

STRATEGIC EBOOK · EXECUTIVE LEADERSHIP

---

FOR CEOS, CSOS, CMOS (CHIEF MEDICAL OFFICERS), HEADS OF R&D, CHIEF DATA OFFICERS AND BOARDS OF THE WORLD'S LARGEST PHARMACEUTICAL COMPANIES

# The *Invisible* Trial.

*How Fully Homomorphic Encryption redefines clinical research, drug discovery, real-world evidence and pharmaceutical intellectual property — and why the first house to understand this will redesign the economics of the industry.*

---

VOLUME I · EDITION 2026 · CONFIDENTIAL

# What you will *read*.

*This eBook is a decision document. It was written to be read in an executive committee meeting, on a plane ride, or on a Saturday morning before a long-term investment decision.*

00	<b>Executive Summary</b> <i>The argument in one page</i>	3
I	<b>The Industry of Data That Heals</b> <i>Why pharma became—without realizing it—the world's largest operation of sensitive clinical data</i>	5
II	<b>The Regulatory Tightening</b> <i>HIPAA, GDPR, ICH-GCP, FDA, EMA, ANVISA—and the end of performative anonymization</i>	9
III	<b>FHE in Executive Language</b> <i>What it is, without the math—and why it changes industry economics</i>	13
IV	<b>Use Cases by Area</b> <i>Clinical trials, RWE, discovery, oncology, rare disease, biologics, manufacturing, pharmacovigilance</i>	17
V	<b>The Economics of the Invisible Study</b> <i>Real costs, ROI, and where billions emerge</i>	29
VI	<b>Competitive Advantage and Positioning</b> <i>Why this is regulatory architecture, not IT</i>	35
VII	<b>24-Month Roadmap</b> <i>From learning phase to first public study</i>	39

VIII Risks, Mitigations, and Pitfalls

*What can go wrong and how to avoid it*

---

PT EN

IX Manifesto

*A letter for the next decade of data-driven medicine*

---

49

· Appendices

*Glossary, vendors, 30 questions for the CSO/CIO*

---

53

# The argument in *one* page.

*If you are only going to read one thing from this eBook, read this.*

**T**he global pharmaceutical industry is facing, simultaneously, the largest scientific opportunity and the biggest operational bottleneck in its history. On one side, the explosion of clinical, genomic, real-world and imaging data creates conditions for unprecedented precision medicine. On the other, that same data is so sensitive, so regulated and so proprietary that it becomes structurally unusable outside the silos where it was created.

Every major pharmaceutical company in the world operates today on a central contradiction: **it has more data than it has ever had, and can use less of it than it should.** Clinical trials cost two billion dollars and take a decade because crossing data between centers is legally heavy. Real-world evidence is the regulatory obsession of the decade, but it depends on payer, hospital and pharmacy data that no one can nominally combine. Collaborative drug discovery is the obvious future, but it requires two pharma companies to share libraries that neither one trusts to the competitor.

At the same time, three forces are converging:

## 1

### TIGHTENING REGULATION

HIPAA, GDPR, LGPD (Brazilian data protection law), ICH-E6(R3), European EHDS — all requiring mathematical proof of minimization and provenance, not promises of governance.

## 2

### RWE AS MANDATE

FDA, EMA and ANVISA (Brazilian health authority) accepting real-world evidence as regulatory support — but demanding quality and governance that almost no one delivers today.

## 3

### MATURE FHE

## 4

### OPEN WINDOW

Lattigo, OpenFHE, Concrete in production.  
Real use cases (MELLODDY, Owkin, Roche)  
already demonstrating technical and  
economic viability.

No top 10 has yet anchored structured  
internal private computation capability.  
Whoever arrives first defines the standard  
the others follow.

PT EN

The thesis of this eBook is direct:

---

*“The next decade of pharma will be defined by those who first learn to collaborate without trusting — on clinical, genomic and proprietary data, without ever decrypting it.”*

---

FHE — *Fully Homomorphic Encryption* — is the technology that makes this mathematically possible. It is not a governance promise. It is a theorem. The server that computes on encrypted patient data **cannot** see the data, even if it wants to, even if it is breached, even under court order.

The computational cost is high and will be discussed honestly. But for an industry where a single phase III trial costs two billion, where a single molecule saves or destroys entire quarters, and where competitive advantage lives in intellectual property and proprietary data — the math is already settled. The question for the board is no longer "whether" to invest in private computation. It is "which will be our first anchor use case, and when do we start".

#### THE DECISION

It is not "whether" the pharmaceutical industry will adopt private computation. It is "who" will lead the definition of the sector standard and capture the regulatory, scientific and competitive advantage of the next decade.

# The Industry of *Data That Heals*.

*How a century-old industry of molecules and trials transformed — without decree and without manifesto — into one of the largest operations of clinical, genomic and biometric data on the planet. And how that same data became, at the same time, central asset and largest operational liability.*

In 2005, when one of the great pharmaceutical houses conducted a phase III clinical trial, the work was essentially an exercise in medical logistics. Hundreds of sites distributed across the world, thousands of recruited patients, physical paperwork, monitoring through on-site visits, data consolidated in spreadsheets, statistical analysis at the end. Costly, yes. Slow, yes. But technically understandable, legally clear, and operationally well understood.

In 2025, the same trial is a real-time data operation. Wearables send continuous patient data directly to the study system. MRI images are processed automatically by AI models. The genome is sequenced in the first 30 days. Biomarkers are monitored by telemedicine. Each patient generates, over the course of the trial, between 50 and 500 times more data than twenty years ago. And each of these data points is potentially identifiable, eternally sensitive, and legally loaded.

This transformation happened without the governance models, the technical architecture, and the relationship between pharma, hospital, payer and patient keeping pace. The result is an industry that has more data than it has ever had, operates under increasingly aggressive regulation, and still depends on a privacy architecture that was designed for a world that no longer exists.

## What changed in the value chain

Pharma has always been an intellectual property industry sustained by clinical trials. What has changed in the last ten years is not the structure of that chain — it is that each link in it became an industrial-scale data operation. The five new assets:

ASSET	WHAT IT IS	RISK
<b>Electronic health record</b>	Clinical history of patient in trial or in RWE study	Special category · fragile chain of custody
<b>Genomics and omics</b>	Sequencing of DNA, RNA, proteome, metabolome, microbiome	Eternal identification · reveals non-consenting relatives
<b>Medical imaging</b>	MRI, CT, PET, retinography, digital pathology	Re-identification by facial features or rare patterns
<b>Real-world data</b>	Claims, dispensations, outpatient records, wearables	Cross-referencing leaked sources reconstructs the patient
<b>Molecule libraries</b>	Chemical structures, screening data, ADMET, targets	Intellectual property worth hundreds of millions

None of these five assets existed, at scale, in pharmaceutical operations in 2010. In 2026 they are the five pillars of the scientific operation of any top 20.

