring-fence liability and let different investors fund different stages. They may create transfer-pricing, priority and coordination issues.

The preferred perimeter should match the parties' strategic commitments. If a health system seeks sovereign control of data infrastructure while an investor seeks exposure to global licensing, the platform can separate the governed data utility from an IP commercialisation company. Rights between them must be durable, performance-linked and auditable.

6. SELECT THE ENTITY ARCHITECTURE

A single company is simple when contributions and risk are aligned. A holding company with operating subsidiaries can separate regulated data operations, research services and programme IP. A contractual alliance may be sufficient for early discovery. A special-purpose company can ring-fence one target, indication or product.

Entity architecture affects tax, transfer pricing, licensing, control, insolvency, funding and exit. It also determines where employees, inventors, laboratory assets and regulatory sponsorship sit. These matters should be mapped before valuation because enterprise value can migrate with contracts and people.

Local-law and cross-border advice is required for the selected jurisdictions. The commercial framework should first identify the operational truth: which entity controls the data environment, employs the scientific team, owns foreground IP, contracts with customers, sponsors regulated studies and receives licensing income.

7. BUILD THE RIGHTS REGISTER

The rights register should list every material asset, owner, contributor, restriction, licence, field, territory, duration, sublicensing right, improvement rule, publication right and exit treatment. It should include patents, applications, software, models, weights, training pipelines, datasets, annotations, laboratory protocols, biomarkers, samples, dossiers, trademarks and confidential know-how.

WIPO describes healthcare AI protection as a layered bundle rather than a single right. Patents can cover qualifying inventions. Copyright may protect code and original compilations. Trade secrets can protect model parameters, curation methods and know-how. Contracts govern access, confidentiality, permitted use and commercial allocation.

The register should connect each right to evidence of title. Employment terms, consultant assignments, university policies, open-source licences and third-party data conditions can change ownership. A platform with incomplete chain of title can produce strong science and weak transaction value.

8. DEFINE BACKGROUND IP

Background IP exists before or outside the collaboration. It can include a technology partner's models, a university's assays, a hospital's software, a pharmaceutical company's targets and a research institute's protocols. Each item should be identified or described with enough precision to avoid later disputes.

The platform needs a licence broad enough to perform the agreed programme. That licence should address modification, hosting, use by contractors, regulatory disclosure, validation, continuity and exit. A narrow research-only licence may be inadequate when the business plan assumes commercialisation.

The contributor should preserve rights outside the agreed field. The agreement can prohibit reverse engineering or competitive model training while still granting the platform access to documentation, escrow or transition support. Financeability improves when essential background rights survive a change of control or contributor distress.

9. ALLOCATE FOREGROUND IP

Foreground IP arises from programme activity. Ownership can follow inventorship, contribution, field, entity or a negotiated allocation. Joint ownership can appear balanced and can create enforcement, licensing and consent problems across jurisdictions.

A platform company can own programme foreground while contributors retain background. Another structure assigns inventions to the party best placed to prosecute and commercialise them, subject to licences and revenue sharing. The correct rule depends on the product pathway and the parties' capabilities.

The agreement should cover invention disclosure, prosecution control, costs, abandonment, enforcement, defence and settlements. Scientific notebooks, model versions and experimental records should support inventorship and priority. Commercial economics should follow enforceable rights and validated evidence rather than an undefined promise to share future IP.

10. GOVERN MODEL IMPROVEMENTS

AI models can improve through platform data, user feedback, fine-tuning, new labels and experimental outcomes. The parties should distinguish general improvements from programme-specific models and clinically significant adaptations.

A technology provider may retain improvements to its general platform. The venture may require a perpetual licence to improvements generated through its data and funding. Programme-specific weights, prompts, features or decision rules can remain exclusive to the venture. The agreement should address whether outputs can train models for other customers.

Model lineage should record the base model, data version, code, parameters, evaluation set, changes and approvals. This record supports scientific reproducibility and regulatory credibility. It also enables a buyer to identify which performance survives if a vendor relationship ends.

11. CREATE A DATA-ACCESS LICENCE

The data-access licence should define purpose, users, systems, duration, data classes, linkage, re-identification restrictions, output review, publication, retention, deletion, audits and consequences. It should align with ethics approval, patient consent or another lawful basis.

Access can be non-exclusive while programme outputs or defined uses are exclusive. Exclusivity over broad disease areas can impede public-health goals and future research. A narrower field, target, indication or use case can provide commercial confidence while preserving other beneficial activity.

The licence should state that approval to access data does not automatically grant ownership of discoveries, models or products. Those allocations belong in the IP and commercial schedules. This separation makes governance clearer and improves valuation discipline.

