The Face-Value Question: Accuracy, Privacy and Deployment Rights in Biometric M&A

An M&A framework for linking biometric performance to lawful and transferable deployment

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Biometric businesses are often valued through algorithm benchmarks, contracted revenue, enrolled identities and claimed fraud reduction. Those indicators can matter, but none establishes that a buyer can lawfully continue the target's deployed use cases after closing. A facial-recognition model can perform well in a laboratory while failing in the target's operating environment. A customer contract can remain in force while consent, purpose limitation, dataset provenance, controller allocation, retention or transfer terms prevent continued processing. A valuable model can also carry deletion exposure when its training data or derived embeddings were obtained without sufficient rights.

This paper develops a Biometric Deployment Rights Framework for boards, investors and acquirers. It joins four evidence systems that are commonly reviewed in separate workstreams: technical performance, operational outcomes, data and model provenance, and legal deployment authority. The framework begins with each material use case and follows the full chain from collection, enrolment and template creation through matching, threshold decisions, human review, adverse action, retention, deletion and change-of-control transfer. It converts that evidence into a deployment-rights matrix, a performance-to-economics reconciliation, a remediation reserve, transaction protections and a first-one-hundred-days plan.

Primary sources show why the work must be use-case specific. NIST evaluates face-recognition performance across one-to-one verification, one-to-many identification, demographic effects, image quality and presentation attack detection. NIST digital-identity guidance requires defined biometric accuracy, facial presentation attack detection and sensitive-data safeguards in covered authentication contexts. The UK Information Commissioner's Office explains that biometric recognition processes special-category biometric data, requires a lawful basis and separate special-category condition, and generally requires a data protection impact assessment. The EU Artificial Intelligence Act imposes restrictions and obligations for biometric categorisation and remote biometric identification, including narrow conditions for certain law-enforcement deployments. U.S. enforcement and state law show separate exposures around consent, retention, sale, disclosure and derived models. The Federal Trade Commission's Everalbum order required deletion of affected work product derived from improperly obtained data; its Rite Aid action addressed inadequate safeguards and operational harm. Illinois law restricts collection and profit from biometric information, while Texas enforcement has demonstrated material financial exposure.

A wholly hypothetical acquisition case applies the framework to a biometric identity company serving financial institutions, airports, digital platforms and employers. The target reports USD 31.0 million of revenue, 22.0 million enrolled identities and strong benchmark results. Diligence finds that only 68 percent of revenue sits in use cases with verified current deployment rights, representative field performance, complete retention controls and transferable customer terms. The remainder depends

on remediation, renewed consent, customer approval, dataset replacement or discontinued use. Every company, amount, rate and operating assumption in the case is illustrative management input. It is not observed market data, a valuation opinion, a financing proposal or a forecast.

The central conclusion is that biometric value resides in lawful, accurate and transferable outcomes. Algorithm accuracy is an input. Deployment value requires the right population, capture conditions, threshold, safeguards, purpose, jurisdiction, human process and contractual chain. An acquirer should price verified contribution from transferable use cases as current economics; assign conditional value to remediable deployments; reserve for deletion, retraining, consent refresh and customer migration; and connect deferred consideration to evidence that survives change of control.

JEL Classification: G34, K24, L86, O32, O33

Keywords: biometrics, facial recognition, biometric M&A, accuracy, privacy, deployment rights, data provenance, model valuation, consent, due diligence

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 ACQUISITION DECISION

The acquisition decision is whether the buyer can own and continue a set of biometric outcomes that customers are permitted to deploy. A target may possess code, embeddings, datasets, customer contracts and benchmark certificates. The buyer still needs evidence that each material workflow has a lawful purpose, valid authority, suitable performance, effective safeguards and transferable rights.

The analysis should begin with revenue-bearing deployments. For each one, the buyer asks what biometric decision is made, whose data is processed, where collection and comparison occur, which entity determines purpose and means, which threshold is used, what happens after a match or rejection, and which rights survive a change of control.

This approach turns biometric diligence into an operating and valuation exercise. It prevents a strong laboratory result from being treated as a universal licence to deploy. It also gives the investment committee a direct bridge from rights and performance evidence to revenue, contribution, remediation and deal terms.

Table 1. Biometric value layers and acquisition evidence

Value layer	Minimum evidence	Principal exposure	Transaction response
Technical performance	Representative FMR, FNMR, failure-to-acquire and attack testing	Benchmark does not predict field outcome	Reperform tests and condition value
Operational outcome	Funnel from capture to decision and human resolution	Error creates harm or hidden manual cost	Price verified contribution
Data provenance	Collection source, notice, authority, purpose and lineage	Deletion, retraining or claim exposure	Reserve and remediation covenant
Deployment authority	Lawful basis, special-category condition, DPIA and use limits	Use case cannot continue	Exclude or make closing condition
Contract rights	Customer, licensor and processor terms transferable	Consent or licence terminates at closing	Obtain consent before closing
Governance	Threshold ownership, monitoring, appeals and incident evidence	Drift or adverse action is unmanaged	Strengthen controls and warranties
Economics	Revenue and contribution by verified deployment	Enrolment volume without durable value	Underwrite deployment contribution

Each layer requires use-case-specific evidence and a clear transaction response.

2. DISTINGUISH BIOMETRIC TECHNOLOGY FROM A DEPLOYABLE SYSTEM

A biometric algorithm compares features. A deployable system includes cameras or sensors, capture software, quality controls, presentation attack detection, template generation, reference galleries, thresholds, exception handling, human reviewers, customer configuration, data stores and deletion services. Performance and compliance arise from the complete system.

This distinction matters in M&A because the target may control only some components. A licensed recognition engine can rank highly in independent testing while a reseller controls capture, a cloud provider stores templates, a customer chooses thresholds and an outsourced reviewer resolves alerts. The target's revenue may depend on elements it cannot transfer or govern.

The buyer should map ownership, control and evidence for every component. Intellectual property ownership does not establish deployment authority. A licence to use an algorithm does not establish rights to train it on customer images. A right to process templates for one customer purpose does not automatically support another purpose, geography or buyer.