## What no one told the patient

The patient who signs an informed consent form for a clinical trial today does not understand, in any practical sense, what they are handing over. They believe they are authorizing a controlled medical intervention — and they are. But they are also authorizing their genome to be stored for decades on a multinational's server, their medical images to be used to train AI models, their wearable to transmit continuous data to systems on three different continents. The consent form says this in some way. But saying is not understanding, and the industry operates, today, on that gap in understanding.

This is the hard part. And it is the part that will collapse first, because it is the part European regulators have already started to move on.

The pharmaceutical industry operates, today, depending on a chain of trust between patient, doctor, sponsor, CRO, technology vendor, and cloud — which has dozens of points where the data exists in the clear. Each of these points is a point of failure. And regulation is, finally, starting to count the points.

## The historical parallel — and why pharma is different

The social media industry has already been through its Cambridge Analytica. The financial industry has already been through its Equifax. The healthcare industry has not yet been through its event, but the accumulated pressure is greater than both combined — because clinical and genomic data are, legally and morally, the most protected category that exists.

The difference is that the pharmaceutical industry has, unlike the other two, a structural advantage: **the relationship with the patient is, ultimately, about healing.** Pharma can choose to be demonstrably trustworthy. It can choose to be the first to say "we coordinated research across five hospitals without ever seeing the individual patient of any". This move — made in time — is simultaneous competitive and regulatory advantage. Made late, it is a chapter in a crisis book.

---

*“The question for any pharma executive committee is not whether the current data architecture is sustainable. It is how much time remains until it ceases to be.”*

---

# The Regulatory *Tightening*.

*Four continents, dozens of regulations, one common direction: the end of the era in which anonymization and contractual governance were sufficient defenses.*

There is a comfortable misconception shared by most legal leadership at large pharma companies: that current clinical data regulation, although heavy, can be managed on the basis of traditional defenses — well-written informed consent, technical anonymization, CRO contracts, international transfer clauses, competent DPO. This view is correct for 2020. It is wrong for 2026.

## HIPAA — when anonymization stops being enough

US HIPAA allowed, for decades, clinical data to be shared in "de-identified" form — with the 18 direct identifiers removed. This defense worked in a world where re-identification required substantial effort and external databases were rare. That world is over. Several studies published in the last five years have demonstrated re-identification of "anonymized data" at rates above 80% using only three to five quasi-identifier attributes cross-referenced with public databases.

The practical result: HHS has already opened multiple investigations against medical centers and technology vendors for use of "anonymized data" that was re-identified. The defense "we removed the 18 identifiers" is no longer a defense — it is evidence of attempted minimization without understanding the problem.

## GDPR and the European interpretation

GDPR is more demanding. Health data is a special category under article 9, requiring specific and robust legal basis. European national authorities — French CNIL, Italian Garante, Spanish AEPD, UK ICO — have already produced dozens of decisions sanctioning pharmaceutical research projects for fragile legal

basis or insufficient anonymization. Fines in millions of euros, and more importantly: **orders to cease processing** that halted ongoing studies.

PT EN

In 2024, the EHDS — European Health Data Space — came into force, allowing secondary use of clinical data for research, but **requiring a technical governance architecture** that almost no pharma has implemented today. EHDS is the instrument that, within three years, will redefine who can operate clinical research in Europe. Those without privacy-preserving analytics capability will be excluded due to lack of technical qualification, not lack of money.

## ICH-E6(R3) — the operational turning point

The ICH-E6(R3) revision, published in 2023, marks a fundamental change in global Good Clinical Practice: for the first time, the technical and privacy documentation of a trial must demonstrate **verifiable** minimization, not declared. Regulators have started asking, in inspections, "how do you *prove* the sponsor did not access individual data it should not have"? The answer "we have a policy" is no longer enough. The correct answer will be, in three years, "we have an auditable cryptographic architecture".

## FDA, EMA, ANVISA and the acceptance of RWE

FDA, EMA and ANVISA (Brazilian health authority) today accept real-world evidence as support for regulatory decisions — including indication expansion and post-approval confirmatory studies. This unlocks billions in value for sponsors, but with a counterpart: RWE must have sufficient scientific quality and data governance to sustain the decision. The three authorities have published, over the last two years, guidance documents detailing expectations — and the expectations converge on a single principle:

---

*“It is not enough to say the data was handled with care. You must mathematically prove it could not have been handled otherwise.”*

---

This is exactly where FHE stops being a technical curiosity and becomes a **structural regulatory advantage**. A pharmaceutical company that runs its RWE study on encrypted data can demonstrate to the FDA, EMA or ANVISA — and to the ethics committee — that the sponsor never accessed individual data, that the local investigator never saw data from other centers, that the statistical analysis happened without anyone being able to intervene in the data before lock. This is not a policy. **It is a theorem verifiable by a third party.**

## Brazil, China, India — regulatory fragmentation

Brazil's LGPD (Brazilian data protection law) has article 11 placing health data as a special category — comparable to GDPR in formal rigor, although still immature in case law. The Brazilian Data Protection Authority (ANPD) began signaling, in 2025, specific enforcement on clinical research. China has PIPL, perhaps the world's most aggressive law on international transfer of health data — any data collected from a Chinese patient must pass national security assessment to leave the country. India has the Digital Personal Data Protection Act, new and in the regulation phase, but clearly following the European model.

For a global pharma, this means the current operation depends on maintaining **five to eight parallel regulatory architectures**, each with its specific governance, its data localization, its transfer restrictions. It is operationally expensive and strategically fragile — any change in one jurisdiction creates a wave of rework in the others.

### THE STRATEGIC WINDOW

Regulators are seeking "exemplary" use cases of privacy-preserving research to cite in guidelines and decisions. The first top 10 to present an auditable FHE architecture for RWE or a clinical trial becomes an official citation in the three major authorities — regulatory protection + competitive positioning in a single move.

# The cost of not acting, in numbers

PT EN

RISK	5-YEAR PROBABILITY	TYPICAL IMPACT
GDPR/LGPD fine for secondary use without legal basis	High	2–4% of global revenue
Halt of ongoing trial by European DPA	Medium	6–18 month program delay
Rejection of regulatory submission due to data governance failure	Medium-high	12–24 month approval delay
US class action over data use without robust consent	Medium	USD 500M–2B (precedents)
Exclusion from European consortia due to technical incapacity	High in EU	Loss of access to EHDS and centers of excellence
Breach of clinical data from trial patient	Medium	24–48 month reputational crisis + litigation costs

# FHE in *Executive* Language.

*No math. No jargon. Just what senior management needs to understand to make a US\$ 50 million decision.*

Imagine a transparent vault. You can see there is something inside, but you cannot see what it is. Now imagine that you can, from outside the vault, with mathematical gloves, manipulate the contents: add two things in there, multiply, compare, compute regressions, train models. You execute operations on the contents of the vault without ever opening it. When you finish, you return the closed vault to the key owner, who opens it and sees the result. This is Fully Homomorphic Encryption, in one sentence.