12. SET CONSENT AND PURPOSE CONTROLS

Consent and purpose limitations affect which research, training, validation and commercial uses are permitted. A platform should not assume that de-identification resolves every legal, ethical or contractual issue. Genetic and longitudinal health data can remain sensitive and can create group-level concerns.

The governance team should map each dataset to its lawful basis, consent language, ethics approval and regulator conditions. The map should identify whether commercial development, international collaboration, model training, secondary research and return of results are covered.

When the purpose expands, the platform should follow an approved change process. The change may require new review, additional consent, technical controls or a different data product. A financing milestone should depend on confirmed usable scope rather than the theoretical volume of records.

13. ENGINEER INTEROPERABILITY

Interoperability affects both science and economics. Common identifiers, coding standards, laboratory units, terminology and metadata enable datasets to be linked and models to be validated across institutions. Weak interoperability increases curation cost and can produce hidden bias.

Abu Dhabi's health-information exchange standards specify coding, units, minimum datasets and implementation expectations. A platform should treat conformity as an operating capability. Data completeness, semantic consistency and change management should be measured by source and period.

The integration plan should include mapping logic, exception handling, source-system changes and reconciliation. These artefacts can become valuable know-how. They should be documented without exposing sensitive data and should remain available to the platform if a systems integrator is replaced.

14. MEASURE DATA FITNESS

Record count is a weak proxy for value. Data fitness depends on representativeness, completeness, accuracy, longitudinal depth, outcome availability, missingness, label quality, linkage, provenance and relevance to the scientific question.

The platform should publish an internal data card for each approved use. It should describe the population, time period, sources, transformations, exclusions, known limitations and allowed decisions. Drift should be monitored when clinical practice or source systems change.

Financing should follow usable cohorts and reproducible analyses. A dataset that cannot support external validation or regulatory reliance may still have research value. The valuation should reflect the actual decision it can support.

15. DEFINE THE MODEL CONTEXT OF USE

The context of use states what a model does, for whom, with which data, under what conditions and for which decision. The same algorithm can carry different risk when used for hypothesis generation, patient selection, dosing, diagnosis or regulatory evidence.

FDA's 2025 draft guidance proposes a risk-based credibility framework for AI models used to support regulatory decisions. FDA and EMA's 2026 principles emphasise human-centred design, clear context, data governance, performance assessment and lifecycle management. WHO also stresses intended use, transparency, external validation and monitoring.

The transaction model should connect each context to evidence, liability and approval rights. Commercial value depends on validated use rather than generic model capability.

16. ESTABLISH HUMAN ACCOUNTABILITY

The platform should assign accountable human owners for data access, scientific design, model approval, clinical interpretation, regulatory interaction and commercial release. Automation can recommend an action. Governance determines who may rely on it and who can override it.

High-impact uses require documented review, escalation and incident management. The platform should record model limitations, uncertainty and circumstances in which output must not be used. Users need training appropriate to their role.

Human accountability also affects warranties and indemnities. A technology provider should not be made responsible for clinical decisions it does not control. A healthcare entity should not accept model-development risk outside its capability. The contract should follow actual control.

17. VALIDATE ACROSS POPULATIONS

A model trained on one population may not perform consistently elsewhere. The GCC can support valuable research into regional genomic and clinical patterns. Global commercialisation requires evidence that separates population-specific insight from generalisable performance.

The validation plan should test relevant demographic, clinical, geographic and operational subgroups. It should use independent data and prospectively generated evidence where the context demands it. Data leakage and repeated tuning against a holdout set should be controlled.

A platform can license a region-specific model, a globally validated model or a method that partners adapt locally. Each product has different evidence needs and market rights. Valuation should reflect the supported scope.

18. LINK MODELS TO EXPERIMENTAL EVIDENCE

Bio-AI value increases when computational predictions are tested in laboratories and translated into reproducible biological evidence. The platform should define the handoff from model output to assay, replication, mechanism, candidate selection and development.

Experimental protocols should identify samples, controls, endpoints, acceptance thresholds and independent review. Negative results should remain available because they can improve models and prevent repeated failure. The data and IP schedules should address those results.

The economic model should measure conversion between stages. Prediction volume has limited meaning without validation yield, cycle time, cost and decision impact. Funding can be released when evidence crosses a pre-agreed threshold.

