PT EN

STRATEGIC EBOOK · EXECUTIVE LEVEL

FOR CEOS, SCIENTIFIC DIRECTORS, CMOS AND BOARD MEMBERS OF THE LARGEST CLINICAL DIAGNOSTIC LABORATORIES AND DIAGNOSTIC MEDICINE GROUPS

The *Blind* Analysis.

How Fully Homomorphic Encryption enables the laboratory to operate genomics, telemedicine, diagnostic AI and collaborative research — without ever decrypting what belongs to the patient.

VOLUME I · EDITION 2026 · CONFIDENTIAL

TABLE OF CONTENTS

PT EN

What you will *read*.

Executive decision document. Written for executive meetings or Saturday mornings before making long-term architecture decisions.

PT EN

The argument in *one* page.

If you read only one thing from this eBook, read this.

The clinical diagnostic lab is the point in the healthcare chain where the most sensitive personal data is generated per unit of time, and where the least has been done to rethink the privacy architecture that governs that data. Every processed sample produces a result that is, at once, a clinical decision, a permanent record, and — increasingly — an input to AI algorithms, international registries, clinical trials and observational research. In some segments (genomics, liquid biopsy, microbiome), the result is a biological identifier as unique as any digital one.

Traditional operations protect this data with familiar layers: TLS in transit, AES at rest, user-level access control, audit logging, rigorous vendor contracts. These layers work for the simplest scenario — delivering the result to the ordering physician. They do not work when the lab wants to do what it increasingly does: offer telemedicine on top of the result, train its own AI over cohorts, participate in international registries, sell aggregated insights to pharma, collaborate in multicenter research, conduct epidemiological surveillance.

Each of these new operations requires the data to be, at some point, decrypted in an environment not fully controlled by the lab — AI vendor, research partner, telemedicine center, scientific repository. Each decryption is a point of legal and reputational failure. And regulation across three converging jurisdictions (LGPD / the Brazilian Data Protection Authority (ANPD), HIPAA/HHS, GDPR / European Data Protection Authorities) is moving rapidly toward requiring **mathematical proof** of minimization, not just governance statements.

1

GENOME IS FOREVER

DNA sequence is a permanent identifier that reveals non-consenting relatives. De-identification is mathematically impossible.

2

AI REQUIRES SCALE

Diagnostic models only become good with volume — and volume requires data sharing between labs, today impossible without exposing patients.

3

PHARMA WANTS REAL DATA

RWE became a regulatory asset accepted by FDA/EMA. The lab is the natural source — but only if it can sell without handing over the patient.

4

FHE IS MATURE

Lattigo, OpenFHE, Concrete are viable. Real cases (Owkin, Tune Insight, Lifebit) already in production at European labs.

PT EN

The thesis of this eBook is direct:

> *"The next decade of diagnostics will be defined by which labs can offer what no one else offers — collaborative science, applied AI and pharma partnerships — without ever decrypting what belongs to the patient."*

The computational cost of FHE is absorbable by the IT budget of any reference lab. The cost of inaction is much greater: being progressively commoditized by AI vendors who extract value from the lab's data, losing the precision medicine game, and being left out of international research consortia. The first to move sets the technical standard the others will have to follow.

THE DECISION

PT EN

The question for the board is not "whether" to invest in private computation. It is "who in your category will lead — and what does it cost to wait until you discover it was someone else".

CHAPTER I

PT EN

The *Results* Industry.

The laboratory has quietly become the point in the healthcare chain where the most clinical data is generated per unit of time and where the least has been done to rethink how that data should be protected.

In 2010, when a patient had a complete blood count, the cycle ended with a PDF delivered to the ordering physician. In 2025, the same test feeds a sepsis prediction algorithm, updates a chronic disease registry, contributes to a proprietary AI model, and — in some labs — is integrated into a longitudinal profile that crosses with genome, microbiome and wearables. Each of these uses is defensible in isolation. The aggregate result, without anyone having decreed it, is an operation of sensitive clinical data at industrial scale mediated by dozens of technology vendors that no one fully controls.

