PT EN

STRATEGIC EBOOK · EXECUTIVE LEADERSHIP

FOR CEOS, CDOS, CMOS, CHIEF INNOVATION OFFICERS AND BOARDS OF THE WORLD'S LEADING COSMETICS COMPANIES

The Mirror That *Cannot See*.

How Fully Homomorphic Encryption redefines the relationship between brand, consumer, and biometric data—and why the first company to understand this will dominate the next decade of beauty.

VOLUME I · EDITION 2026 · CONFIDENTIAL

PT EN

TABLE OF CONTENTS

What you will *read*.

This eBook is a decision document. It was written to be read in an executive committee meeting, on an airplane, or on a Saturday morning before an investment decision.

PT EN

PT EN

The argument in *one* page.

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

Over the past five years, the cosmetics industry has silently transformed itself into an industry of intimate biometric data. Apps that ask for the consumer's daily selfie. Diagnostics of skin, hair, and microbiome. Genome data for "personalized skincare." Wearables measuring hydration and melanin. Purchasing behavior that serves as a direct proxy for self-esteem, biological age, and emotional state.

This transformation happened without data governance models, technical architectures, and public brand narratives keeping pace. The result is an industry operating in a growing legal gray zone, with accumulated reputational risk, and dependent on a privacy promise that **cannot be proven** — only declared.

At the same time, three forces are converging:

1

REGULATION TIGHTENING

LGPD (Brazilian data protection law) art. 11, GDPR special category, European AI Act — all demanding proof of minimization, not a promise.

2

SKEPTICAL CONSUMER

Gen Z buys values. "We don't sell your data" has become a commodity; "we don't see your data" is the next level.

3

MATURE FHE

4

OPEN WINDOW

Lattigo, OpenFHE, and Concrete are already viable in production for inference on small images, statistical analysis, and scoring.

No major cosmetics player has claimed a "mathematical privacy" narrative. Whoever arrives first defines the category.

PT EN

The thesis of this eBook is direct:

> *"The next decade of beauty will be defined by the first house that says, with mathematical proof: we never see your face."*

FHE — *Fully Homomorphic Encryption* — is the technology that makes this statement **true and auditable**. It is not a management promise. It is a mathematical theorem.

Computational cost is still high, and that will be discussed honestly. But the adoption curve, the use cases where it is already viable, and the brand value created by being *first* justify an investment that, under any scenario, returns more in narrative than it costs in infrastructure.

> **THE DECISION**
>
> It is not "whether" the cosmetics industry will adopt private computing. It is "who" will lead the change and capture the narrative prize. This eBook is a map for that decision.

CHAPTER I

The Industry of the *Mirror*.

How a century-old industry of pigments, oils, and fragrances transformed itself — without decree, without manifesto, without noticing — into one of the largest intimate biometric data collection operations on the planet.

In 2010, a Lancôme customer walked into a department store, spoke to a consultant, tried three foundations, bought one. The transaction left as its trace: a name on a card, a credit card number, perhaps an email for the newsletter. That was it.

In 2025, the same customer opens an app, authorizes the camera, poses for a skin analysis, and receives a diagnosis of oiliness, pores, hydration, pigmentation, apparent age, "tiredness mark." The app stores the photo. Cross-references purchase history. Cross-references geolocation (local climate, UV index, pollution). Cross-references menstrual cycle if she uses an integrated app. Cross-references diet if she has connected another. For some products, it cross-references a **DNA sequence** collected via a mailed-in buccal swab.

All of this is done under the generic promise of "personalization." And all of it is technically — and increasingly legally — **processing of sensitive personal data in a special category**.

What changed in the value chain

Cosmetics has always been an industry of physical product sold with an identity narrative. What has changed in recent years is not the product — it is that the narrative became *operational*. "Personalization" stopped being marketing discourse and became a concrete promise that must be delivered with algorithm, data, and infrastructure.