3. CLASSIFY VERIFICATION, IDENTIFICATION AND CATEGORISATION

One-to-one verification asks whether a person matches a claimed identity. One-to-many identification searches a gallery to determine whether a person matches any enrolled reference. Biometric categorisation uses characteristics to place people into categories. These functions have different performance metrics, social effects and regulatory treatment.

The buyer should reject a single label such as facial AI. A remote identification product used in a public space creates different exposure from device-based account verification. A workforce access-control system differs from retail loss prevention. A model that estimates age or emotion may fall under separate rules from one that confirms identity.

Each revenue line should be assigned to a defined function and deployment context. The same code can support several functions, but rights and safeguards may not travel with it. Product roadmaps should therefore be valued by authorised use case rather than claimed technical capability.

4. BUILD THE DEPLOYMENT-RIGHTS GRAPH

The deployment-rights graph links the data subject, collector, controller, processor, technology provider, customer, deployment operator and decision maker. It also connects the capture source, reference database, model, threshold, purpose, jurisdiction, retention rule and contractual permission.

The graph exposes dependencies that a contract list misses. A target may have a direct customer contract while using a third-party dataset under a research-only licence. A customer may control the lawful basis while the target determines model improvement purposes and becomes a controller for secondary use. A distributor may lack authority to approve a change of control.

Concentration matters. A single dataset, cloud service, public-sector approval or platform integration can enable many deployments. It can also create common-mode withdrawal risk. The buyer should value the network of enforceable rights, then stress the removal of each critical node.

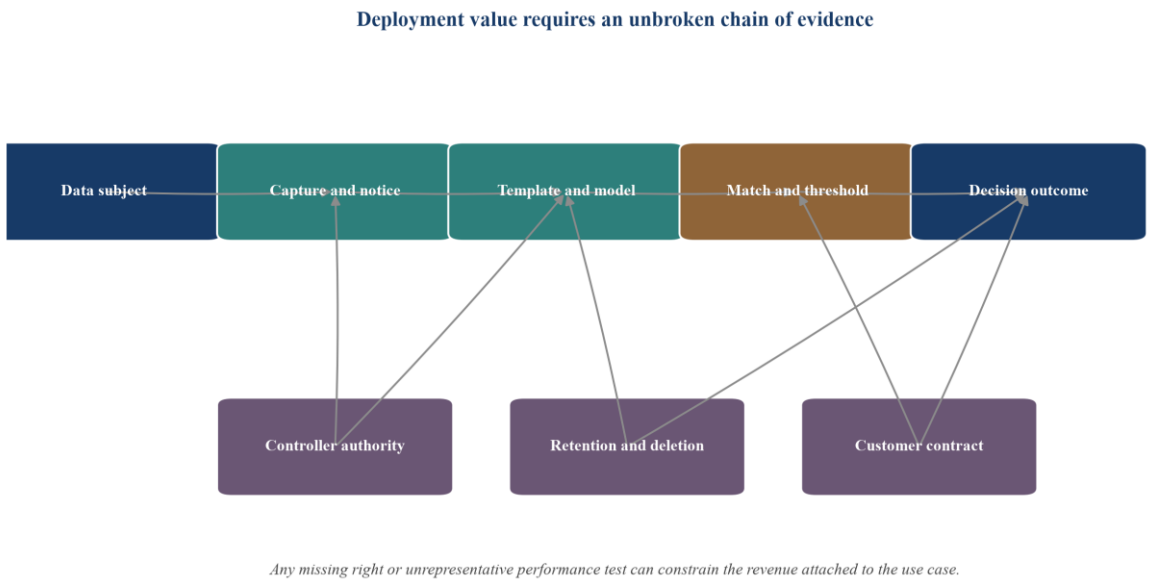

Figure 1. Biometric deployment-rights graph

A valuable deployment connects authority, performance, safeguards and economics through the complete decision chain.

5. DEFINE THE OPERATING POPULATION

Performance depends on the people, images and environments within the deployment. Age, sex, skin tone, image quality, camera position, illumination, motion, occlusion and gallery

composition can affect error. NIST's ongoing evaluation programme separates tasks and reports demographic and quality effects because a single headline accuracy number is insufficient.

The buyer should define the operating population for each material contract. Evidence should include enrolment demographics where lawful and appropriate, capture conditions, devices, geography, gallery size, transaction frequency and consequence of error. The target should explain which groups and conditions were absent from testing.

A representative test does not require indiscriminate collection of sensitive attributes. It requires a documented method, lawful authority, appropriate safeguards and enough evidence to assess the intended deployment. Unknown coverage is a diligence gap, not proof of equal performance.

6. RECONSTRUCT THE RECOGNITION FUNNEL

The recognition funnel begins before matching. A person may fail to present, fail to acquire, fail quality checks, fail liveness, fail template creation, fall below a match threshold or require manual review. Each stage affects conversion, cost, exclusion and customer outcomes.

Targets often report match accuracy while excluding failure to acquire, abandonment and human-review outcomes. That can overstate commercial performance. The buyer should reconcile unique attempts through every stage and segment results by device, customer, geography, use case and relevant population.

The economic denominator should be eligible attempted transactions, not only successfully captured samples. Revenue may be priced per verification while customers bear abandonment and support cost. A buyer that understands the full funnel can distinguish algorithm value from capture quality, workflow design and operational subsidy.

7. SEPARATE FALSE MATCHES FROM FALSE NON-MATCHES

A false match accepts or identifies the wrong person. A false non-match rejects the correct person. Threshold changes alter the trade-off. The appropriate balance depends on consequence, prevalence, gallery size, review process and available fallback.

One error rate cannot represent both risks. In authentication, a false match can enable account takeover while a false non-match creates friction or exclusion. In one-to-many identification, a candidate list may require human confirmation, and the false-positive identification rate must be understood at the chosen gallery size and threshold.

The buyer should obtain detection-error trade-off curves and operating points, then reproduce performance with representative data. Management should show who owns the threshold, how it changes, what approvals are required and how customer configurations are monitored. Uncontrolled thresholds create hidden product variants and uncertain liability.