What Changed in the Value Chain

The lab has always been in the business of delivering results to physicians. What changed is that the result stopped being the end of the process and became raw material for five new operations, each with its own technical and legal requirements:

OPERATION	WHAT IT IS	RISK
Diagnostic AI	Models on medical imaging, pathology, ECG	External vendor needs the test in plaintext
Telemedicine	Remote specialist interpreting results	Fragile chain of custody between institutions
Clinical registries	International disease databases for research	Linkage complicates granular consent
Genomics and omics	Sequencing, transcriptomics, proteomics, microbiome	Permanent identification; reveals non-consenting relatives
Epidemiological surveillance	Early outbreak detection via positivity patterns	Reporting without exposing the individual patient is difficult

The Invisible Assets of the Laboratory

It is important to name what is at stake. Reference labs today accumulate a set of assets whose scientific density has no parallel in any other health institution:

- **Longitudinal cohorts per test** — millions of CBCs, lipid panels, glucose readings, hormone panels followed for years on the same patient
- **Digital anatomic pathology** — high-resolution scanned slides, with paired specialist reports
- **High-quality medical imaging** — MRI, CT, mammography, ultrasound with technical standardization few hospitals match
- **Genomic sequences** — SNP databases, pathogenic variants, pharmacogenomics data
- **Microbiome** — bacterial profiles almost as identifiable as DNA
- **Rare biomarkers** — unique cohorts for specific conditions, often not replicable elsewhere

Almost none of these assets can be used outside the lab's silo without massive friction. This is the central contradiction of diagnostics in 2026: **there has never been so much data available, and it has never been so hard to make use of it**.

What no one told the patient

PT EN

The patient who hands over a sample today does not understand, in any practical sense, what they are handing over. They believe they are ordering a test. They are also (potentially) contributing to scientific cohorts, AI model training, registries, and — in some cases — insights sold to pharma. The lab's informed consent form says so somehow. But saying is not understanding, and the industry operates on that difference.

> **THE SILENT PROBLEM**
>
> The current operation depends on a chain of trust between patient, physician, lab, technology vendor and cloud that has dozens of points where data exists in plaintext. Each is a failure point. Regulation is, finally, starting to count the points.

“The question for the lab's board is not whether the current data architecture is sustainable. It is how long until it stops being.”

CHAPTER II

PT EN

The Regulatory *Siege*.

Three converging jurisdictions, one common direction: the end of the era in which technical anonymization and vendor contracts were sufficient defenses.

Clinical data regulation in the laboratory is perhaps the most complex in the health sector because it simultaneously involves the data regulator, the sanitary regulator, the professional council, and — in mature jurisdictions — a specific regulator for diagnostic medicine. Those who operate labs know that each new technical feature goes through a review process that takes months. What many do not perceive is that this pathway is becoming, year after year, more demanding — and that traditional defenses are starting to fail.

LGPD and article 11

Health data is a special category under LGPD (Brazilian data protection law). From the law's standpoint, the lab is a controller of that data — with joint liability for whatever happens with any vendor that has access. In 2025, the Brazilian Data Protection Authority (ANPD) opened the first formal investigation against a large Brazilian lab for data sharing with an AI vendor without a demonstrable legal basis. It was the first of an expected series.

The detail few labs have internalized: responsibility does not end when data is "anonymized". The Brazilian Data Protection Authority (ANPD) follows the European interpretation: real anonymization requires a guarantee of irreversibility, and most traditional techniques do not meet that bar.

HIPAA and the obsolescence of Safe Harbor

Labs operating in the United States are subject to HIPAA. For decades the central defense was Safe Harbor de-identification — removal of the 18 identifiers. This defense was effective while re-identification required substantial effort. Several studies published over the last five years demonstrated re-identification of "HIPAA-compliant" data at rates above 80% via linkage with public datasets.

HHS has already opened multiple investigations against labs and medical centers over use of data that proved re-identifiable. The "we removed the 18 identifiers" defense is no longer a defense — evidence of a minimization attempt that does not understand the current problem.