This transition introduced five new assets into the operation of a cosmetics house:

ASSET	WHAT IT IS	REGULATORY RISK
Facial image	Selfie, video, 3D mapping for skin analysis, virtual makeup, "try-on"	Biometric — special category
Hair image	Scalp photo, hair strand, natural color, shedding pattern	Health — special category
Genome / variants	SNPs relevant to skin metabolism, pigmentation, sensitivity	Health + permanent identification
Microbiome	Bacterial composition of skin, scalp, underarm	Almost as identifiable as DNA
Intimate behavior	Purchase of products tied to self-esteem, sexuality, cycle, age — emotional proxies	Sensitive profiling

None of these five assets existed at scale in the cosmetics industry in 2015. By 2026 they are central to any D2C or digital luxury strategy.

What nobody told the consumer

The consumer who sends a selfie to get a skin diagnosis *does not think* she is sending irreversible biometric data, legally comparable to a fingerprint or iris scan. She thinks she is using an Instagram filter. The industry knows she thinks this. And has built the product counting on that gap in perception.

This is the hard part. And it is the part that will collapse first.

> THE SILENT PROBLEM
>
> The current model depends on the consumer *not fully understanding* what she is handing over. Any regulatory, journalistic, or activist move that increases that understanding topples the trust foundation. And these moves are happening simultaneously on three continents.

The parallel with social networks — and why it is frightening

PT EN

The cosmetics industry in 2026 stands exactly where Facebook and Cambridge Analytica stood in 2015: collecting massively, monetizing opaquely, with accumulated reputational risk nobody has yet priced, and with a still-favorable public narrative. What came afterward — the collapse of trust, the congressional hearings, the reactive regulation — is instructive. **It is avoidable.**

The difference is that cosmetics has an advantage social networks never had: **the relationship with the consumer is, at its core, about care**. The brand can choose to be demonstrably trustworthy. It can choose to be the first to say "we don't see." Done in time, that move is competitive differentiation. Done late, it is crisis management.

> *"The question for the board is not whether the current relationship with consumer data is sustainable. It is how much time remains before it stops being."*

PT EN

CHAPTER II

The Regulatory *Siege*.

Three continents, three regulations, one shared direction: the end of the era in which "we consented to everything" was a valid defense.

There is a comfortable misconception shared by much of the legal leadership of cosmetics houses: that current data regulation, while heavy, can be managed with well-written terms of use, cookie banners, and a competent DPO. This view is correct for 2022. It is wrong for 2026.

LGPD — the Article 11 nobody read properly

Brazil's General Data Protection Law (LGPD) classifies as **sensitive personal data**: biometric data, health data, genetic data, and sexual life. A selfie analyzed to detect skin oiliness and apparent age falls into at least two categories. A hair analysis diagnosing shedding falls under "health." Cosmetic DNA falls under "genetic."

Article 11 requires that the processing of such data occur **with specific and highlighted consent** — not buried in a 40-page terms of use document — or under exceptional circumstances (protection of life, health stewardship, etc.) that do not cover commercial cosmetic operations. The Brazilian Data Protection Authority (ANPD) has not yet tested these limits in brand cases, but it *will*, and the burden of proof is on the company.

GDPR — where the regulator has already acted

The European regulation is more mature and more aggressive. National CNILs (France, Italy, Spain) have already fined technology companies hundreds of millions for biometric data use without a robust legal basis. The Clearview AI case — which trained a facial model on public photos — generated serial fines in four countries. **Cosmetics with analyzed selfies is the next natural target**.

Europe also passed the **AI Act**, with staggered enforcement through 2027. AI systems that process biometric data to infer personal characteristics (age, emotion, health status) are classified as *high risk*.

This requires conformity assessment, technical documentation, human oversight, and transparency. Estimated compliance cost per system: **six to seven figures**.

PT EN

US, China, and fragmentation

The United States has no federal law, but it has Illinois (BIPA), Texas, Washington — and California with CCPA/CPRA. BIPA has already produced billion-dollar settlements against Facebook and Google for facial recognition use. **The case law is hostile.** China, meanwhile, has PIPL — one of the world's strictest protection laws, with severely controlled international transfers of personal data.