8. TREAT FAILURE TO ACQUIRE AS A COMMERCIAL METRIC

Failure to acquire occurs when the system cannot produce a usable sample. It may result from lighting, camera placement, movement, disability, age, device quality, network interruption or user guidance. A model can achieve excellent comparison accuracy on accepted images while excluding a material share of attempted users.

The buyer should measure acquisition success at the start of the funnel and segment it. Repeated attempts, fallback to documents, contact-centre intervention and abandoned transactions should

be included in unit economics. These costs can sit outside the target's reported margin when customers absorb them.

A deployable system needs a humane fallback. NIST digital-identity guidance calls for an alternative non-biometric authentication option in covered contexts. The diligence question is whether fallback exists, works, is accessible and preserves customer economics.

9. TEST DEMOGRAPHIC DIFFERENTIALS

NIST has documented demographic differentials across face-recognition algorithms and explains that false-negative effects can be influenced by image quality while false-positive differences can remain even with high-quality images. The target should present results by relevant groups and conditions using a method suitable for the deployment.

The buyer should inspect absolute error, relative disparity, sample size, confidence intervals and operational consequence. A small absolute difference can matter at high transaction volume. A large ratio can be misleading when both rates are extremely low. The investment committee needs both scale and impact.

Mitigation may require better capture, a different algorithm, adjusted workflow, additional review or a non-biometric alternative. Changing thresholds separately by demographic group can raise legal, ethical and operational problems. Evidence should show a fixed, governed approach and outcome monitoring.

10. VALIDATE PRESENTATION ATTACK DETECTION

Presentation attacks use artefacts or manipulated inputs to defeat the system. Examples include printed photographs, replayed video, masks and digitally injected media. Presentation attack detection is a distinct component with its own error rates and attack coverage.

NIST has evaluated passive software-based detection and reported that performance varies by algorithm and attack. NIST digital-identity guidance requires facial presentation attack detection in covered authentication systems and specifies a target impostor attack presentation accept rate for deployment testing.

The buyer should request attack inventories, laboratory protocols, field incidents, device coverage and update cadence. A vendor statement that liveness is enabled is weak evidence. The system should show which attacks were tested, how novel attacks are handled, what happens after failure and whether customers can disable the control.

11. MODEL INJECTION AND DEEPFAKE EXPOSURE

Remote identity systems can be attacked after the camera, including through virtual cameras, compromised devices, API injection or synthetic media. Presentation attack testing at the sensor does not cover the entire path.

The buyer should trace sample integrity from capture to comparison. Device attestation, signed metadata, channel binding, replay protection, integrity checks and anomaly monitoring can reduce exposure. Logs should distinguish live capture, uploaded image, document portrait, stored template and synthetic input.

Deepfake risk should be connected to actual workflows. A model marketed as deepfake-resistant needs defined attack classes and measured performance. The investment committee should avoid

valuing a broad claim. It should price verified resistance at the deployment boundary and reserve for unsupported channels.

12. EXAMINE MORPHING RISK

Face morphing combines characteristics from more than one person so a credential image may match multiple contributors. The risk is relevant to identity-document issuance, border processes and remote enrolment. Detection performance depends on image source, compression, print-scan cycles and attack method.

The target should identify whether it creates, receives or verifies credential portraits and whether morph detection is part of the product. The buyer should inspect testing protocols, false alarms, customer procedures and escalation. A high false-alarm rate can create substantial manual cost.

Standards and evaluation work continue to develop. Draft or emerging methods should be described accurately and valued as future capability until adoption, field performance and customer acceptance are demonstrated.

13. RECONCILE BENCHMARK AND FIELD PERFORMANCE

Independent evaluation is valuable because it provides common data and measurement. It does not guarantee performance in a different population, gallery, device or workflow. The buyer should build a bridge from benchmark task to actual deployment.

The bridge identifies algorithm version, submission date, task type, dataset, threshold, image quality, demographic coverage, gallery size, attack assumptions and hardware. It then maps each item to the customer environment. Any mismatch becomes a test requirement or valuation limitation.

Management should reconcile updates after the published benchmark. A target may market an old strong result while deploying a newer untested model. The buyer should hash production binaries or model artefacts and connect them to release records, test evidence and customer configurations.

14. ESTABLISH MODEL AND DATASET LINEAGE

Model lineage should identify every material training, tuning, calibration and evaluation dataset; the purpose and authority for each use; transformations; exclusions; licences; retention; and derived artefacts. The lineage must extend to face embeddings and synthetic or augmented data.

The FTC's Everalbum matter demonstrates why derived work product matters. The final order required deletion of affected models and algorithms developed from improperly obtained data. In M&A, dataset exposure can therefore affect the model itself, not only the raw images.

The buyer should classify data as owned, licensed, customer-controlled, publicly sourced, consent-based, statutory or otherwise authorised. Public availability does not automatically establish unrestricted biometric use. Missing lineage should trigger targeted testing, legal analysis and a retraining cost estimate.

15. TEST CONSENT AND NOTICE EVIDENCE

Where consent is relied upon, the buyer should examine the actual notice, interface, affirmative action, scope, withdrawal mechanism and audit record. Consent should be specific to the processing purpose and should not be inferred from general service use.

The evidence room should contain versioned notices, consent language, timestamps, jurisdiction, product version and downstream disclosures. The buyer should test whether consent covered model training, service improvement, fraud prevention, third-party sharing, cross-border transfer and acquisition.

Consent can be withdrawn. The operating model therefore needs a way to locate and delete affected samples, templates and derived records where required. A consent count without deletion capability is incomplete. If the use case depends on consent that cannot be refreshed or transferred, value should be conditional.

16. IDENTIFY LAWFUL BASIS AND SPECIAL-CATEGORY CONDITION

In the UK, a biometric recognition system processes special-category biometric data for unique identification. The ICO states that an organisation needs a lawful basis and a separate condition for processing special-category data. The appropriate basis depends on purpose and relationship.