GDPR and the European EHDS

GDPR is more mature and more aggressive. European labs already face regular fines for weak legal bases in sharing. EHDS — the European Health Data Space — comes into force in phases through 2027 and will create the common European infrastructure for secondary use of clinical data in research. EHDS will require technical capability for privacy-preserving analytics as a prerequisite for participation. European labs without that capability will be excluded due to technical incapacity, not bad faith.

For Brazilian labs, EHDS matters because it defines the technical standard regulators worldwide will adopt in the coming years.

Genomics — the special case

Genomic data is treated in some jurisdictions as a category beyond the special category — a "super-category" that requires renewed consent for each secondary use, specific storage restrictions and a right to be forgotten that conflicts directly with science. California, several European states and the United Kingdom have specific laws. Brazil does not yet, but the trend is to follow.

A lab that performs sequencing and stores results lives under regulatory threat that most other labs do not. FHE is not a luxury here — it is the only architecture that allows defensible scientific operation over the medium term.

The conceptual shift

PT EN

> *“Policy is not enough. Technical proof that data could not have been handled otherwise is required. ”*

It is precisely here that FHE stops being a technical curiosity and becomes a regulatory defense tool. A lab processing data under FHE can demonstrate to the Brazilian Data Protection Authority (ANPD), HHS, the European authority or its own ethics committee that individual data was never accessible to vendors, partners or cloud. It is not policy — it is a theorem verifiable by a third party.

The cost of inaction

PT EN

RISK	PROBABILITY 5 YEARS	IMPACT
LGPD fine for sharing without a robust legal basis	High	2% of revenue or USD 50M+
Brazilian Data Protection Authority (ANPD) / HHS / CNIL investigation over a vendor failure	Medium-high	Reputation + expensive remediation
Class action over genomic data leak	Low-medium	Hundreds of millions
Exclusion from European consortia post-EHDS	High in EU	Loss of cohort access
Loss of contracts with global pharma	Medium-high	USD 5M–50M annually per contract

CHAPTER III

PT EN

FHE in Executive *Language*.

No mathematics. No jargon. Only what the board needs to understand to make a long-term decision.

Imagine a transparent safe. You can see there is something inside, but you cannot see what it is. Now imagine you can manipulate the contents from outside the safe, with mathematical gloves: add, multiply, compare, compute regressions, train models. You execute operations on the safe's contents without ever opening it. When finished, you return the sealed safe to the key owner, who opens it and sees the result. This is FHE, in one sentence.

The conceptual leap

All the cryptography your lab uses today protects data in *two* of the three states:

- **In transit** — between devices. Solved by TLS.
- **At rest** — stored. Solved by AES.
- **In use** — while the server processes. *Here the data has to be in plaintext.*

The third state is the Achilles heel of every privacy architecture in history. When an AI vendor runs inference on a pathology slide, it *has* to see the image. When a researcher runs analysis on a cohort cross-referenced with another lab, someone *has* to see both sets. FHE eliminates the third state. The server processes the data **without ever having access to plaintext**.

How it works

PT EN

> ANALOGY
>
> The lab "locks" its patients' data in a mathematical box using a key that only it holds. It sends the sealed box to a neutral server (cloud, AI vendor, research partner). That server — which never receives the key — executes all desired calculations on the sealed box. The result is a new box, still sealed, returned to the lab. Only the lab opens it.

FHE vs alternatives

TECHNOLOGY	PROMISES	FAILS
De-identification	"We removed identifiers"	Trivial re-identification via linkage
TEE (hardware enclave)	"The chip isolates"	Trusts the manufacturer; side-channel attacks
Federated Learning	"Data stays in the lab"	Gradients leak individual data
Differential Privacy	"We added noise"	Poor for individual clinical decision-making
Synthetic Data	"Generated data"	Does not capture the long tail
FHE	**"Server never sees in plaintext"**	**High computational cost — but decreasing**

FHE is the only technology on this list whose guarantee is mathematical and auditable by a third party.