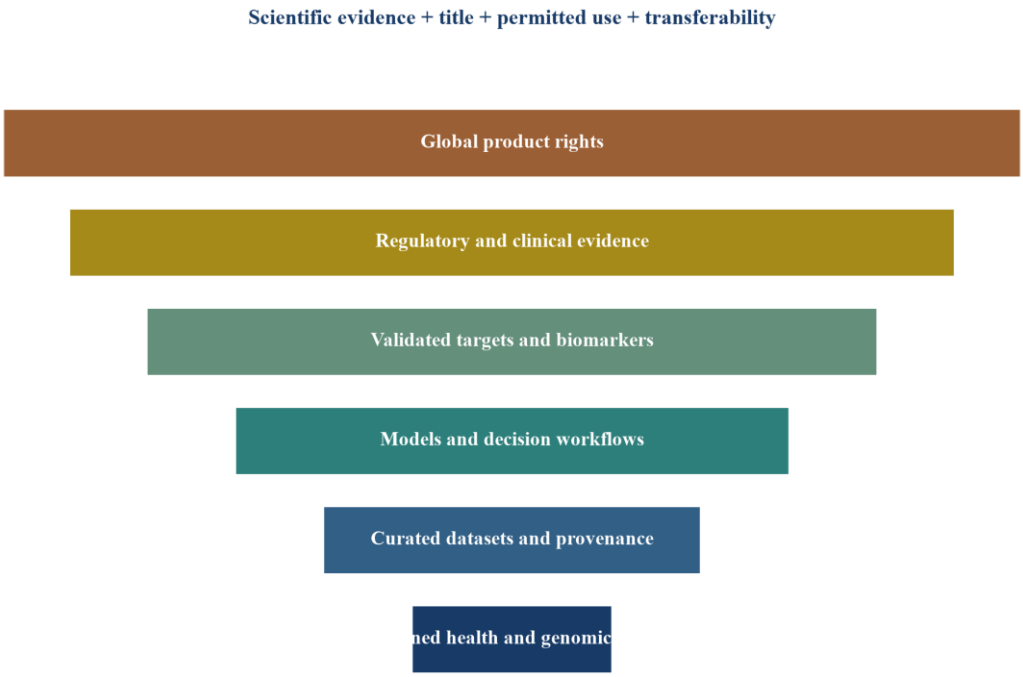

Figure 2. Evidence and rights stack for a GCC bio-AI platform

Commercial value requires both scientific evidence and an enforceable rights chain.

19. DESIGN THE PUBLICATION POLICY

Scientific publication supports credibility, recruitment and public benefit. Premature disclosure can damage patentability or reveal confidential methods. The publication policy should define review periods, permitted delay, attribution, data disclosure, preprints, conference abstracts and regulator submissions.

Academic partners need freedom to publish. Commercial partners need time to protect inventions and remove confidential information. A defined review process can accommodate both. Indefinite vetoes can undermine the research mission and talent proposition.

The policy should also address model and dataset transparency. Full release may be inappropriate for sensitive data or dual-use methods. Reproducibility can be supported through controlled access, protocol disclosure, evaluation benchmarks and independent review.

20. ALLOCATE REGULATORY SPONSORSHIP

The regulatory sponsor controls interactions, submissions, safety reporting and evidence commitments for a product. A research platform may generate evidence without becoming the sponsor. A pharmaceutical or diagnostic partner may take that role after an option or licence.

The transaction should define who prepares records, who can rely on them, who answers regulator questions and who bears remediation cost. Data and model traceability should remain available throughout the product lifecycle.

Regulatory transfer provisions matter on exit. A buyer or licensee needs access to study data, code versions, validation reports, adverse-event information and correspondence. The platform should maintain these assets before a transaction begins.

21. DEFINE COMMERCIAL FIELDS

Commercial rights should be divided by product, target, indication, modality, territory, customer segment or application. Broad field definitions can create dead zones when a licensee does not pursue all opportunities. Narrow definitions can fragment development.

The agreement should include diligence obligations, development plans, minimum spend, milestone dates and reversion. Sublicensing economics should distinguish genuine development partners from intra-group transfers. Rights can expand after performance.

Government or health-system contributors may require access rights, local supply, affordability commitments, research licences or regional capacity. These conditions should be explicit and valued as part of the transaction.

22. STRUCTURE OPTION-TO-LICENSE PATHWAYS

An option allows a strategic partner to fund research before committing to a full licence. The option should define the programme, exclusivity period, information rights, exercise trigger, pre-agreed economics and consequences of non-exercise.

The platform retains leverage when the option is narrow, time-limited and funded. A broad unpaid option can block other partnerships. The partner gains value when diligence access and licence terms are sufficiently clear.

Exercise can occur after a validated target, nominated candidate, accepted biomarker or regulator interaction. The trigger should be objectively evidenced. The parties should avoid a milestone that can be withheld through subjective dissatisfaction.