The buyer should review the target's documented basis for every material use case, rather than accepting a global privacy statement. Employment, customer onboarding, public security and fraud prevention can involve different power relationships and conditions. Explicit consent may be unsuitable or unavailable in some contexts.

If no valid condition supports the purpose, the deployment may need to stop. The valuation model should therefore map revenue and contribution to current authority, remediation feasibility and alternative workflows.

17. REVIEW DATA PROTECTION IMPACT ASSESSMENTS

A data protection impact assessment should describe purpose, necessity, proportionality, data flows, risks, safeguards, consultation and residual risk. The ICO states that a DPIA is required before using a biometric recognition system in the circumstances covered by its guidance.

The buyer should assess whether each DPIA matches the live system. Generic templates, unsigned drafts or assessments completed after deployment provide limited comfort. The DPIA should reflect current algorithms, populations, vendors, jurisdictions, thresholds, retention and decision consequences.

A change of control can alter parties, purposes, transfers and security arrangements. The buyer should decide which assessments require refresh before closing, at closing or during integration. Required consultation with a regulator or customer should be treated as a timetable dependency.

18. SEPARATE PURPOSE FROM PRODUCT CAPABILITY

A biometric model can perform many tasks. Legal and contractual authority is usually narrower. Data collected for account verification may not be available for employee monitoring, marketing, emotion analysis or general model training.

The buyer should compare each stated purpose with actual processing, product telemetry and sales material. Secondary uses can emerge through feature reuse, analytics, centralised data lakes or internal experimentation. A broad technical roadmap can therefore create rights debt.

Purpose limitation should be enforced through architecture and governance. Dataset access, experiment approval, feature flags and customer configuration should reflect authorised purposes. Where controls depend on policy alone, the remediation plan should include technical separation and deletion.

19. MAP CONTROLLER, PROCESSOR AND JOINT-CONTROL ROLES

Contract labels do not decide responsibility by themselves. The buyer should examine who determines purpose and essential means, who selects data, who sets thresholds, who retains templates and who uses outputs.

A target may act as processor for customer verification and controller for model improvement or fraud intelligence. It may become a joint controller when it influences purpose or combines customer data. These roles affect notice, lawful basis, rights requests, security, breach response and transfer terms.

The transaction team should build a role matrix for each data flow and compare it with agreements and operations. Misalignment can require contract amendment, product separation or revised notices. Value should follow the role that the buyer can defend and operate.

20. TEST RETENTION AND DELETION

Biometric data requires a defined retention purpose and deletion process. The buyer should inspect schedules for raw images, templates, probes, embeddings, watchlists, logs, reviewer notes, backups and training extracts.

Deletion should be demonstrated, not described. Samples should be traced through primary stores, replicas, analytics platforms, caches and model-development environments. Customer termination, consent withdrawal, failed enrolment and account closure should each trigger the correct process.

Backups create practical latency, but indefinite retention promises create exposure. The buyer should measure deletion completion, exceptions and evidence. A remediation reserve should include engineering, customer communication, data discovery and possible model retraining.

21. REVIEW SALE, DISCLOSURE AND TRANSFER RESTRICTIONS

Illinois law states that a private entity may not sell, lease, trade or otherwise profit from a person's biometric identifier or information. It also establishes notice and written-release requirements for collection in covered circumstances. Other jurisdictions use different definitions and enforcement structures.

An M&A transaction does not automatically resolve whether biometric data, templates or related rights can transfer. The buyer should review statutory rules, privacy notices, customer contracts, dataset licences and consent language. Asset deals can raise distinct assignment and controller questions.

The purchase agreement should allocate responsibility for consents, notices, deletion and excluded data. If a dataset cannot transfer, the buyer should value the remaining model, retraining path and customer continuity rather than assuming the full data asset moves at closing.

22. MAP EU AI ACT RESTRICTIONS AND OBLIGATIONS

The EU Artificial Intelligence Act distinguishes biometric uses and imposes restrictions on certain practices, including narrow conditions for real-time remote biometric identification in publicly accessible spaces for law enforcement. High-risk use can trigger provider and deployer obligations, registration, monitoring and impact assessment.

The buyer should classify each system and intended purpose under the applicable rules and implementation timetable. A product sold as general-purpose infrastructure may become high-risk through its intended use, marketing, configuration or integration.

The diligence file should contain classification reasoning, technical documentation, conformity plans, post-market monitoring, incident processes and customer instructions. Revenue linked to a restricted or unprepared deployment should be separated from compliant current value.

23. EVALUATE OPERATIONAL SAFEGUARDS

Accuracy alone does not prevent harm. Safeguards include notice, accessible alternatives, human review, challenge, reason codes, threshold controls, watchlist quality, reviewer training, escalation, incident response and redress.

The FTC's Rite Aid case addressed facial recognition used for security or surveillance and alleged failures in reasonable procedures and safeguards. The case illustrates that deployment practices, vendor oversight and harm controls matter alongside the model.

The buyer should sample adverse outcomes from alert to resolution. It should inspect how people were notified, whether reviewers had sufficient evidence, how long flags persisted and whether the system learned from correction. Weak safeguards can convert a technical error into regulatory, litigation and reputational exposure.

24. QUANTIFY HUMAN REVIEW

Human review is often described as a safeguard. It can also be a hidden cost or weak control. Reviewers may anchor on algorithm scores, work under time pressure, lack training or see low-quality images.

The buyer should measure referral rate, reviewer agreement, overturn rate, time per case, escalation and downstream outcome. It should test whether review is genuinely independent and whether the person can challenge the result before harm occurs.

Unit economics should include reviewer cost and delay. A lower threshold may increase recall while flooding the review queue. The target's gross margin can change materially when customer operations are brought into the valuation perimeter.

25. CONNECT ACCURACY TO CUSTOMER ECONOMICS

Technical metrics should be translated into conversion, fraud loss, support, review, abandonment and customer retention. A false non-match can create lost sales or exclusion. A false match can enable fraud or adverse action. Failure to acquire can drive people to expensive fallback.

The buyer should build a cost function for each use case using observed volumes and outcomes. It should avoid assigning the same cost to every error. Consequence depends on transaction value, available review and reversibility.