The three flavors that matter

PT EN

CKKS

ML, STATISTICS, IMAGING

For cohort analysis, AI applied to medical imaging, predictive models. Lattigo and OpenFHE.

BFV/BGV

CLINICAL DATABASES

For exact queries over test databases, counting, stratification.

TFHE

LOGICAL DECISIONS

For eligibility algorithms, rule-based clinical decisions. Concrete (Zama).

Hybrid

IN PRACTICE

Real systems combine two or three per workflow.

The cost myth

FHE has dropped two to three orders of magnitude in seven years. For a lab, most relevant cases (statistics over cohorts, inference on a single test, queries) today run in seconds to minutes — perfectly compatible with non-emergency clinical workflows. For very high-volume cases (training over millions of images), start with selective cases where the ROI already closes.

For a lab with annual revenue above USD 500M, total investment in FHE — initial capex plus annual opex — lands below 0.7% of the IT budget. It is less than many labs spend on a single LIS migration.

CHAPTER IV

PT EN

Use Cases by *Line*.

What concretely changes in each vertical of the laboratory. Genomics, telemedicine, AI, research, surveillance, biobank.

Genomics and precision medicine

Genomics is the case where FHE stops being a luxury and becomes practically mandatory. A DNA sequence is a permanent identifier that reveals non-consenting relatives. De-identification is mathematically impossible — any clinically useful sequence length is unique in the world. And the scientific value of genomic data is so high that the whole industry wants access.

1. PHARMACOGENOMIC TESTS ON ENCRYPTED GENOMES

The patient sequences their genome. The lab stores it encrypted, with the patient's key. When the physician wants to check whether the patient has variants affecting a specific drug's metabolism, the algorithm runs over the ciphertext — without the lab ever decrypting it. The result goes straight to the patient/physician. **The lab never has the genome in plaintext**, even as the operator of the service. This is the only architecture defensible in the long term for genomics.

2. POLYGENIC RISK UNDER CIPHER

Polygenic risk scores (PRS) are the frontier of predictive medicine. Computing them requires looking at variants across the whole genome. Under FHE, the calculation happens over the encrypted genome, and the patient receives only the final score.

3. INTERNATIONAL GENOMIC COHORTS

To discover rare variants associated with disease, you need to look at large cohorts. Each cohort is legally complex. A Brazilian lab with a significant cohort can contribute by encrypting locally and participating in international studies without the data ever leaving the Brazilian server — solving both the scientific and sovereignty problems.

AI applied to medical imaging

PT EN

1. DIAGNOSTIC INFERENCE WITHOUT SENDING THE IMAGE

An AI vendor offers a model for breast detection, retinopathy, melanoma, pathology. Today, this requires sending the image in plaintext to the vendor's server. Under the correct architecture, the feature-extraction CNN (ResNet-50 or equivalent, ~25M parameters) runs *locally at the lab*, in plaintext, over the high-resolution image. It produces a 2,048-dimensional embedding that captures all relevant diagnostic information. Only the embedding is encrypted and sent. The vendor runs the final linear classifier over the encrypted embedding. **The vendor never sees the image or the embedding in plaintext — this is the real production pattern used by Owkin in partnerships with Sanofi and Roche, Lifebit in genomics, Mozaic in general radiology.**

2. TRAINING PROPRIETARY MODELS WITHOUT MOVING DATA

The lab wants to train a proprietary model over its own images. The compute infrastructure is in the cloud (cheaper, faster). Under FHE, training happens over ciphertexts in the cloud. Data never leaves the lab's cryptographic control.

3. EXTERNAL VALIDATION OF AN INTERNALLY DEVELOPED MODEL

Lab A developed a pathology classification model. To publish, it needs to validate on an external cohort. Lab B is willing, but does not want to expose its cohort. FHE allows validation under cipher: A's model runs over B's encrypted data, the performance metric is computed without either side seeing what belongs to the other.