## The conceptual leap

All the cryptography your company uses today — TLS on the portal, AES on backups, HTTPS on APIs — protects data in *two* of the three possible states:

- **In transit** — between devices. Solved by TLS.
- **At rest** — stored. Solved by AES.
- **In use** — when the server processes. *Here the data must be in the clear.*

The third state is the Achilles heel of every clinical privacy architecture in history. When the trial system calculates the partial statistic, it *needs* to see patient data. When the AI model runs on the exam, it *needs* to see the image. It is at that moment that the data is vulnerable to a dishonest employee, to intrusion, to misconfigured backup, to improperly rotated log, to subpoena, to careless CRO, to fragile integration between vendors.

FHE eliminates the third state. The server processes data **without ever having access to plaintext**. This is a phase change, not an incremental improvement.

# How it works, in an analogy that holds up to scrutiny

PT EN

The actual mathematical mechanism involves lattices and the RLWE problem — the same problem on which the next generation of post-quantum cryptography is built. But the executive intuition is as follows:

## ANALOGY

The hospital "locks" its patients' data into a mathematical box using a key only it has. It sends the closed box to the multicenter study's server. The server — which *never* receives the key — executes the entire statistical analysis algorithm (mean, variance, regression, survival, AI model) *on the closed box*, producing a new box with the result, still closed. It returns it to the sponsor. The sponsor, who has another key (generated via threshold cryptography), opens only the aggregate result. At no point in the chain does anyone — not the server, not the sponsor, not the remote hospital — see the individual data of any patient.

# What FHE offers that nothing else offers

PT EN

TECHNOLOGY	WHAT IT PROMISES	WHAT FAILS
HIPAA Anonymization	"We removed the 18 identifiers"	Trivial re-identification via cross-referencing; legally fragile
TEE (hardware enclave)	"The chip isolates"	Trusts the manufacturer; several side-channel attacks published
Federated Learning	"Data stays at the site"	Gradients leak individual data; already demonstrated in clinical research
Differential Privacy	"We add noise"	Good for aggregate statistics, bad for individual efficacy decisions
Synthetic Data	"Artificially generated data"	Does not capture the long tail; fragile causal inference; does not replace real data
<b>FHE</b>	<b>"Server never sees in the clear"</b>	<b>High computational cost — but decreasing</b>

FHE is the only technology on this list whose guarantee is **mathematical and auditable by a third party**. For FDA, EMA, the Brazilian Data Protection Authority (ANPD) or a European ethics committee, it is the difference between "trusting governance" and "verifying mathematically".

## The three flavors that matter

### CKKS

APPROXIMATE · REAL NUMBERS

The flavor for machine learning, statistical analysis, medical image processing. Allows

### BFV/BGV

EXACT · INTEGERS

The flavor for encrypted clinical databases, exact counts, stratification. When the result must be identical to plaintext.

multiplications over large vectors.  
Implemented by Lattigo and OpenFHE.

PT EN

## TFHE

BOOLEAN · ULTRAFINE

Bit-by-bit logic, comparisons, arbitrary programs. Slower per operation, but the most flexible. Ideal for clinical eligibility decisions.

## Hybrid

IN PRACTICE

Real systems combine two or three.  
Radiological inference in CKKS, event counts in BFV, eligibility in TFHE.

## The cost myth: what you need to know

The standard argument against FHE is "it is too expensive". In 2018, this was true. In 2026, it is a half-truth that needs to be carefully dismantled:

- **Cost has dropped two to three orders of magnitude** in seven years. The curve is constant.
- **For statistical analysis on clinical cohorts** (means, variances, logistic regression, Kaplan-Meier on thousands of patients), latency today is minutes to a few hours — commercially viable for any study whose interim analyses happen in weekly cycles.
- **Hardware accelerators** (Intel HEXL, FPGAs, specific ASICs) promise another 10–100× in the next 24 months.
- **The right model is selective:** use FHE at the sensitive core (cross-site matching, federated gradient aggregation, eligibility scoring) and keep trivial operations in the clear locally. This reduces total cost to a fraction of the current clinical IT budget.

The cost of FHE is not a barrier. It is a **design variable**. For an industry that routinely invests US\$ 2 billion in a single phase III trial, the computational overhead of FHE in cases where it makes a real difference is, literally, noise on the spreadsheet.

# Use Cases by *Area*.

*What concretely changes in each vertical of the pharmaceutical operation. Clinical trials, RWE, discovery, oncology, rare disease, biologics, pharmacovigilance, manufacturing — each with its specific opportunity and each with its distinct ROI.*

## Decentralized clinical trials

Decentralized trials (DCTs) are the industry's largest operational transformation in the last five years. The patient at home, the wearable, telemedicine, the PRO app. All of this generates rich, continuous, and legally loaded data at unprecedented volume.

FHE use cases:

### 1. CONTINUOUS MONITORING UNDER ENCRYPTION

The patient's wearable encrypts the data locally before sending. The study system computes pre-defined alerts (adverse events, clinical deterioration, imminent dropout) over the ciphertext. The study monitor receives only the relevant alerts, never raw data. **The sponsor never sees continuous individual patient data** — only aggregate metrics and the alerts that justify intervention. This solves, at once, the "too many sponsor people can see everything about patients" problem.

### 2. INTERIM ANALYSIS UNDER ENCRYPTION

Futility and interim efficacy analyses are critical moments in a trial. Today they require partial database unblinding, with all the governance complexity. With FHE, interim analysis happens on the ciphertext, and only the final statistic is decrypted (via threshold key shared between DSMB and sponsor). Result: more frequent interim analyses, with less risk of involuntary unblinding.

### 3. PRO AND EMOTIONAL DATA

Patient-reported outcomes capture psychological, behavioral and quality-of-life data. This data is especially sensitive and patients tend to under-report when they know they will be seen. Collecting

PROs under encryption assures the patient that no one will see their individual response — which increases truthfulness and adherence. Several preliminary studies show that private collection raises response rate by 20–40%.