23. USE PROGRAMME-SPECIFIC VEHICLES

A programme-specific vehicle can hold one discovery programme, licence or product. It can receive focused capital and isolate risk. The platform contributes rights and services; investors fund defined milestones; strategic partners obtain options or licences.

The vehicle documents should address platform access, personnel, data, models, laboratories, IP prosecution and overhead. A thin entity without durable access can be difficult to finance or sell. The services agreement should survive control changes long enough to complete the programme.

Programme vehicles also support portfolio financing. Investors can choose exposure by indication or stage. The parent platform can retain common infrastructure and a share of downstream economics.

24. BUILD THE GOVERNANCE MATRIX

Reserved matters should cover data use, new programmes, budget, capital, licences, publications, regulatory submissions, model deployment, related-party arrangements, cybersecurity, material disputes and exit. Approval thresholds should reflect consequence and expertise.

A scientific committee can review evidence. A data-governance committee can approve use and monitor compliance. The board controls capital and strategy. These bodies should have defined authority and escalation.

Deadlock mechanisms should preserve patient safety and asset value. Interim operating rules, expert determination and narrow buy-sell rights can prevent paralysis. A party should not gain commercial advantage by blocking required compliance.

Table 2. Rights and control matrix

Decision	Data steward	Platform company	Technology partner	Product partner
Data access and permitted purpose	Approve and audit	Implement and evidence	Comply	Comply
General model improvement	Review permitted use	Receive venture rights	Own or co-own as agreed	No automatic right
Programme-specific model	Oversight for data conditions	Own or control	Defined licence or share	Option or field licence
Foreground inventions	Consent where required	Coordinate and own by rule	Assign contributor inventions	Receive exercised rights
Clinical and regulatory strategy	Health-system input	Support evidence	Technical support	Sponsor and control
Global commercial licence	Protect regional obligations	Negotiate and approve	Consent only if background affected	Acquire defined rights

The final allocation should follow actual capability, law and transaction objectives.

25. PRICE CONTRIBUTIONS

Contribution pricing should distinguish cash, data access, infrastructure, people, technology, laboratory work, regulatory capability and market rights. Each contribution has a cost, risk and alternative source.

Data-access value should reflect fitness for the defined use, scarcity, lawful scope, quality and the cost of building an equivalent cohort. It should not be calculated by multiplying record count by an arbitrary price. Model value should reflect incremental performance and portability. Scientific work should be measured by accepted output.

Contributions can earn equity, fees, milestones, royalties or preferred returns. The structure should avoid paying twice for the same input through both an inflated valuation and full service fees.

26. BUILD THE PROGRAMME ECONOMICS

Programme economics should connect spend to evidence and optionality. Costs include data engineering, compute, software, laboratory validation, clinical work, regulatory support, IP, manufacturing and programme management.

Revenue can arise from research funding, access fees, options, licence upfronts, milestones, royalties, services and product economics. Timing and probability differ. Contracted research is closer to cash than a future royalty.

The model should show programme-level cash burn, stage conversion, partner concentration and shared-platform allocation. It should separate recurring platform revenue from contingent pipeline value. This supports financing and valuation.

27. DESIGN MILESTONE FINANCE

Milestone finance releases capital after defined evidence. Early funding can build the governed environment and complete data curation. Later tranches can follow validated model performance, replicated biological findings, accepted candidates and signed commercial rights.

The milestone schedule should define evidence, reviewer, timing, cure and consequence. Scientific uncertainty should not be disguised as a covenant breach. A failed hypothesis can still be a valid result if the programme followed the approved protocol.

Funding conditions should focus on platform readiness and decision-quality evidence. They should avoid rewarding only positive results. Portfolio capital can tolerate programme attrition when the platform learns and reallocates resources.

28. MATCH CAPITAL TO RISK

Sponsor equity should fund governance, core infrastructure and early uncertainty. Strategic capital can fund shared capability and preferred access. Grants can support public-interest research or capacity. Partner research payments can fund programmes. Venture debt may enter after contracted revenue, liquidity and enterprise evidence exist.

Equipment finance can support identifiable laboratory or computing assets with title and recovery value. It should not be underwritten against speculative IP. Working-capital facilities can finance accepted receivables from creditworthy partners.

The capital stack should preserve runway through scientific delays. Debt service should not force premature licensing of high-potential programmes. Liquidity covenants and draw conditions should reflect the platform's development cycle.

29. UNDERWRITE STRATEGIC CONCENTRATION

A platform can depend on one health system, technology provider, laboratory or pharmaceutical partner. Concentration can accelerate execution and can weaken negotiating leverage and continuity.