Telemedicine and second opinion

Cross-institutional telemedicine is a case where the traditional chain of custody is inherently fragile. A remote specialist receives an image by email, opens it, analyzes it, returns a report. Each of these steps creates copies and logs that no one fully controls. Under FHE, the test circulates encrypted, the specialist performs analysis under cipher, the report returns without any of the original test having existed in plaintext outside the originating lab.

Multicenter clinical research

This is the area of highest scientific value. Epidemiological studies, biomarker validation, longitudinal cohort analysis — anything that requires crossing data between centers. Today it requires months of

DUAs and ethical approvals. Under FHE, the cross-referencing happens over ciphertexts, with ethical governance that is **more robust**, not less. Time to first analysis: 60–70% reduction.

Private epidemiological surveillance

The lab is often the first to detect an outbreak — a spike in positivity for a specific pathogen in a region. Today, reporting that information to a health authority or sharing it with other labs is legally delicate. Under FHE, it is possible to operate a surveillance network where labs contribute encrypted data, abnormal patterns are detected statistically over ciphertexts, and alerts are issued without exposing individual patients or revealing specific labs.

Biobank and stored samples

Biobanks are the lab's long-term asset. The content (physical samples + associated data) only gains scientific value if it can be used for future research — including research that does not yet exist today. Broad consent is legally fragile. Under FHE, it is possible to operate a biobank in which data associated with samples stays permanently encrypted, and new research runs over the ciphertexts. Each new study does not require a new broad consent, because no individual data is technically exposed.

Aggregated insights for the pharmaceutical industry

Pharma wants to buy lab data to understand real-world drug usage, disease patterns, treatment effectiveness. Today, this market exists via anonymized data (IQVIA/Close-Up model) — legally fragile. Under FHE, pharma sends an encrypted query, the lab computes over the encrypted base, and returns only aggregate statistics. Pharma pays, the lab books revenue, the patient was never exposed. **It is a new recurring-revenue market that is blocked today by the absence of technical architecture.**

Clinical decision-making under cipher

PT EN

Clinical decision algorithms (cardiovascular risk calculators, sepsis scores, therapy eligibility) can run over the patient's encrypted data, returning an answer without the central calculator ever having seen the patient's parameters.

CHAPTER V

PT EN

The Economics of the *Encrypted Lab*.

The real numbers. How much it costs, how much it returns, and where the returns appear.

The cost of doing it (initial capex)

COMPONENT	INVESTMENT
Founding team (1 senior cryptographer, 2 ML engineers, 1 clinical PM, 1 legal)	USD 4M – 6M / year
Licenses and tooling	USD 200k – 800k / year
Infrastructure (GPU, AVX-512, optional FPGA)	USD 1.5M – 3M initial
Strategic consulting	USD 800k – 2M
Regulatory study	USD 400k – 1M
Integration with LIS, RIS, PACS, sequencers	USD 1M – 3M
Total year 1	**USD 8M – 16M**

The cost of operating (annual opex)

PT EN

ITEM	ANNUAL ESTIMATE
Compute	USD 2M – 6M
Maintenance team	USD 4M – 7M
Audit	USD 500k – 1.5M
Stabilized annual opex	**USD 6.5M – 14.5M**

For a lab with revenue above USD 1B (Dasa, Fleury, Hermes Pardini, Sabin), this represents between **0.6% and 1.3% of revenue** — comparable to the cost of a single regional expansion or an LIS migration.

The return — six vectors

1. New revenue from aggregated insights for pharma

The RWE market from labs to pharma is estimated today at USD 5B globally, dominated by IQVIA. The entry of labs with auditable FHE architecture captures premium segments (oncology, rare diseases, expensive therapies) where the client wants more assurance. Estimate for a Brazilian reference lab: **USD 20M–80M in incremental annual revenue within three years**.

2. New revenue from diagnostic AI under FHE

The lab can offer "guaranteed-private AI" as a premium service, charging a premium per test. In high-end segments (oncology, radiology), the premium patient accepts paying 20–40% more for that guarantee. Estimate: **USD 10M–50M annually**.