PT EN

## Real-world evidence (RWE)

RWE is the regulatory holy grail of the decade. FDA, EMA and ANVISA (Brazilian health authority) accept RWE as support for decisions — as long as data quality and governance are sufficient. This is the area where FHE has the greatest immediate ROI, because it unlocks cases that are legally impossible today.

### 1. CROSS-REFERENCING BETWEEN SPONSOR, HOSPITAL AND PAYER

To prove a drug works "in the real world", you must cross-reference dispensation (pharmacy/payer), utilization (hospital record), and outcome (clinical). Each belongs to a different entity, and none of the three can deliver nominal data to the others. FHE enables the matching without any of the parties seeing the others' individual data. Result: RWE studies that do not happen today because there is no legal basis — start happening with solid legal basis.

### 2. COMPARATIVE EFFECTIVENESS

Comparing two drugs in the real world requires adjustment for observed confounders — age, sex, comorbidities, prior therapies. Each of these adjustments requires individual data. Under FHE, statistical adjustment (propensity score, IPTW, doubly robust) happens on encrypted data, without anyone needing to see the patient.

### 3. POST-MARKETING STUDIES REQUIRED BY REGULATORS

FDA and EMA frequently require, as a condition of approval, post-marketing safety studies. These studies are expensive and slow today because they involve multiple centers and payers. FHE drastically reduces time to first analysis, because it eliminates the need for bilateral DUAs (data use agreements) negotiated over months.

## Drug discovery

### 1. SHARED CHEMICAL LIBRARIES

The most celebrated case of private computation in pharma is MELLODDY — a consortium of ten pharma companies (Janssen, Bayer, Boehringer, Novartis, Servier, AstraZeneca, GSK, Merck KGaA,

Amgen, Astellas) that trained bioactivity models via federated learning on their combined libraries, without any of them seeing the others' data. MELLODDY used federated learning without FHE which still allows partial leakage via gradients. **The next generation of this kind of consortium will be FHE-secured**, and the first house to position itself as technical leader of this movement captures the central relationship with all the others.

## 2. VIRTUAL SCREENING IN THE CLOUD WITHOUT EXPOSING IP

Pharma wants to run heavy molecular simulation (docking, MD, QSAR) on cloud GPUs (AWS, Azure, GCP), but the formula is the central asset. Today: either run on-prem (expensive, slow) or trust the cloud provider. FHE allows outsourcing of computation **without revealing either input or output**. CKKS is especially good for numerical simulation.

## 3. SHARING OF NEGATIVE DATA

Data on molecules that *did not* work is almost as valuable as data on those that did, and the industry wastes billions duplicating negative experiments because no one shares. An FHE consortium for "negative data" is the kind of initiative that saves money without handing over IP.

# Oncology — the most sensitive and most valuable case

Oncology is where everything converges: tumor genomic data, imaging (digital pathology, MRI), comorbidities, medication, long-term outcome. Serious disease, patient identifiable by rare characteristics, extremely high scientific value, maximum legal sensitivity.

## 1. MULTICENTER TUMOR BOARDS UNDER ENCRYPTION

International tumor board on a complex oncology case: 4 oncologists in 4 countries discuss the case. Today: exchange of de-identified emails, with a fragile chain of custody. With FHE: the case circulates encrypted, each oncologist annotates observations over the ciphertext, no one needs to see the patient's name. Applicable to both individual cases and cohorts.

## 2. AI INFERENCE WITHOUT SENDING THE EXAM

AI models for immunotherapy response prediction, druggable variant identification, and risk stratification already exist. Today they require the patient or hospital to send the exam in the clear to the server hosting the model. In the right architecture, the feature-extraction network (heavy CNN) runs *locally at the hospital*, producing a dense embedding that captures all the diagnostic information. Only the embedding is encrypted and sent. The vendor executes the final classifier under FHE and returns the

encrypted score. **The vendor never sees the exam or the embedding in the clear** — it is the pattern used by Owkin, Lifebit and Mozaic in production today.

PT EN

### 3. COLLABORATIVE GENOMIC STUDIES IN RARE CANCER

Rare cancer has patients in such small volumes that no single hospital has a sufficient cohort on its own. Cross-referencing databases is legally heavy in any jurisdiction. FHE enables the cross study without moving data, and — more importantly — without needing secondary consent for "cross-referencing" because technically the matching happens under encryption.

## Rare diseases

Rare diseases share with rare cancer the problem of a tiny cohort. The difference is that regulation for rare diseases is more favorable (orphan drug status, fast track), and pharma companies that dominate the category have direct relationships with patient associations. There is a specific window here:

FHE enables the creation of **international rare disease registries** where patients contribute encrypted data via an app, and pharma/researchers access only aggregate statistics. The patient association holds the key. This simultaneously solves the scientific problem (sample size) and the ethical problem (patient control). Several associations have already sought such partnerships — what is missing is a technical counterpart able to deliver.

## Biologics, biosimilars and advanced therapy

Gene therapy, CAR-T cell therapy, biosimilars — segments where the patient is rare, the product is very expensive, and evidence must be accumulated on each individual patient over years. Long-term RWE here is not a luxury, it is a regulatory requirement.

FHE enables longitudinal follow-up of a CAR-T patient for 10 years without anyone having to store an accumulated record in the clear. The patient holds the key, contributes data periodically, receives follow-up, and the sponsor obtains the aggregate evidence required by the FDA — without violating anything.

## Collaborative pharmacovigilance

Detecting a rare adverse event depends on correlating data across countries, hospitals, payers and manufacturers. No single source has enough sample. Today pharmacovigilance operates on voluntary

reports and manual cross-referencing — systematically slow to detect rare signals.

PT EN

FHE enables a global pharmacovigilance network where data sources contribute encrypted data to a neutral server (perhaps under WHO or ICH), and safety signals are detected statistically over the ciphertext. This drastically reduces time to detection of rare events — and saves measurable lives.

## Manufacturing and supply chain

A less obvious but important case. Pharmaceutical manufacturing involves proprietary process data (fermentation parameters, chromatography conditions, impurity profiles) that are extremely high-value intellectual property — and that must be partially shared with vendors (CDMOs), regulators (FDA, EMA), and quality partners.

FHE allows the CDMO and sponsor to compute on process data without either of them revealing their full IP. And it allows quality data to be reported to regulators in encrypted aggregate form, with auditing possible through a shared key.

## Financial evidence and value-based contracting

Outcome-based contracts (the payer only pays if the patient responds) are the trend of the decade. They require continuous matching between dispensation (payer), utilization (hospital) and outcome (clinical). FHE allows this matching to happen without the sponsor or payer needing to see the individual patient — fulfilling the contract without violating privacy.