Customer value should be evidenced through cohort performance, not a universal fraud-reduction claim. Strong deployments can support premium pricing and retention. Weak deployments may survive because switching is difficult, which creates revenue durability and liability at the same time.

26. MEASURE DEPLOYMENT-RIGHTS COVERAGE

Deployment-rights coverage is the share of material revenue or contribution supported by current lawful authority, representative performance evidence, required safeguards and transferable contracts. It is a valuation denominator, not a compliance score.

The buyer should grade each use case as verified, remediable, conditional or unsupported. Verified means the evidence is current and reconciled. Remediable means a defined action can close the gap within an evidenced cost and timetable. Conditional depends on a third-party approval or future rule. Unsupported lacks a defensible path.

Coverage should be calculated by contribution as well as revenue. High-margin unsupported deployments can dominate value. The board should also review the lowest-control material use case because average coverage can hide severe exposure.

Table 2. Deployment-rights grading matrix

Grade	Rights evidence	Performance evidence	Economic treatment
Verified	Current authority and transferable terms	Representative field test and monitored outcome	Include in current value
Remediable	Gap has feasible cure and accountable owner	Testing can be completed before or after closing	Deduct cost and defer value
Conditional	Approval, consent or rule remains external	Performance partly evidenced	Probability-weighted value
Unsupported	No defensible authority or transfer path	Benchmark or claim only	Exclude and reserve exposure

Grades apply to a defined use case, jurisdiction, customer and system version.

27. BUILD CUSTOMER AND USE-CASE COHORTS

The buyer should group contracts by biometric function, jurisdiction, customer role, deployment channel, algorithm version and rights grade. Revenue retention should then be analysed by cohort.

This reveals whether growth depends on a narrow use case or permissive jurisdiction. It also shows whether remediation threatens renewal, conversion or margin. A customer may accept the technology while refusing assignment of biometric data or a change in subprocessors.

Cohorts should reconcile contract value, active volume, enrolled identities, support, incidents and contribution. Enrolment counts without active use can overstate network value. The economic unit is a lawful, accurate and retained deployment.

28. TEST CUSTOMER ASSIGNMENT AND CONSENT

Change-of-control and assignment clauses should be reviewed with privacy and operating obligations. A contract may transfer while a customer-controlled dataset, API credential, approval or policy does not.

The buyer should identify customers requiring consent, notice or security review. It should also inspect public-sector procurement rules, approved-vendor status and local hosting requirements. The closing timetable should reflect the longest critical consent path.

Customer communication can itself create churn. The seller should not approach counterparties outside the agreed transaction process. The purchase agreement can use closing conditions, holdbacks or transition services to protect value where approval remains outstanding.

29. VALUE PROPRIETARY DATA CAREFULLY

Biometric datasets can support model performance, testing and product development. Their value depends on lawful acquisition, permitted purpose, representativeness, quality, documentation, exclusivity, refresh and transferability.

The buyer should avoid valuing raw volume. Millions of images can have little use if licences prohibit commercial training, consent is absent, labels are weak or the population does not match deployment. A smaller governed dataset may be more valuable.

Value should be connected to an evidenced improvement, replacement cost and remaining rights period. Where lineage is incomplete, the model should include discovery, deletion and retraining. Derived embeddings and model checkpoints should be included in the analysis.

30. EVALUATE MODEL PORTABILITY

Portability asks whether the model can operate after closing across the buyer's infrastructure, customers and jurisdictions. Dependencies can include licensed libraries, restricted datasets, cloud services, device SDKs and vendor-specific accelerators.

The buyer should reproduce the build, deploy the model in a clean environment and compare outputs. It should confirm access to weights, source code, feature extractors, calibration data, evaluation scripts and signing keys.

Portability also requires rights. A seller may own code while lacking permission to deliver training data or third-party components. The purchase agreement should specify artefacts, licences, assistance and acceptance tests. Missing portability reduces strategic option value.

31. PRICE REGULATORY ACCESS

Regulatory access is the ability to sell and operate a use case in a jurisdiction with the required evidence, approvals and controls. It can create economic value when customers face high compliance costs and trust the provider's documentation.

The buyer should identify which revenue requires registration, conformity assessment, sector approval, data localisation or customer audit. Evidence should include current status, scope, expiry, system version and transfer implications.

Access should not be valued as a permanent moat. Rules, guidance and enforcement change. The model should include maintenance cost, monitoring and potential redesign. A credible

compliance system can support retention and speed to market; unsupported claims should receive no premium.

32. TEST SECURITY AND BREACH EXPOSURE

Biometric characteristics cannot be rotated like passwords. Templates can be protected and transformed, but compromise can create persistent risk. The buyer should review encryption, key separation, access, segmentation, logging, incident response and template protection.

Security testing should cover sensors, endpoints, APIs, data stores, development environments and customer integrations. Central comparison can increase scale of attack, while local comparison can limit exposure and reduce server-side data.

The buyer should inspect incidents, near misses, vulnerability reports and remediation. Insurance, contractual caps and notification duties should be mapped. A breach scenario belongs in the valuation downside and purchase-agreement allocation.

33. CONSTRUCT THE HYPOTHETICAL CASE

The hypothetical target provides face verification, watchlist identification and workforce access products. It reports USD 31.0 million of annual revenue, USD 23.4 million of recurring revenue and USD 18.2 million of gross contribution before product development and central overhead. It serves 84 customers and supports 22.0 million enrolled identities.

Management reports a false match rate of one in 100,000 on a selected verification benchmark and a false non-match rate of 1.2 percent. Field evidence varies. Customer-specific thresholds exist, and three older deployments lack complete failure-to-acquire and demographic monitoring.

Every amount and rate in this case is hypothetical. The case demonstrates the framework and does not describe a named company, market benchmark, expected result or valuation opinion.

34. RECONCILE THE HYPOTHETICAL REVENUE BASE

Diligence maps 68 percent of revenue to verified use cases. These deployments have current authority, representative field evidence, documented safeguards, transferable customer terms and tested deletion.