3. Access to international research funding

Labs with the technical capability to participate in international consortia capture NIH, Wellcome and Horizon Europe funding they do not reach today. Estimated increase: **USD 5M–30M annually within**

36 months.

PT EN

4. Enabling defensible clinical genomics

Genomics is a high-growth segment but structurally held back by privacy concerns. FHE enables an aggressive offering. For a lab entering genomics hard with FHE as a differentiator: **USD 30M–150M of value enabled over 5 years**.

5. Reduction of regulatory risk

Expected exposure to fines/litigation: USD 30M–100M NPV over 5 years. FHE as hedge: **USD 12M–50M in insured value**.

6. Position in AI consortia

AI vendors need cohorts to train. Labs with FHE become preferred partners for these vendors — rather than commoditized suppliers. Margin capture: **USD 5M–25M annually**.

Summary business case

~USD 12M YEAR 1 INVESTMENT	~USD 10M STABILIZED ANNUAL OPEX
USD 100M+ VALUE ENABLED OVER 5 YEARS	10×–20× EXPECTED 5-YEAR ROI

PT EN

“For any reference lab with scientific ambition, FHE is the digital transformation investment with the highest return asymmetry available in 2026.”

CHAPTER VI

PT EN

Competitive Advantage and *Positioning*.

A lab that adopts FHE first becomes an elite technical partner — for hospitals, pharma, academic hospitals and the premium patient segment.

The clinical testing industry has, for twenty years, been dominated by a commoditization logic: winning on price, scale, speed and operational excellence. That remains true for most of the market. But there is a segment — small in volume, large in margin — that does not compete on price: the premium segment, the scientific segment, the rare-disease segment, the genomic segment, the applied-AI segment. In that segment, elite technical partners win. And that is exactly the segment FHE unlocks.

The three possible postures

Posture 1 — The Lab Leading in Defensible Genomics

Focus on genomics and precision medicine. FHE as a technical prerequisite. Direct communication with the patient: "your genome is never decrypted, not even by us". Works best for labs with a strong bet on precision medicine.

Posture 2 — The Pharma Partner Lab

Focus on RWE and insights for the pharmaceutical industry. FHE as the technical differentiator that unlocks premium partnerships. Recurring pharma revenue becomes a pillar. Works for labs with enough scale to generate interesting cohorts.

Posture 3 — The Academic-Scientific Lab

Focus on collaborative research, publishing, international funding. FHE as infrastructure fo participation. Works for labs tied to a university hospital or with scientific aspirations.

The three reinforce one another. A robust strategy combines at least two.

The cost of not taking a position

A scenario that must be made explicit at the board: what happens if none of the Brazilian labs adopt FHE structurally in the next 36 months? Answer: **AI vendors and RWE platforms will capture the value**. Owkin, ConcertAI, Tempus, PathAI will offer "AI-as-a-service" and "RWE-as-a-service", extracting value from Brazilian labs' data and turning the labs into commoditized suppliers. Within five years, the gap will be structural.

> *"A lab that leads with FHE is not merely adopting technology. It is protecting its place among the labs that will still matter a decade from now."*

CHAPTER VII

PT EN

24-Month *Roadmap*.

From the board's decision to the first commercial offering under FHE.

01

MONTHS 1–6 · LEARN

Foundation and capability

Hire a founding crypto engineer or partner with a university (USP, Unicamp, EPFL via Tune Insight). Identify three candidate use cases with clear ROI. Align with the institutional Research Ethics Committee (CEP/CONEP) and privacy counsel.

Output: documented architecture, three selected cases, favorable legal-ethical opinion.

02

MONTHS 7–14 · BUILD

Internal pilot

Build one end-to-end use case. Recommendation: encrypted statistical analysis over an internal cohort, OR AI inference on an encrypted test. Validate latency, integration with LIS/PACS, key governance.

Output: functional demo, metrics validated by a third party, documentation ready for an external partner.

PT EN

03

MONTHS 15–20 · FIRST CLIENT

First commercial offering under FHE

Launch the first commercial product using FHE. Suggestion: a private pharmacogenomic test service, or an RWE contract with pharma as a pilot. Marketing aimed at the premium/scientific segment.