# The Economics of the *Invisible Study*.

*The real numbers. How much it costs, how much it returns, and where capital meets value — at pharmaceutical industry scale.*

Every senior management investment decision must pass through three sieves: capex, recurring opex, and net present value discounted at a realistic cost of capital. FHE is no exception. But unlike marketing-scale or e-commerce technologies, in pharma the calculation must be made against the correct backdrop: the industry where a single phase III trial costs US\$ 2 billion and where a single molecule defines entire quarters.

# The cost to build (initial capex)

PT EN

COMPONENT	TYPICAL INVESTMENT
Founding team (1 senior crypto engineer, 2 ML engineers, 1 clinical PM, 1 clinical privacy counsel)	USD 1.5M – 2.5M / year
Licenses and tooling (Lattigo open, Concrete commercial, OpenFHE)	USD 80k – 350k / year
Infra: GPUs and CPUs with AVX-512, optional FPGA, validated GxP environment	USD 500k – 1.5M initial
Strategic consulting (Zama, Duality, Inpher, Owkin) for architecture	USD 300k – 800k
Regulatory study with a clinical-research specialized firm	USD 250k – 700k
CSV (computer system validation) for GxP use	USD 200k – 500k
<b>Total year 1</b>	<b>USD 3M – 6M</b>

# The cost to operate (recurring annual opex)

PT EN

ITEM	ANNUAL ESTIMATE
Compute (FHE 100x–1000x more expensive than plaintext in the core operation)	USD 1.5M – 4M
Maintenance team (5–8 engineers + 1 regulatory advisor)	USD 2.5M – 4M
Annual security audit and revalidation	USD 300k – 800k
<b>Stabilized annual opex</b>	<b>USD 4.3M – 8.8M</b>

For a top 20 pharma, this represents between **0.02% and 0.08% of revenue**. To put it in perspective: it is less than most houses spend on *a single regional sales force meeting*. For a global top 5, it is a rounding error in the R&D budget.

## The return — where the money reappears

The ROI of FHE in pharma comes from six vectors that must be modeled separately, each with magnitude far exceeding cases in other sectors:

### 1. Reduction of clinical trial time

The opportunity cost of each month of delay in a phase III trial is typically **USD 30M–100M in lost revenue** (especially in oncology and rare diseases). FHE reduces time to first interim analysis, reduces time negotiating DUAs between sites, and eliminates governance rework. Conservative estimate: **1–3 months per trial**. For a top 20 with 30 active phase III trials, this is **USD 1B–10B in present value**.

### 2. Enabling RWE as a regulatory asset

RWE submitted to FDA/EMA is today rejected in 40–60% of cases due to insufficient governance or quality. FHE raises quality and provenance, increasing acceptance rates. Each successful RWE submission is worth between USD 100M and USD 2B in indication expansion or accelerated approval.

Conservative capture: 2–5 additional submissions per year in a top 20 = USD 200M–10B in incremental value.

PT EN

### 3. Access to consortia and EHDS

The European Health Data Space, once fully operational, will require technical privacy-preserving analytics capability to access secondary data. Those without the capability will be excluded. EHDS gives access to cohorts of hundreds of millions of Europeans. Estimated value: **USD 500M–3B in enabled research over the next 5 years.**

### 4. Protection of chemical library IP

The annual risk of library data leakage (internal or via vendor) is estimated at **USD 50M–500M in expected present value.** FHE for HPC in the cloud drastically reduces this risk. Treatment as a hedge: **USD 30M–300M in insurance value.**

### 5. Reduction of regulatory governance costs

Each multicenter trial today spends **USD 500k–2M on data governance and DUAs.** FHE eliminates most of this cost by replacing contractual governance with technical governance. For a top 20, estimated annual savings: **USD 30M–100M.**

### 6. Advantage in strategic partnerships

Tier 1 academic hospitals (MD Anderson, Mayo, Charité, Karolinska, Hospital das Clínicas) are increasingly cautious in partnerships requiring nominal sharing. Sponsors with FHE capability become viable partners in collaborations competitors cannot enter. Worth **5–10 exclusive strategic partnerships in 36 months,** each with a value between USD 20M and USD 200M.

## The business case, summarized

~USD 5M

YEAR 1 INVESTMENT

~USD 7M

STABILIZED ANNUAL OPEX

USD 1B+

VALUE ENABLED OVER 5-YEAR  
HORIZON

100x+

EXPECTED ROI OVER 5-YEAR  
HORIZON

PT EN

---

*“In any honest modeling for a top 20 pharma, FHE is the digital transformation investment with the greatest return asymmetry available in 2026.”*

---

Not because a high return is certain. But because the downside is trivially small (known cost, perfectly budgetable) and the upside is structurally asymmetric — each of the six vectors above, in isolation, justifies the investment. The six combined make it one of the obvious investments of the decade.

# Competitive Advantage and Positioning.

*In pharma, FHE is not narrative first — it is regulatory and scientific advantage first. Narrative comes later, and is a consequence.*

There is an important difference between the FHE thesis for cosmetics and the FHE thesis for pharma. In cosmetics, the central value is brand narrative — "we never see your face" is a promise that sells expensive cream. In pharma, narrative is secondary. The central value is **operational and regulatory capability that unlocks cases impossible today**. The first house to master FHE in pharma will not sell more because of a campaign. It will approve more drugs, faster, in more indications, with less regulatory friction. This is worth exponentially more than any narrative.

## The three types of advantage FHE unlocks

### Regulatory advantage

Whoever is first to submit an RWE submission with an auditable FHE architecture becomes a regulatory reference. FDA, EMA and ANVISA (Brazilian health authority) need exemplary cases to cite in guidelines — and the first example shapes the sector standard. It is not an exaggeration: the first top 20 to do this will influence the content of the next three rounds of guidance documents from the three authorities.

### Scientific advantage

Whoever masters FHE can run collaborative studies competitors cannot run. Rare cancer in a five-hospital cohort across three continents, without moving data, without years of ethical approval, without institutional friction. Research that takes 5 years today comes out in 18 months. This translates into more publications, more citations, more scientific reputation — and more talent attracted.

## Direct competitive advantage

PT EN

Access to partnerships with elite academic hospitals. Access to European consortia via EHDS. Access to shared chemical libraries. Access to international rare disease registries. Each of these doors opens only to those with the technical capacity to operate under private computation.

## The three possible positionings

### Positioning 1 — The Regulatory Standard

Focus on becoming a sector reference. Heavy investment in publishing architecture, contributing to guidance documents, participating in multi-stakeholder initiatives with FDA, EMA, EHDS. The house becomes a mandatory citation. Works best for global top 5 with significant regulatory muscle.

### Positioning 2 — The Scientific Leader

Focus on elite collaborations with top 10 global academic hospitals. The house becomes the preferred partner for sensitive multicenter studies. Works best for houses with strong academic presence (Roche, Novartis, Merck KGaA, GSK).