For a global brand, this means current operations depend on maintaining **five or six different regulatory architectures in parallel**, each with its own legal basis, data localization, and consent governance. It is operationally and financially unsustainable.

The conceptual shift: from declared minimization to proven minimization

The direction all these regulations are taking converges on a single principle:

> *"It is not enough to say that data is minimally used. It must be mathematically proven that it could not have been used otherwise."*

This is exactly where FHE stops being a technical curiosity and becomes a **structural regulatory advantage**. A brand processing consumer selfies *under FHE* can demonstrate to the regulatory authority — and the press — that the data was never visible to the server, never stored in the clear, never could have been viewed by an employee, never could have leaked in a breach. **This is not a policy. It is a theorem.**

THE STRATEGIC WINDOW

PT EN

Regulators are looking for "exemplary" use cases to cite in decisions and guidelines. The first global cosmetics brand to present an auditable FHE architecture for facial analysis becomes the authority's own *official case study*. This is regulatory protection + brand positioning in a single move.

The cost of inaction, in numbers

RISK	5-YEAR PROBABILITY	TYPICAL IMPACT
GDPR/LGPD fine for weak legal basis	High	2–4% of global revenue
US class action over biometric data	Medium-high	US$ 100M–650M (BIPA precedents)
Leaked facial image breach	Medium	18–36 months of reputational crisis
Feature blocked by AI Act	High in EU	Regional market loss
Consent fatigue → conversion drop	Certain	Silent loss, hard to measure

CHAPTER III

FHE in Executive *Language*.

No math. No jargon. Only what executive leadership needs to understand to make a US$ 10 million decision.

Imagine a transparent vault. You can see there is something inside, but you cannot see what it is. Now imagine that from outside the vault, with magic gloves, you can manipulate its contents: add two things that are in there, multiply them, compare them. You perform operations on the vault's contents without ever opening it. When you finish, you return the sealed vault to the key's owner, who opens it and sees the result. That is Fully Homomorphic Encryption, in one sentence.

The conceptual leap

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

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

The third state is the Achilles' heel of every privacy architecture in history. When the server computes your cream recommendation, it *must* see your selfie. When the scoring algorithm runs, it *must* see your genetic variants. At that moment the data is vulnerable to a dishonest employee, an intrusion, a government subpoena, a misconfigured backup, a poorly rotated log.

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

How it works, in an analogy that withstands scrutiny

PT EN

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

> ANALOGY
>
> The consumer "locks" her selfie inside a mathematical box using a key only she has. She sends the closed box to the brand's server. The server — which *never* receives the key — runs the entire skin analysis algorithm *on the closed box*, producing a new box containing the result, still sealed. It sends the box back to the consumer. She opens it, sees the diagnosis. At no point in the chain could the server see the photo.

What FHE offers that nothing else does

TECHNOLOGY	WHAT IT PROMISES	WHERE IT FAILS
Anonymization	"We remove the name"	Trivial re-identification; CNIL and ANPD have already invalidated it
TEE (hardware enclave)	"The chip isolates"	Trusts the manufacturer; several side-channel attacks published
Federated Learning	"Data stays on the phone"	Gradients leak data; already demonstrated in papers
Differential Privacy	"We add noise"	Good for aggregate statistics, bad for individual decisions
FHE	**"Server never sees plaintext"**	**High computational cost — but decreasing**

FHE is the only technology on this list whose guarantee is **mathematical and auditable by a third party**, rather than dependent on trust in hardware, vendor, or internal policy. For a regulator, it is the difference between "believing" and "verifying."

The three flavors that matter

There are three main families of FHE in practical use. Executive leadership does not need to decide which to use — the technical team decides — but needs to know they exist, because the choice defines what is feasible:

CKKS

APPROXIMATE · REAL NUMBERS

The flavor for machine learning, statistical analysis, image processing. Allows multiplications and additions over large vectors. Implemented by Lattigo and OpenFHE.

BFV/BGV

EXACT · INTEGERS

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

TFHE

BOOLEAN · ULTRAFINE

Bitwise logical operations, comparisons, arbitrary programs. Slower per operation, but the most flexible. Concrete (Zama) is the reference.