Diligence should test substitution time, contractual rights, technical portability, data continuity and personnel dependence. The platform should maintain exportable records, modular interfaces and transition support.

Strategic exclusivity should be priced. A partner that receives broad access or blocking rights should provide meaningful funding, development capability and performance commitments. Reversion should follow inactivity.

30. PROTECT CYBERSECURITY AND MODEL INTEGRITY

Health and genomic platforms face confidentiality, integrity, availability and model-manipulation risks. Controls should include identity management, least privilege, encryption, segmentation, logging, monitoring, incident response, vendor review and tested recovery.

Model integrity requires controlled code, data lineage, signed versions, evaluation records and deployment approval. The platform should detect unauthorised changes and data poisoning. Third-party models and libraries should be inventoried.

Cyber obligations belong in data, technology and service agreements. Liability caps, insurance and remediation should reflect control and consequence. A breach can affect patient trust, regulatory permission and enterprise value.

31. PLAN CROSS-BORDER COMMERCIALISATION

Global licensing can transfer patents, know-how, software, validation evidence and product rights without exporting the underlying sensitive dataset. Partners can receive approved outputs, model artefacts or access to a controlled environment.

Cross-border design should consider data protection, health-data rules, export controls, sanctions, tax, competition and local regulatory requirements. The transaction should identify where processing, development, regulatory sponsorship and commercial exploitation occur.

Territorial licences can preserve GCC rights while enabling global development. The platform should ensure that regional evidence can be used where lawful and scientifically appropriate. Reciprocal rights to global learning can improve local health outcomes.

32. PRESERVE REGIONAL VALUE

Regional value can include health-system access, training, research capacity, local validation, manufacturing options, priority supply, affordable access and reinvestment. These objectives should be translated into measurable obligations.

A licence can reserve GCC rights or grant them back to the platform. It can require a regional development plan, local studies or knowledge transfer. Economics can differ by market while preserving global incentives.

Vague localisation promises are difficult to enforce. The agreement should define deliverables, dates, budgets, responsible entities and remedies. Public-interest conditions should remain compatible with the product's development pathway.

33. BUILD VALUATION BY LAYER

Infrastructure and contracted services can be valued through replacement cost, utilisation and cash flow. Software and models require evidence of performance, rights and portability. Programme value requires probability-adjusted cash flow and explicit development cost. Platform optionality should be assessed without duplicating programme value.

Data access should support value only when it is lawful, usable, durable and relevant. The valuation should not treat sensitive data as freely transferable inventory. A buyer may value the relationship, governance capability and reproducible pipeline more than possession of data.

Scenario analysis should vary access continuity, conversion rates, partner terms, development cost, time and probability. The result is a range, not a single precise number.

34. RUN A HYPOTHETICAL FINANCING CASE

Assume a GCC bio-AI platform seeks USD 120 million over four years. Three health systems provide governed access to approved clinical and genomic cohorts. Two technology partners contribute models and engineering. A pharmaceutical partner funds selected programmes and holds narrow options.

Management assumes five discovery programmes, two validated targets, one development candidate and one global licence by year four. These are hypothetical assumptions. The financing case discounts them through evidence gates, programme attrition, delays and partner concentration.

The platform spends USD 28 million on secure data infrastructure and curation, USD 24 million on software and models, USD 32 million on experimental validation, USD 18 million on translational and regulatory work and USD 18 million on people, IP and programme operations.

Table 3. Hypothetical capital stack and permitted use

Capital source	Amount	Primary use	Release evidence
Sponsor equity	45	Core platform, governance and runway	Rights package, team and approved budget
Strategic capital	25	Technology integration and shared programmes	Durable access, interoperability and contribution agreements
Grants and programme funding	20	Public-interest research and validation	Approved protocols and accepted deliverables
Milestone-linked venture debt	15	Expansion after operating evidence	Contracted revenue, liquidity and validated output
Partner-funded research	15	Defined pharmaceutical programmes	Signed scope, budget and acceptance criteria
Total	120	Integrated platform and programmes	Controlled draw schedule

All figures are illustrative assumptions in USD millions.

35. STAGE THE DRAWDOWN

The first draw funds incorporation, governance, security, rights diligence and data-zone commissioning. The second follows approved data access, interoperable pipelines and reproducible cohort construction. The third follows independently reproduced model and laboratory evidence.

Venture debt becomes available after the platform has contracted partner funding, sufficient runway and accepted output. It is not used to finance the first proof of scientific validity. Partner research payments follow programme budgets and deliverables.

Draw conditions should be assessed by qualified reviewers. The platform should retain flexibility to stop or replace a programme when evidence weakens. Capital reallocation should require board and scientific approval.