### Positioning 3 — The Rare Disease Specialist

Focus on the rare disease and rare cancer segment, where FHE unlocks genuinely impossible cases. Builds a direct relationship with patient associations and international registries. Works best for houses with rare disease portfolios (Takeda, Sanofi Genzyme, Pfizer Rare Disease, Alexion).

The three are not mutually exclusive. A robust strategy combines all three at different levels of the organization: positioning 1 at the corporate level, positioning 2 in R&D units, positioning 3 in specific business units.

## The cost of not positioning

There is a scenario that must be made explicit at the board: what happens if none of the top 20 adopt FHE structurally in the next 36 months?

Answer: **tech players will do it.** Owkin, Lifebit, Mozaic, ConcertAI, and new entrants will offer "FHE-as-a-service for pharma" and capture the intermediate position. Top 20s will end up buying these services at high multiples — handing over data, operational dependency and value capture to vendors that could have been built internally. It is exactly what happened with clinical AI infrastructure over the last five years: those who outsourced early pay dearly today. Those who do not outsource now pay even more dearly in three years.

PT EN

---

*“The choice is not between adopting FHE or not. It is between having internal capability or renting it at rising tech multiples.”*

---

# 24-Month *Roadmap*.

*From the board decision to the first clinical study submitted with auditable FHE architecture. Four phases, clear milestones, exit metrics for each.*

## 01

MONTHS 1-6 · LEARN

### **Foundation and internal capability**

Hire the founding senior crypto engineer. This hire is the real bottleneck — there are perhaps 200 qualified people in the world. Recommendation: partial acquisition of a startup (Zama, Owkin, Inpher, Tune Insight) is often the fastest shortcut. In parallel: engage strategic consulting for initial architecture. Reproduce public benchmarks. Identify three candidate use cases with clear ROI and select them for pilot.

**Exit metric:** documented technical architecture, three selected cases (ideally one RWE, one clinical trial and one discovery), internal legal opinion validating feasibility under GDPR/HIPAA/LGPD (Brazilian data protection law).

## 02

MONTHS 7-14 · BUILD

### **Internal pilot in a controlled environment**

Build a single use case, end to end, in a controlled environment. Recommendation: encrypted statistical analysis on a synthetic cohort mimicking an RWE study. This is the most technically mature case and the one with the greatest demonstrative value.

Validate latency, cost, precision, key flow, audit chain. In parallel: begin structuring the material for regulatory dialogue.

**Exit metric:** functional demo on a synthetic cohort, metrics validated by an independent third party, technical documentation ready for IRB and regulator submission.

## First real study, under controlled ethical scrutiny

Launch the first real clinical or RWE study using the architecture. Suggestion: a small-scale study in a segment well dominated by the house. Ideally: a rare disease or rare cancer, where the value is most obvious. Patients explicitly consent to the new architecture. IRB receives complete technical documentation. The result is validated against a parallel study in traditional architecture.

In parallel: begin formal dialogue with FDA, EMA and ANVISA (Brazilian health authority) about using the architecture in future submissions. Engage pre-submission meetings.

**Exit metric:** first real study completed, favorable IRB opinion, first positive regulatory feedback from at least one of the three authorities.

## Adoption as an institutional capability

Make the architecture available to multiple ongoing programs. Train R&D teams on how to propose FHE in new designs. Publish a technical whitepaper. Present at DIA, ASCO, ESMO, regulatory conferences. Submit the first formal study using the architecture as part of the regulatory package.

This is the phase where the investment from the previous 20 months begins to generate returns at scale. Done correctly, it generates sustained sector recognition and a preferential position in future collaborations.

**Exit metric:** 3+ internal programs using the capability, first formal regulatory submission incorporating FHE architecture, recognition in at least one public FDA/EMA guidance.

# Milestones the board should demand

PT EN

MILESTONE	WHEN	ACCOUNTABILITY
Founding crypto engineer hired or acquisition defined	Month 4	Without this, there is no project
Use case selected and validated by clinical legal	Month 6	Without legal alignment, endless friction
Functional technical demo on synthetic cohort	Month 14	Proof of real feasibility
First real study started with IRB approval	Month 18	Simultaneous ethical and technical validation
First positive regulatory feedback	Month 22	Sign that the investment converts into regulatory advantage

# Risks, Mitigations and *Pitfalls*.

*What can go wrong, in decreasing order of probability and severity.*

## Risk 1 · Failing to hire founding talent

**Probability:** high. **Impact:** blocking.

There are perhaps 200 people in the world qualified to lead an FHE operation in production. Almost none work in pharma today — they come from Zama, Duality, Inpher, Tune Insight, Owkin, IBM Research, or academia. Convincing one to enter pharma is difficult because the GxP validation cycle is hostile to the profile.

**Mitigation:** treat it as a strategic acquisition, not a hire. Acquihire of a startup is often the fastest path (see Owkin, bought by Sanofi in 2025; Lifebit; Mozaic). Higher cost, team already formed, significantly lower time-to-capability.

## Risk 2 · CSV validation and GxP environment

**Probability:** medium-high. **Impact:** manageable but time-consuming.

Pharma operates under validated environments (GxP, CSV, 21 CFR Part 11). FHE libraries were not designed for this standard. Validating the stack for GxP use is a non-trivial project.

**Mitigation:** start with non-GxP use (discovery, pre-clinical, retrospective RWE). Validate the stack in parallel with non-critical cases. Migrate to GxP only when the stack is mature and the team comfortable.

## Risk 3 · Computational cost not dropping at the expected pace

**Probability:** medium. **Impact:** manageable.

The FHE cost-reduction curve depends on algorithmic and hardware advances. If the pace slows, cases with very large patient volumes may be temporarily unfeasible.

**Mitigation:** start with low-volume, high-value cases (rare diseases, oncology, biologics). The case in these segments stands on its own and does not depend on further cost reductions.

## Risk 4 · Regulator misunderstands or rejects the architecture

PT EN

**Probability:** low-medium. **Impact:** high.

There is a chance that FDA, EMA or ANVISA (Brazilian health authority) interpret the architecture as insufficiently auditable, or ask for additional guarantees that make the use impractical.

**Mitigation:** engage regulators **before** use, in pre-submission meetings. Present the architecture. Seek a prior opinion. Co-build the audit protocol. Regulators who give favorable opinions become allies, not adversaries.

## Risk 5 · A competitor announces first

**Probability:** medium. **Impact:** high.

Roche, Novartis, Pfizer, Sanofi, Janssen — any of them may have a similar project underway. Unlikely to be at the same stage, but possible.