Output: first contract closed, first patient served, first revenue attributable to the architecture.

04

MONTHS 21–24 · INSTITUTIONAL CAPABILITY

Adoption as a strategic pillar

Multiple products on the infrastructure. Commercial team training. Structured external communication. Presence at congresses (DAI, ASCO, SBPC). Possible first consortium with another lab.

Output: 3+ active commercial products, first external consortium, sector recognition.

CHAPTER VIII

PT EN

Risks, Mitigations and *Pitfalls*.

What can go wrong, in order of probability.

1 · Inability to hire talent

High probability. **Mitigation:** a university partnership or capability acquisition via a specialized consultancy. The lab does not need to have talent in-house from day one.

2 · Integration with legacy systems

A typical lab runs 10–30 systems (LIS, RIS, PACS, sequencers, ERPs). **Mitigation:** integrate only the 2–3 that matter for the first use case.

3 · Internal cultural resistance

Clinical operations are averse to technical novelty. **Mitigation:** treat it as a cultural project before a technical one. Engage medical champions early.

4 · Computational cost at high volume

Training AI over millions of images is still expensive. **Mitigation:** start with selective cases. Do not try to replace everything.

5 · CEP or legal blocking

An innovative proposal may be rejected out of excess caution. **Mitigation:** co-design with the Research Ethics Committee (CEP/CONEP) / legal from the start. FHE as a strengthening tool, not a bypass.

6 · Competitor announces first

Mitigation: speed. Every month of delay is a month of risk.

Pitfall 1 · Treating it as an IT project

Reporting to the CIO instead of the Scientific Director or CMO. Result: technical delivery, z commercial impact. **FHE should report to the Scientific Director/CMO with CEO sponsorship.**

Pitfall 2 · Starting with the most ambitious case

Trying to start with an international consortium. Mistake. Start internally, validate, expand.

Pitfall 3 · Forgetting key governance

FHE protects during computation. Key management across the lab, patient, CEP and external partner is half the project.

A letter for the next decade of *diagnostics*.

For the Scientific Directors, CEOs and Board Members of labs that can still choose to lead.

The industry you lead was built on a simple, old promise: that the technical analysis of a biological sample can help heal a person. That the result of a test, delivered on time and with quality, is the oldest and most reliable instrument of evidence-based medicine. That the lab is the point in the healthcare chain where technical precision meets clinical care, and where the patient's trust is translated into medical decisions.

This promise lasted for decades. It survived equipment transformations, automation, industrial scale, consolidation. It survived because it was — and largely still is — true. The patients who hand over samples every day trust, deeply, that laboratory with something extraordinarily intimate: pieces of their own body, and the biological story those pieces tell.

But in the last fifteen years, without anyone having decreed it, what happens to that sample has changed completely. The result stopped being a PDF handed to the physician. It became input to algorithms, registries, AI models, pharma partnerships, epidemiological surveillance, collaborative research. Each of these uses is defensible in isolation. The aggregate result is an industrial-scale clinical data operation, mediated by dozens of vendors, where the data exists in plaintext at points that no one fully controls.

It is possible to turn back. More than that: it is strategically preferable to turn back. FHE is the first technology in decades that allows offering all the new services — AI, telemedicine, collaborative research, pharma partnerships — **without ever decrypting** what belongs to the patient. It is possible to keep doing everything the modern lab needs to do. It is possible to do all of it while the institution keeps, with verifiable mathematical proof, that the individual patient was never exposed.

What is at stake is not a technical feature. It is the possibility for the lab to become again, unambiguously, what it has always claimed to be: a space of technical precision and respect for the patient, now with an instrument worthy of the modern challenge.

> *“Within three years, someone will lead. Th question is whether it will be your lab, or the one you will have to look to as a reference. ”*

There is a window. It is short. It is real. Whoever reads this eBook holds a map. The rest is courage.

— End of Volume I

APPENDIX A

PT EN

Executive *Glossary*.