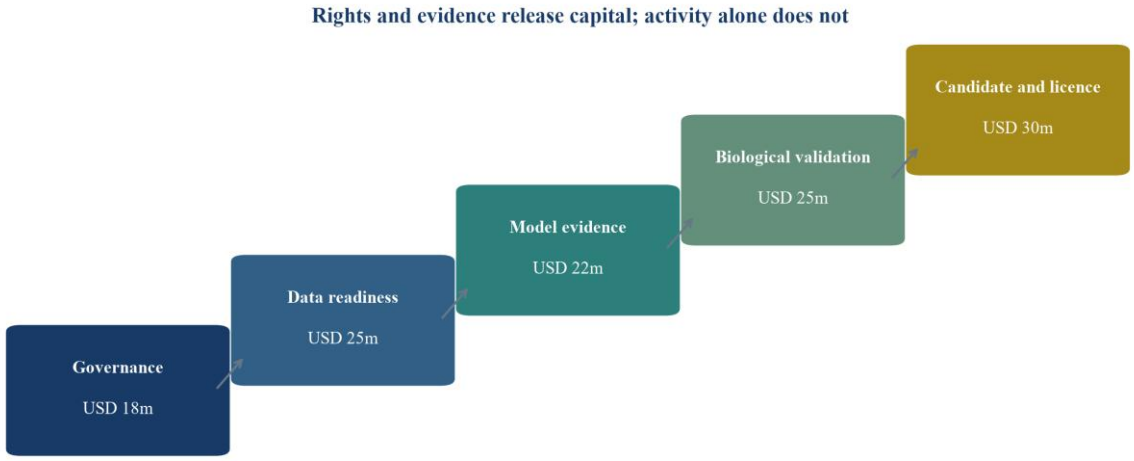

Figure 3. Evidence-linked financing sequence

Funding follows governed access, reproducibility, validation and commercial commitment.

36. TEST THE DOWNSIDE

The downside case assumes data access is narrower than expected, integration takes twelve months longer, one technology provider must be replaced, model performance does not generalise and no global licence is signed during the financing period.

Management should reduce discretionary programmes, preserve the governed platform, protect core IP and extend contracted research. Equity and grant capital absorb early uncertainty. Debt availability is delayed until cash evidence exists.

The board should identify minimum liquidity, termination rights, transition services and asset sale options. A platform can retain strategic value through compliant infrastructure, curated datasets, trained teams, reusable models and validated methods even when a product programme fails.

37. DEFINE COVENANTS

Financial covenants should focus on liquidity, leverage and contracted cash. Operating covenants can address security, insurance, data permissions, IP maintenance, material licences, key-person coverage and regulatory standing.

Scientific covenants should require process and reporting rather than guaranteed outcomes. The platform can report reproducibility, validation status, programme conversion and exceptions. It should not promise discovery success.

Consent rights should apply to broad new data uses, disposal of core IP, material exclusivity, change of control and related-party licences. Ordinary scientific decisions should remain with qualified management under approved budgets.

38. DESIGN EXIT PATHWAYS

Exit pathways include a platform sale, programme acquisition, licence portfolio, strategic merger, public listing or continued cash-generating research business. The structure should support more than one route.

A buyer will test title, access continuity, data permissions, model lineage, regulatory evidence, security, key people and change-of-control clauses. The platform should maintain a transaction-ready data room from inception.

Public or health-system contributors may require consent or preserved regional obligations. These conditions should be disclosed early. Hidden restrictions reduce transaction certainty and value.

39. PREPARE THE DILIGENCE DATA ROOM

The data room should contain entity documents, contribution agreements, data maps, approvals, security evidence, IP registers, assignments, licences, model cards, validation reports, laboratory records, publications, regulatory correspondence, contracts, budgets and cap tables.

Sensitive patient-level data should remain outside the transaction data room. Buyers can review governance artefacts, schemas, quality reports, approved demonstrations and controlled diligence environments. Access should be logged.

Red flags should be linked to remediation owners and dates. Chain-of-title gaps, undocumented model training, expired permissions and inconsistent exclusivity should be addressed before a financing or sale.

Table 4. Transaction-readiness scorecard

Domain	Minimum evidence	Principal risk	Decision
Governance	Approved purposes, access logs and accountable owners	Unlawful or unusable data	Restrict, remediate or stop
IP	Register, assignments, licences and prosecution status	Weak title or blocked commercialisation	Cure before value attribution
Models	Lineage, context, validation and monitoring	Performance or transfer failure	Limit use or require evidence
Science	Protocols, reproducibility and independent review	Non-reproducible claims	Re-run or discount programme
Commercial	Signed scope, options, licences and diligence duties	Concentration or dormant rights	Reprice or add reversion
Finance	Runway, programme cash, obligations and downside plan	Forced financing or premature licence	Restructure capital

Evidence should be current, attributable and consistent across legal, scientific and technical records.