**Mitigation:** speed. Every month of delay is another month of risk exposure. Consider partnerships with FHE vendors as acceleration rather than building everything from scratch.

## Risk 6 · Vulnerability discovered in the FHE scheme

**Probability:** very low. **Impact:** medium-high.

Modern FHE schemes are based on well-studied lattice problems — the same problems on which NIST standardized the next generation of post-quantum cryptography. Mathematical confidence is high. But unforeseen theoretical advances are always possible.

**Mitigation:** use conservative parameters. Monitor the literature. Have a migration plan between schemes. Do not use FHE as the sole defense layer — combine it with TLS, AES, key segregation, auditing.

## Pitfall 1 · Treating it as an IT project

The most common mistake in adopting transformative technology is placing it under CIO/CTO instead of under CSO/Chief Medical Officer/Chief R&D Officer. Result: perfect technical delivery and zero scientific impact. **FHE in pharma should report to R&D or the CMO, with CEO sponsorship.**

## Pitfall 2 · Starting with the wrong case

PT EN

The temptation is to start with the most visible case — often an ongoing phase III trial. It is a mistake. Critical cases do not tolerate experimentation. Start with discovery, retrospective RWE, or investigator-initiated studies, where the risk is low and learning is fast.

## Pitfall 3 · Forgetting key governance

FHE protects during computation. But key management is where most implementations fail. For pharma, the right structure involves threshold cryptography shared between sponsor, IRB and independent partner — so that no party alone can decrypt. **Key governance design is half the project.**

## Pitfall 4 · Underestimating clinical legal

Clinical pharma legal is its own specialty, distinct from commercial privacy legal. Engaging the wrong people or too late results in massive rework. Legal must be in the project from month 1.

# A letter for the next decade of *data-driven medicine*.

*For the CEOs, board members, Chief Scientific Officers and Chief Medical Officers of the houses that can still choose to lead.*

The industry you lead was built on an old and noble promise: that science applied to disease can reduce human suffering at scale. That discovering a molecule, validating it in rigorous trials, and bringing it to the patient is an act of collective care that deserves the capital, talent and trust society invests in the industry. That the patient who enters a clinical trial is not raw material, but a partner — and that the relationship between sponsor and patient, mediated by doctor and investigator, is a trust relationship in the deepest sense of the word.

This promise crossed a century. It survived ethical tragedies (Thalidomide, Tuskegee, Estonia), it survived technical transformations (from small molecules to biologics, from biologics to gene therapies), it survived the partial commoditization by generics. It survived because it was — and largely still is — true. The patients who enter clinical trials today have a real relationship of respect and expectation with the industry, mediated by doctors who trust the system.

But in the last ten years, without anyone decreeing it, the relationship between industry and patient has changed in nature. The patient stopped being the person who signs a specific consent form for a specific experiment, and became a continuous source of data: record, genome, image, wearable, behavior, geolocation, microbiome. Each of these data points was born with a local justification — a single feature-inclusion decision, a single protocol, a single integration. But the aggregate result is something no individual executive would consciously design: an industrial-scale operation to collect intimate clinical data, mediated by a chain of vendors and providers where the data exists in the clear at dozens of points, and where the only defense is a fragile chain of contractual trust.

It is possible to reverse course. More than that: it is strategically preferable to reverse course. Not because regulation demands it — although it does. Not because an eventual breach is inevitable — although it is, eventually, somewhere. But because the original relationship, based on real respect for the patient, was more valuable, more sustainable, and more aligned with the industry's central mission.

FHE — Fully Homomorphic Encryption — is the first technology in decades that allows reversing course **without losing capabilities**. It is possible to keep doing RWE studies, keep training AI models

on medical imaging, keep running decentralized trials, keep doing collaborative discovery, keep building international registries. **It is possible to do all of this without ever seeing the individual patient**

PT EN

This sentence seems, on a first reading, paradoxical. On a second reading, it seems too technical. On a third, it seems the most obvious argument the industry has ever had at its disposal. How is it possible to research without seeing? The mathematics answers — and the answer is elegant, old, and finally viable in production.

What is at stake is not a technical feature. It is the possibility of an industry becoming again, unambiguously, what it has always said it is: an industry of healing. An industry that respects the patient as a full partner, with mathematically verifiable privacy and not with a management promise. An industry able to collaborate with itself without handing over IP, able to share knowledge without handing over competitiveness, able to fulfill the regulatory contract with mathematical transparency.

This possibility is open, today, to a small number of global houses. The first to understand what is in their hands. The first to make the right hire, call the right lawyers, present the case to the right board. The first to publish the first regulatory submission using the architecture, the first to engage FDA and EMA in formal conversations, the first to sustain the position in interview, in conference, in article, for two consecutive years without yielding to the temptation to dilute the message.

In three years, this position will be taken. Some player will lead — internal or tech entrant. Pharma has always had the technical capacity and capital to lead its own transformations; what was often missing was the courage to move before the problem became a crisis. This is a different window. It is the first window in which the industry can choose to lead before being pressured — because the regulatory problem is still forming, but the technical solution is already ready.

---

*“The question is not whether the industry will change. It is whether we will change leading, or being led.”*

---

There is a window. It is short. It is real. It is historically rare. Entire industries spend decades waiting for windows like this — and most installed houses lose them through excess operational prudence. Pharma has been through others: the transition to biologics (and some houses lost to specialized entrants), the transition to gene therapy (and some lost to startups that became unicorns). This is the next one. It is

possibly the last of this decade with capacity to redefine the central architecture of how the industry researches, collaborates and regulates itself.

PT EN

Whoever reads this eBook holds a map. The map is not complete, not without risks, not without costs. But it is clear. And it is, at this moment, in front of the right people to make the right decision.

The rest is courage.

— *End of Volume I*

---

# Executive *Glossary*.

*The terms you will hear from the CSO/CIO. In plain English.*

## **FHE — FULLY HOMOMORPHIC ENCRYPTION**

Encryption that allows computations to be executed on encrypted data without decrypting it. The result, when decrypted, is equal to what would be obtained on the original data.

## **RLWE — RING LEARNING WITH ERRORS**

The mathematical problem on which most modern FHE schemes are based. It is the same problem as the post-quantum cryptography standardized by NIST (ML-KEM, ML-DSA). Resistant to classical and quantum computers.

## **CKKS, BFV, BGV, TFHE**

The four main FHE schemes in practical use. CKKS for ML and statistics; BFV/BGV for exact integers; TFHE for flexible boolean logic.

## **RWE — REAL-WORLD EVIDENCE**

Clinical evidence derived from real-use data (claims, records, registries) instead of controlled trials. FDA, EMA and ANVISA (Brazilian health authority) accept RWE as regulatory support when the quality is sufficient.