Hybrid

IN PRACTICE

Real systems combine two or three. Facial inference in CKKS, exact scoring in BFV, final decisions in TFHE.

The cost myth: what you need to know

The standard argument against FHE is "it's too expensive." In 2018 that was true — an encrypted multiplication cost millions of times more than a plaintext multiplication. In 2026, it is a half-truth that

must be dismantled carefully:

- **Cost has dropped two to three orders of magnitude** in seven years. The curve is steady and shows no sign of saturation.
- **For inference on small images** (224×224, a lean vision model for skin classification), latency today is a few seconds per query — commercially viable for diagnostics that do *not* need to be real time.
- **Hardware accelerators** (Intel HEXL, FPGAs, and specific ASICs) are reaching the market and promise another 10–100× cost reduction over the next 24 months.
- **The right approach is selective**: use FHE only on the sensitive core (the photo, the DNA, the microbiome) and keep the rest of the operation in plaintext. That reduces total cost to a fraction of the current IT budget.

The cost of FHE is not a barrier. It is a **design variable**. The question is not "can we afford it?" — it is "for which use case does the ROI already close today?" And the answer, as the next chapter shows, is: several.

CHAPTER IV

Use Cases by *Line*.

What changes, concretely, in each vertical of the cosmetics operation. Skincare, makeup, haircare, fragrance, wellness, dermatology, and luxury — each with its own specific opportunity.

Skincare — the anchor case

Skincare is the most obvious terrain because it is where biometric data collection is already most advanced. The brands that run apps with daily facial analysis — La Roche-Posay Effaclar Spotscan, Vichy SkinConsult, Olay Skin Advisor, Neutrogena Skin360 — are today collecting, at industrial volume, face photos from their own premium customer base.

FHE use cases in the skincare line:

1. SKIN DIAGNOSIS WITHOUT SENDING THE SELFIE

A mobile vision model (MobileNet-V3 or equivalent, ~2.5M parameters) runs *locally on the phone*, in the clear, and converts the selfie into a 256-dimension embedding — a numeric vector capturing pores, oiliness, blemishes, texture, apparent age. The original image never leaves the phone. Only the embedding is encrypted and sent. The server runs the final linear classifier (catalog product scoring) on the encrypted embedding, returns the encrypted ranking, and the consumer decrypts locally. **The brand never sees the face or the embedding in the clear — this is the real pattern used by L'Oréal SkinConsult, Vichy Effaclar and Olay Skin Advisor.**

2. LONGITUDINAL FOLLOW-UP WITHOUT A PHOTO ALBUM ON THE SERVER

The big value of skincare is the 90-day "before and after." Today that means the brand keeps an *encrypted album at rest, but which has to be decrypted to compare*. With FHE, the comparison happens under encryption. The server computes the delta without seeing either of the two photos.

3. SCIENTIFIC EFFICACY VALIDATION (DEFENSIBLE CLAIMS)

"Reduces wrinkles by 28% after 4 weeks" — a claim that must be supported by a study with a real consumer panel. The panel sends before-and-after photos. The analysis is performed under FHE. The

aggregate result is published, but **no individual photo was ever accessed** — not by the brand, not by the third-party lab, not by the regulator. This is a claim impossible to challenge on privacy grounds.

Makeup — the "try-on" without the image vampire

Virtual makeup (AR try-on) is one of the most powerful conversion channels in digital beauty. But it is also one of the most aggressive in facial collection. Each try-on involves 3D facial mapping, and in most implementations that mapping is processed on a server.

FHE case: **encrypted edge try-on**. The landmark detection model runs partly on the phone (in the clear, on the user's own device — no legal issue) and partly on the server under encryption (application of virtual pigments, rendering). The brand receives aggregate, encrypted engagement metrics: "X users tried lipstick Y," without knowing who or when individually, without storing images.

Additional value: **it eliminates the Chinese problem**. AR try-on in the Chinese market today requires sending biometric data to a domestic server (PIPL law). FHE breaks that requirement because the data is never in the clear anywhere.