Another 17 percent is remediable. Gaps include incomplete DPIAs, outdated notices, weak field segmentation and customer consent that can be refreshed. Management estimates a twelve-month remediation programme. The estimate requires engineering plans, customer support and evidence.

Nine percent is conditional on customer approval, dataset relicensing or regulatory classification. Six percent is unsupported because the training-data authority cannot be reconstructed and the customer purpose has expanded beyond the documented basis.

35. CALCULATE HYPOTHETICAL OPERATING ECONOMICS

Verified deployments produce USD 20.4 million of revenue and USD 13.1 million of contribution. Remediable deployments produce USD 5.3 million of revenue and USD 3.1 million of contribution before remediation cost. Conditional deployments produce USD 2.8 million of revenue and USD 1.4 million of contribution. Unsupported deployments produce USD 1.9

million of revenue and USD 0.6 million of contribution. The remainder is non-deployment services.

Management estimates USD 2.6 million of remediation expenditure, including field testing, consent refresh, DPIA updates, deletion engineering and customer migration. A separate replacement-data and retraining scenario ranges from USD 3.0 million to USD 6.5 million.

These are hypothetical management inputs. A real transaction requires ledger, contract, telemetry, consent, data-lineage and customer evidence.

36. MODEL THE RIGHTS-PERFORMANCE FUNNEL

The case begins with 31.0 million attempted annual verification events. Capture succeeds for 29.1 million. Quality and presentation controls accept 27.8 million. The recognition system returns an automated decision for 26.9 million, while 0.9 million go to fallback or review.

Management initially reports accuracy on the 26.9 million decided events. The buyer uses attempted events as the denominator for customer economics. It also segments false matches, false non-matches, failure to acquire, attack rejection, human review and abandonment.

This reconciliation lowers the apparent automation rate and reveals customer-specific cost. It also shows that two contracts with similar model accuracy can have different value because capture conditions and review processes differ.

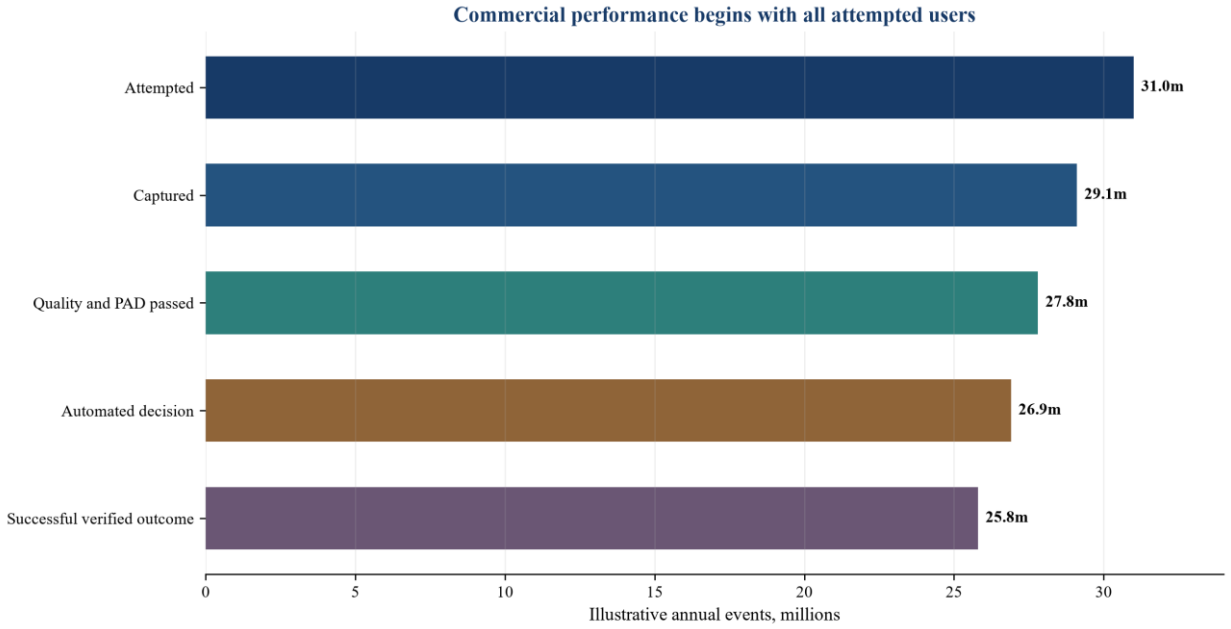

Figure 2. Hypothetical biometric operating funnel

Attempted transactions are the economic denominator; model comparisons cover only part of the customer journey.

37. BUILD AN EVIDENCE-WEIGHTED DEPLOYMENT SCORE

The buyer scores lawful authority, transferability, representative accuracy, demographic evidence, attack resistance, safeguards, deletion, data lineage and operating economics. Current evidence and post-remediation targets are shown separately.

The score is a diligence organiser, not a valuation multiple. Each point must link to a source, period, owner and test. A low score in lawful authority or transferability can outweigh strong technical performance.

The investment committee should review both average and floor. The weakest material deployment may determine regulatory, customer and reputation exposure. Scores should therefore be calculated by use case and contribution, with no unsupported averaging across incompatible workflows.