## **EHDS — EUROPEAN HEALTH DATA SPACE**

European initiative creating common infrastructure for secondary use of clinical data in research. In force in phases through 2027. Will require technical privacy-preserving analytics capability as a prerequisite.

## **ICH-E6(R3)**

Revision of international Good Clinical Practice published in 2023. Marks the turn toward requiring verifiable data governance, not merely declared.

## **CSV — COMPUTER SYSTEM VALIDATION**

Validation process required by the FDA (21 CFR Part 11) and equivalents for computer systems used in GxP operations. FHE in clinical environments must pass CSV.

PT EN

## **GXP**

Umbrella term for Good Manufacturing/Clinical/Laboratory Practice — regulatory standards governing all pharmaceutical operations. Systems used in GxP have additional requirements for validation, auditing and traceability.

## **DSMB — DATA SAFETY MONITORING BOARD**

Independent committee that monitors safety and efficacy in clinical trials. In trials under FHE, the DSMB has a central role in key management for interim analyses.

## **THRESHOLD CRYPTOGRAPHY**

Technique that distributes a cryptographic key among multiple parties, requiring a quorum to use it. Essential for FHE architecture in pharma — no single party alone can decrypt.

## **PSI — PRIVATE SET INTERSECTION**

Protocol derived from FHE/MPC that allows two parties to discover the intersection of their sets without revealing the rest. Useful to identify patients in common between sites without exposing databases.

## **FEDERATED LEARNING**

Distributed training where data stays local. When combined with FHE for gradient aggregation (FL+FHE), it eliminates the leakage that pure FL has.

## **LATTIGO, OPENFHE, CONCRETE**

Main FHE libraries in practical use. Lattigo (Go, Tune Insight); OpenFHE (C++, Duality); Concrete (Rust+Python, Zama).

# Vendors and *Partners*.

*Where commercial capability is, where the talent is, and where there are already active use cases in pharma.*

## Specialized vendors and startups

VENDOR	HQ	FOCUS
<b>Owkin</b>	Paris / NY	FL+FHE for clinical research; partially acquired by Sanofi
<b>Zama</b>	Paris	TFHE, Concrete framework, focus on developer experience
<b>Duality Technologies</b>	USA / Israel	OpenFHE, focus on healthcare and finance, heavy consulting
<b>Inpher</b>	Switzerland / USA	Hybrid FHE + MPC, focus on finance and healthcare
<b>Tune Insight</b>	Switzerland (EPFL)	Lattigo, focus on federated medical research
<b>Lifebit</b>	United Kingdom	Federated genomics platform for pharma
<b>Mozaic / TripleBlind</b>	USA	Privacy-preserving analytics for pharma and payers
<b>ConcertAI</b>	USA	RWE with advanced governance for oncology

# Relevant initiatives and consortia

PT EN

- **MELLODDY** — European consortium of 10 pharmas for collaborative discovery via federated learning (completed; FHE successor under discussion)
- **Trustworthy Federated Data Analytics (TFDA)** — Helmholtz, Germany — privacy-preserving federated clinical research
- **European Health Data Space (EHDS)** — European infrastructure for secondary use of clinical data
- **NIH N3C** — National COVID Cohort Collaborative (applicable governance model)
- **Global Alliance for Genomics and Health (GA4GH)** — standards for genomic data, including private computation

## Academic centers where the talent lives

- **EPFL (Lausanne)** — Christian Cachin, Jean-Pierre Hubaux, origin of Lattigo
- **IBM Research** — historical origin of FHE (Craig Gentry, 2009)
- **Microsoft Research** — SEAL library, Kim Laine
- **NJIT / Duality** — Yuriy Polyakov, Kurt Rohloff
- **Université Paris-Saclay / ENS** — Léo Ducas and the French ecosystem
- **KU Leuven (COSIC)** — Frederik Vercauteren
- **Stanford / MIT / CMU** — several researchers active in clinical ML+FHE

# 30 Questions for the CSO/CIO.

*The list you should take to your next meeting with your scientific and technology team.*

## Strategy and capability

1. Do we have anyone on the team with a background in advanced cryptography? If not, what is the plan?
2. What is our current exposure to clinical data processing, by volume and by jurisdiction?
3. How many ongoing trials process data in environments where the sponsor can access individual patient data?
4. Do we have an updated inventory of which data leaves the investigator site and where it is processed?
5. What is our current legal opinion on the sustainability of this architecture over the next 36 months?

## Priority use cases

6. What is our highest-potential-value RWE program — and what is its regulatory bottleneck?
7. Which ongoing trials could have a more ambitious decentralized design if privacy were solved?
8. In which therapeutic areas is the cohort so small that international cross-referencing would be transformational?
9. At what point in discovery do our chemical libraries leave a fully controlled environment?
10. Do we have libraries worth sharing with competitors in a consortium if privacy were mathematical?

## Technical architecture

11. Which FHE scheme makes most sense for our first case? CKKS, BFV, TFHE?
12. What is the computational depth of our most valuable analytical pipeline?

13. Can we run a basic statistical analysis (means, variances, regression) in FHE with acceptable latency today?
14. What would be the estimated computational overhead for our case?
15. How do we manage keys between sponsor, IRB, DSMB and partners?
16. Who holds the key in the case of an international DSMB?
17. Is our architecture compatible with threshold cryptography?

## Cost and infrastructure

18. Estimated computational cost per analysis, FHE versus plaintext?
19. Which accelerators have we evaluated? FPGA? Specific GPU? AVX-512?
20. Will we build internally or use a vendor (Owkin, Zama, Duality)?
21. What are the expected capex and opex over the next 24 months?
22. How do we integrate with our current eCRF, EDC and CTMS stack?

## Regulation and validation

23. Do security parameters meet 128 bits and are they post-quantum?
24. Who will perform the independent security audit?
25. How do we demonstrate compliance with HIPAA, GDPR, LGPD (Brazilian data protection law), ICH-E6(R3)?
26. Can we publish the architecture for public and regulatory review?
27. Do we already have an open dialogue with FDA, EMA or ANVISA (Brazilian health authority) about the use?
28. How do we pass CSV/21 CFR Part 11 for GxP use?

## Risks

29. Worst-case scenario if a competitor announces before us?
30. Worst-case scenario if the regulator asks for technical guarantees we cannot yet deliver?



## The Invisible Trial

Strategic eBook for the executive leadership of the global pharmaceutical industry.

Volume I · 2026 Edition · Confidential distribution.

Set in lowan Old Style and SF Pro.

Built as a self-contained HTML document.

Print on heavy-weight paper for fidelity to the original layout.

— end —