40. IMPLEMENT THE FIRST 100 DAYS

The first 100 days should confirm entity scope, appoint accountable officers, freeze the contribution schedule and validate every critical right. The platform should commission its data zone, approve the first contexts of use and establish model and experiment lineage.

The board should approve the portfolio, evidence gates, capital plan, security baseline, publication process and transaction-readiness register. Partner contracts should be reconciled against the operating model.

Early delivery should include one end-to-end reproducible workflow. This proves that governance, data, models, laboratory activity and decision records connect. It also exposes gaps before the platform scales.

41. TRACK A BOARD DASHBOARD

The board dashboard should combine capital, governance, science and commercial measures. Capital measures include runway, committed funding and programme cash. Governance measures include approved uses, access exceptions, security incidents and audit closure.

Scientific measures include data fitness, model performance by context, reproducibility, validation conversion and programme decisions. Commercial measures include contracted research, option coverage, licence discussions, concentration and stage-adjusted pipeline value.

Definitions should remain stable and changes approved. A composite score can conceal failure in a critical domain. The dashboard should display each gate separately.

42. RECOGNISE LIMITATIONS

The framework is a transaction and financing method. It does not determine whether a scientific hypothesis is correct, whether a particular data use is lawful or whether a model is clinically safe. Those matters require qualified review in the relevant jurisdiction and context.

GCC legal and regulatory requirements differ and continue to develop. Product pathways vary across therapeutics, diagnostics, devices, software and research tools. International commercialisation introduces additional requirements.

The hypothetical case is not empirical evidence. Actual values, costs, probabilities, terms and timelines require transaction-specific diligence. Every agreement should receive professional legal, privacy, regulatory, scientific, technical, tax and accounting advice.

43. CONCLUDE WITH ALIGNED RIGHTS AND EVIDENCE

The strongest GCC bio-AI platform keeps local health data under accountable governance while creating clearly owned and transferable outputs. Its structure separates access from ownership, background from foreground, general model improvements from programme IP, and scientific activity from decision-grade evidence.

Global value follows a complete rights chain, reproducible science, credible regulatory records and commercial partners with defined performance obligations. Capital should arrive in stages that match those facts. Regional value should be protected through measurable access, capability and market commitments.

The central transaction principle is simple: govern the data where responsibility sits; build evidence where science can be reproduced; place IP where it can be protected and developed; licence markets to partners that can execute. This creates a platform capable of serving local health priorities and participating in global innovation.