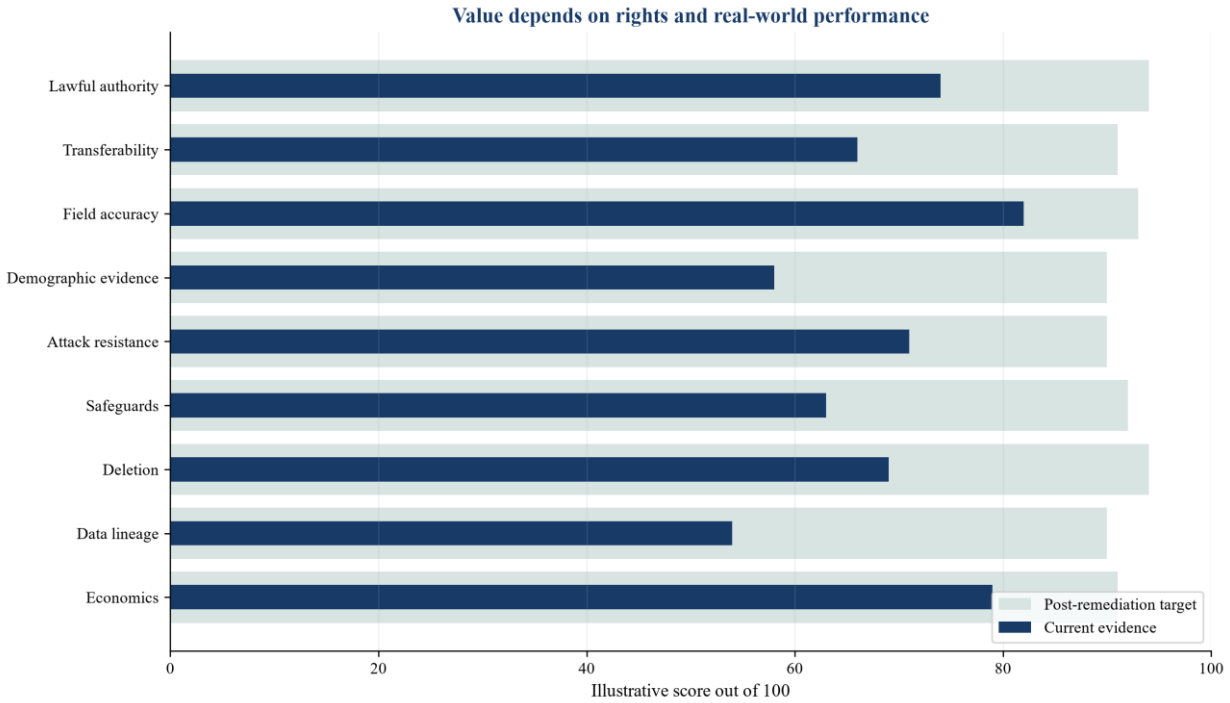

Figure 3. Evidence-weighted biometric deployment score

The hypothetical profile separates current evidence from the post-remediation target.

38. BUILD THE VALUATION BRIDGE

Current value begins with verified recurring contribution. Remediable contribution can be included after deducting cost, delay and execution risk. Conditional deployments should be probability-weighted. Unsupported revenue should be excluded from the base case and assessed for exposure.

The bridge should separately show dataset replacement, model retraining, consent refresh, customer approval, regulatory work and security remediation. Strategic synergies belong after the standalone rights analysis.

Earn-outs and holdbacks can connect price to field evidence, retained customers and completed remediation. The metric should be verified contribution from authorised deployments rather than enrolled identities or raw transaction volume.

Table 3. Hypothetical biometric valuation bridge

Component	Illustrative basis	Value logic	Transaction treatment
Verified contribution	USD 13.1m from supported use cases	Current transferable economics	Cash-flow valuation
Remediable contribution	USD 3.1m before cure cost	Value after evidence and migration	Deferred consideration
Conditional deployments	USD 1.4m contribution	Approval or licence remains external	Probability-weighted earn-out
Unsupported use cases	USD 0.6m contribution	No defensible continuation path	Exclude from base value
Remediation programme	USD 2.6m management estimate	Required to preserve deployments	Purchase-price deduction
Data replacement	USD 3.0m to USD 6.5m scenario	Lineage or licence failure	Reserve and indemnity

All amounts, probabilities and treatments are illustrative management inputs.

39. SELECT TRANSACTION PROTECTIONS

Representations should address dataset provenance, notices, consent, lawful bases, special-category conditions, DPIAs, controller roles, licences, benchmark claims, field performance, thresholds, incidents, complaints, retention, deletion and customer transfer.

Specific indemnities can address identified biometric-law claims, dataset deletion or regulatory investigations. Escrow can fund remediation and retraining. Conditions precedent can require critical customer consent, licence assignment, deletion completion or a clean production test.

Deferred consideration should use measurable outcomes. Suitable metrics include retained verified contribution, completion of agreed field tests, customer approval and closure of defined rights gaps. Raw enrolment or transaction growth can reward volume that increases exposure.

40. BUILD THE DILIGENCE EVIDENCE ROOM

The evidence room should contain data maps, system architecture, model and dataset lineage, licences, notices, consent records, lawful-basis assessments, DPIAs, processor agreements, customer contracts, threshold logs, benchmark submissions, field tests, demographic analysis, attack testing, incidents, complaints, rights requests and deletion evidence.

Samples should cover success, false match, false non-match, acquisition failure, suspected attack, manual review, appeal, consent withdrawal and account deletion. The buyer should trace each sample from capture to final outcome.

Evidence should be graded by source, completeness, version and reconciliation. Slides can explain the system. Production records, executed agreements, customer evidence and reproducible tests should substantiate value.

41. PLAN THE FIRST ONE HUNDRED DAYS

The first phase should freeze new unsupported uses, preserve evidence, protect templates and keys, and establish one inventory of models, datasets, deployments, thresholds, controllers and processors. Critical customer obligations and deletion services require continuity.

The combined company should define one deployment taxonomy, rights matrix and performance denominator. Unsupported data should be quarantined. High-consequence workflows should receive priority for field testing, human review and fallback.

Integration should preserve customer configurations and legal roles until change is assessed. Broad data consolidation can wait until purpose, transfer, retention and security are confirmed. The plan should include owners, milestones and investment-committee reporting.

42. USE A BOARD AND INVESTMENT-COMMITTEE SCORECARD

The scorecard should connect biometric operations to financial and rights outcomes. Measures need precise definitions, stable denominators and reconciliation. A falling false non-match rate can coincide with rising false matches after a threshold change. Higher automation can conceal acquisition failure or weaker review.

Board action should connect to pricing, product investment, customer migration, risk appetite, purchase-agreement protection and capital allocation. The scorecard should show current evidence, target, owner and exception.