FHE

Cryptography that allows computing over encrypted data without decrypting it.

RLWE

Mathematical problem at the base of modern FHE, the same as NIST PQC.

CKKS, BFV/BGV, TFHE

The three main schemes in practical use.

LIS, RIS, PACS

Laboratory Information System, Radiology Information System, Picture Archiving — the typical legacy stack of the lab.

PRS — POLYGENIC RISK SCORE

Polygenic risk score — the frontier of predictive genomic medicine.

RWE — REAL WORLD EVIDENCE

Clinical evidence derived from real-world usage data, accepted as regulatory support by FDA, EMA and the Brazilian Health Authority (ANVISA).

EHDS

European Health Data Space — defines the technical standard global regulators will adopt.

THRESHOLD CRYPTOGRAPHY

Distribution of a key across multiple parties with a required quorum.

PSI — PRIVATE SET INTERSECTION

Lets two parties discover an intersection without revealing the remainder.

FEDERATED LEARNING

Distributed training where data stays local. FL+FHE eliminates gradient leakage.

PT EN

LATTIGO, OPENFHE, CONCRETE

Main FHE libraries in practical use.

APPENDIX B

PT EN

Vendors and *Partners*.

VENDOR	FOCUS
Owkin (Paris/NY)	FL+FHE for clinical research; partnerships with top labs and centers
Tune Insight (Switzerland/EPFL)	Lattigo; focus on multicenter medical research
Lifebit (UK)	Federated genomics for biobanks and labs
Zama (Paris)	Concrete framework, TFHE
Duality (USA/Israel)	OpenFHE, focus on health, consulting
Inpher (Switzerland/USA)	Hybrid FHE+MPC, focus on health and finance
Stickybit (Brazil)	Brazilian technical boutique in FHE/PQC; tailored architecture

Relevant initiatives

- **EHDS** — European Health Data Space
- **GA4GH** — Global Alliance for Genomics and Health
- **SPHN** — Swiss Personalized Health Network
- **NIH N3C** — National COVID Cohort Collaborative

APPENDIX C

PT EN

30 Questions for the *Scientific Director and CIO*.

Strategy

1. Do we have advanced cryptography talent? Acquisition plan?
2. What is our current exposure to processing patient data?
3. How many external vendors have technical access to our data?
4. Is there an up-to-date inventory of which data leaves the lab?
5. Legal opinion on the sustainability of the current architecture in 36 months?

Priority use cases

6. Which commercial areas are most constrained by the inability to cross-reference or apply AI?
7. In which genomic segments would it make sense to offer FHE as a differentiator?
8. Which AI vendors would we like to use but do not because of data exposure?
9. Which pharma partnerships are blocked or fragile today?
10. Which international consortia would be worth joining?

Technical

11. Which FHE scheme for our first case?
12. Can we run statistical analysis over an internal cohort in FHE today?
13. Estimated computational overhead?
14. How do we integrate with LIS, PACS, sequencers?
15. How do we manage keys across lab, patient, partner?
16. Is threshold cryptography compatible with our flow?

Cost

PT EN

17. Cost per FHE analysis vs plaintext?
18. Accelerators evaluated?
19. Build in-house, university partnership, or vendor?
20. Projected capex and opex over 24 months?
21. C-level sponsor confirmed?

Regulation

22. Institutional Research Ethics Committee (CEP/CONEP) engaged?
23. Independent security audit?
24. Demonstrable LGPD/HIPAA/GDPR compliance?
25. Dialogue with the Brazilian Data Protection Authority (ANPD)?
26. Communication to the patient?

Commercial

27. Which premium segments would pay for an FHE guarantee?
28. Differentiated pricing for FHE-backed services?
29. Which pharma partnerships would be unlocked?
30. Worst-case scenario if a competitor announces first?

PT EN

Blind Analysis

Strategic eBook for the senior leadership of clinical diagnostic laboratories.
Volume I · Edition 2026 · Confidential distribution.

Set in Iowan Old Style and SF Pro.
Built as a self-contained HTML document.

— end —