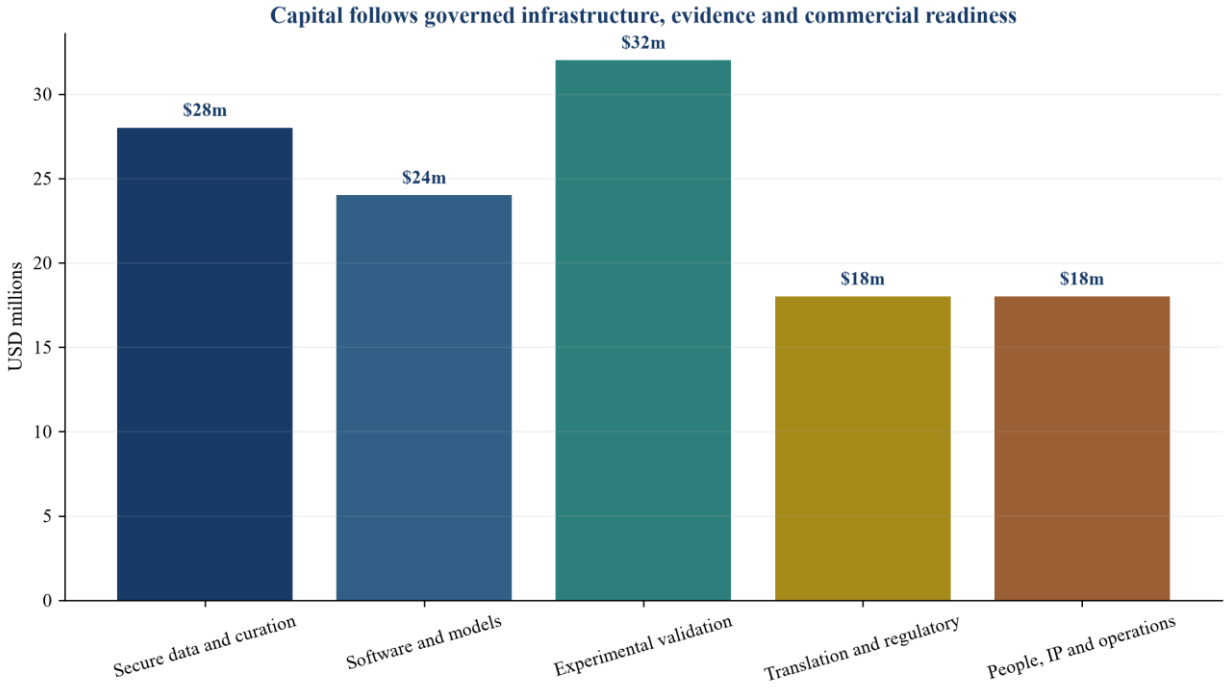

Figure 4. Hypothetical funding allocation by value-building function

Illustrative management assumptions for a USD 120 million platform.

REFERENCES

[1] The Official Platform of the UAE Government. ICT in the health sector: legislations and policies. Updated 8 July 2026. https://u.ae/en/about-the-uae/digital-uae/whole-of-government-approach/digital-health/ict-in-the-health-sector

[2] The Official Portal of the UAE Government. Data protection laws. https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws

[3] Department of Health Abu Dhabi. Health Information Exchange Standards 2026. https://www.doh.gov.ae/en/resources/HIE

[4] Department of Health Abu Dhabi. Data Sharing, Integration and Interoperability Standard. https://www.doh.gov.ae/-/media/E11818DB96604C0DBCC1AEAC6DF16203.ashx

[5] Department of Health Abu Dhabi. Malaffi. https://www.doh.gov.ae/en/featured/malaffi

[6] Dubai Health Authority. Standards for Human Genetic and Genomic Data and Information Governance. 2025. https://dha.gov.ae/uploads/082025/Standards%20for%20Human%20Genetic%20and%20Genomic%20Data%20and%20Information%20Governance2025814423.pdf

[7] Dubai Health Authority. Artificial Intelligence in Healthcare Policy. https://www.dha.gov.ae/uploads/082022/Artificial%20Intelligence%20Policy20228457.pdf

[8] Saudi Vision 2030. The National Biotechnology Strategy. https://www.vision2030.gov.sa/media/iiwlzyo2/national-biotech-strategy-en.pdf

[9] Saudi Ministry of Health. Open Data Policy: Sectorial Health AI and Data Management Office. https://www.moh.gov.sa/en/ministry/opendata/pages/open-data-policy.aspx

[10] Saudi Data and Artificial Intelligence Authority. National Strategy for Data and AI. https://sdaia.gov.sa/en/SDAIA/SdaiaStrategies/Pages/NationalStrategyForDataAndAI.aspx

[11] World Health Organization. Regulatory considerations on artificial intelligence for health. 2023. https://www.who.int/publications/i/item/9789240078871

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

[13] U.S. Food and Drug Administration. Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. Draft Guidance, January 2025. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

[14] U.S. Food and Drug Administration. Guiding Principles of Good AI Practice in Drug Development. 2026. https://www.fda.gov/about-fda/artificial-intelligence-drug-development/guiding-principles-good-ai-practice-drug-development

[15] European Medicines Agency. Reflection paper on the use of Artificial Intelligence in the medicinal product lifecycle. 2024. https://www.ema.europa.eu/en/use-artificial-intelligence-ai-medicinal-product-lifecycle-scientific-guideline

[16] World Intellectual Property Organization. AI-enabled Health Innovation and IP: From Idea to Impact. 2026. https://www.wipo.int/web-publications/ai-enabled-health-innovation-and-ip-from-idea-to-impact/en/

[17] World Intellectual Property Organization. Intellectual Property Rights in Healthcare-related AI. 2 March 2026. https://www.wipo.int/en/web/global-health/w/blogs/intellectual-property-rights-in-healthcare-related-ai

[18] World Intellectual Property Organization. A Primer on Technology Transfer in the Field of Biotechnology. 2025. https://www.wipo.int/publications/en/details.jsp?id=4793

[19] World Intellectual Property Organization. IP Rights Management in Genetic Resources and Data. https://www.wipo.int/en/web/traditional-knowledge/ip-rights-management

[20] World Intellectual Property Organization. Alternative Dispute Resolution for Life Sciences. https://www.wipo.int/en/web/amc/center/specific-sectors/lifesciences/index

[21] IFRS Foundation. IAS 38 Intangible Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[22] IFRS Foundation. IFRS 13 Fair Value Measurement. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

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.