Table 4. Board and investment-committee biometric scorecard

Domain	Core measure	Warning signal	Decision response
Rights	Contribution with verified deployment authority	Revenue depends on absent or stale basis	Suspend or remediate use
Performance	FMR, FNMR and failure to acquire at operating point	Benchmark cannot be reproduced in field	Recalibrate and retest
Fairness	Absolute and relative outcome differences	Material group impact or unknown coverage	Improve capture and safeguards
Security	Attack accept rate and integrity incidents	New attack class bypasses control	Restrict channel and update PAD
Operations	Review, appeal and abandonment	Automation shifts cost to customers	Redesign workflow and pricing
Data	Lineage and deletion completion	Unknown source or orphaned copy	Quarantine, delete or replace
Transfer	Customers and licences cleared for closing	Critical consent or assignment unresolved	Condition closing or exclude value
Economics	Contribution by verified deployment	Volume grows without durable rights	Reprice current and contingent value

Measures should be segmented by use case, jurisdiction, customer, device and relevant population.

43. CONCLUDE WITH DEPLOYABLE VALUE

Biometric technology creates enterprise value when it produces accurate, secure and useful outcomes under authority that can survive ownership change. The model, data, workflow and legal permission form one asset system. Weakness in any part can constrain the associated revenue.

An acquirer should begin with material deployments and trace evidence backward. It should define the biometric function, operating population, capture funnel, threshold, attack boundary, human process, lawful authority, data lineage, customer rights and economics. Independent benchmarks support comparison; field evidence establishes actual outcome. Consent, contract and regulatory classification establish the permission to continue.

The disciplined valuation sequence is to price verified contribution as current value, deduct remediation and replacement cost, probability-weight external approvals, exclude unsupported use cases and connect deferred consideration to evidence. This approach allows a buyer to acquire biometric capability while protecting customer trust, individual rights and durable economics.

REFERENCES

[1] National Institute of Standards and Technology. Face Recognition Technology Evaluation. Updated 26 March 2025. https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt

[2] Grother, P.; Ngan, M.; Hanaoka, K. Face Recognition Vendor Test Part 3: Demographic Effects. NISTIR 8280. 2019. https://doi.org/10.6028/NIST.IR.8280

[3] Grother, P.; Ngan, M.; Hanaoka, K. Face Recognition Technology Evaluation: Demographic Effects. Updated 5 March 2025. https://pages.nist.gov/frvt/html/frvt_demographics.html

[4] Grother, P.; Ngan, M. Face Recognition Technology Evaluation: Demographic Summaries. NISTIR 8429. 2022. https://doi.org/10.6028/NIST.IR.8429

[5] National Institute of Standards and Technology. Face Recognition Technology Evaluation 1:N Identification. https://pages.nist.gov/frvt/html/frvt1N.html

[6] Ngan, M.; Grother, P.; Hom, A. Face Analysis Technology Evaluation Part 10: Performance of Passive, Software-based Presentation Attack Detection Algorithms. NISTIR 8491. 2023. https://doi.org/10.6028/NIST.IR.8491

[7] National Institute of Standards and Technology. Digital Identity Guidelines: Authentication and Authenticator Management. NIST SP 800-63B-4. 2025. https://pages.nist.gov/800-63-4/sp800-63b.html

[8] National Institute of Standards and Technology. Facing the Facts to Keep Our Biometrics Secure. 2 October 2024. https://www.nist.gov/blogs/taking-measure/facing-facts-keep-our-biometrics-secure

[9] Information Commissioner's Office. Biometric Recognition. 23 February 2024; guidance under review following the Data Use and Access Act. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/biometric-data-guidance-biometric-recognition/biometric-recognition/

[10] Information Commissioner's Office. How Do We Process Biometric Data Lawfully? https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/biometric-data-guidance-biometric-recognition/how-do-we-process-biometric-data-lawfully/

[11] Information Commissioner's Office. How Do We Process Biometric Data Fairly? https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/biometric-data-guidance-biometric-recognition/how-do-we-process-biometric-data-fairly/

[12] Information Commissioner's Office. How Do We Demonstrate Compliance with Our Data Protection Obligations? https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/biometric-data-guidance-biometric-recognition/how-do-we-demonstrate-our-compliance-with-our-data-protection-obligations/

[13] European Parliament and Council. Regulation EU 2024/1689 laying down harmonised rules on artificial intelligence. 13 June 2024; consolidated version 27 July 2026. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng

[14] European Data Protection Board. Guidelines 05/2022 on the Use of Facial Recognition Technology in the Area of Law Enforcement. Final version, 17 May 2023. https://www.edpb.europa.eu/documents/guideline/guidelines-052022-on-the-use-of-facial-recognition-technology-in-the-area-of_en

[15] Federal Trade Commission. Everalbum, Inc., Decision and Order. 7 May 2021. https://www.ftc.gov/system/files/documents/cases/1923172_-_everalbum_decision_final.pdf

[16] Federal Trade Commission. FTC v. Rite Aid Corporation. Updated 8 March 2024. https://www.ftc.gov/legal-library/browse/cases-proceedings/2023190-rite-aid-corporation-ftc-v

[17] Illinois General Assembly. Biometric Information Privacy Act, 740 ILCS 14. https://www.ilga.gov/Legislation/ILCS/Articles?ActID=3004&ChapterID=5

[18] Office of the Attorney General of Texas. Attorney General Ken Paxton Secures USD 1.4 Billion Settlement with Meta Over Unauthorized Capture of Personal Biometric Data. 30 July 2024. https://www.texasattorneygeneral.gov/news/releases/attorney-general-ken-paxton-secures-14-billion-settlement-meta-over-its-unauthorized-capture

[19] Washington State Legislature. Chapter 19.375 RCW, Biometric Identifiers. https://lawfilesext.leg.wa.gov/Law/RCW/RCW%20%2019%20%20TITLE/RCW%20%2019%20.375%20%20CHAPTER/

[20] Washington State Legislature. RCW 43.386.080, Restrictions on facial-recognition use for surveillance and identification. https://lawfilesext.leg.wa.gov/Law/RCW/RCW%20%2043%20%20TITLE/RCW%20%2043%20.386%20%20CHAPTER/RCW%20%2043%20.386%20.080.htm

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.