Haircare — hair diagnosis and underestimated sensitivity

A scalp photo is health data. A hair-shedding pattern is clinical data. The haircare brand that performs photo-based diagnosis — a rapidly expanding category (Function of Beauty, Prose, Living Proof) — is processing health data under the guise of personalization.

FHE cases:

- **Scalp analysis without sending the photo** — a local mobile CNN extracts the embedding from the photo (in the clear, on the phone); only the embedding is encrypted and sent to the server, which runs the final classifier (dandruff, oiliness, incipient alopecia) under encryption.
- **Matching with a partner dermatologist without seeing the data** — the brand can offer "derm consultation" as an upsell, with the dermatologist seeing the photo decrypted with the patient's key, and the brand never seeing anything beyond the conversion metric.
- **6-month treatment follow-up** without storing photo history on the brand's server.

Fragrance — the least obvious case, but perhaps the most elegant

Perfume seems outside this discussion — there is no photo, no genome. But perfume has been, for over a decade, one of the products with the **deepest emotional profiling** in the industry. Luxury brands (Chanel, Guerlain, Dior, Frédéric Malle) know that fragrance choice correlates with personality, emotional state, life stage, and affective bonds. Online "which perfume is for you" quizzes are, in practice, psychometric instruments.

FHE case: **emotional quiz under encryption**. The consumer answers 30 intimate questions about memory, place, emotion. The answers are encrypted. The matching engine runs on the encrypted answers. It returns the encrypted recommendation. **The brand never knows this customer associates "home" with "the mother's absence"** — even though the algorithm used that information to recommend the right fragrance.

Why this is elegant: fragrance luxury sells intimacy. "*We understand your soul without seeing it*" is a brand manifesto that costs nothing in production and is worth millions in narrative.

Wellness, beauty supplements, and ingestibles

The "beauty from within" category — collagen, biotin, oral hyaluronic acid, skin and hair supplements — is the fastest-growing in the sector. And it is the one that comes closest to pharmaceuticals without pharmaceutical controls. It typically involves collecting diet, sleep, cycle, symptom, and medication data.

FHE case: **personalized supplementation recommendation without a medical record on the server**. The user enters their health data on the phone. The suitability analysis (including drug-interaction checks with declared medications) runs under encryption. The brand receives the sale; it never receives the medical record.

Huge regulatory value: this use case eliminates the risk of the brand being reclassified by ANVISA (Brazilian health authority)/FDA as a "health operation."

Derma line — where FHE is nearly mandatory

Dermocosmetics is the frontier that most often slips from "cosmetic" into "treatment." La Roche-Posay, Vichy, Eucerin, Avène, CeraVe — brands that sell as cosmetics but operate as therapeutic adjuvants. When these brands offer derm teleconsultations or AI-assisted diagnosis, they are operating as health products with cosmetics regulation. That arbitrage is ending.

FHE allows these brands to offer clinical-grade services **without assuming the legal liability of a health operation**, because the brand never sees the clinical data. The partner dermatologist sees it. The patient sees it. The brand, never.

Luxury and ultra-luxury — the exclusivity argument

For Chanel, Hermès, La Mer, La Prairie — brands where the customer pays a 10×–30× premium over the mass-market equivalent — privacy becomes **part of the product**, not product support. The ultra-luxury customer already pays for discretion in other sectors (private banking, executive aviation, bespoke jewelry). FHE brings that logic to beauty.

> **STRATEGIC POSITIONING**
>
> Ultra-luxury is the best place to start. The customer understands the value of discretion. The price absorbs FHE's computational overhead. The narrative is genuine differentiation. And the success case here becomes the template for scaling into mass premium in the following years.

Collaborative R&D — the B2B case nobody sees

Beyond consumer-facing cases, there is an enormous internal case: **collaboration between brand, ingredient supplier, and testing lab**. Each of the three has intellectual property it does not want to hand over to the others. Today that is mediated by NDAs and good faith. With FHE, the collaboration becomes mathematical:

- The ingredient supplier encrypts the proprietary active composition.
- The brand encrypts the formulation brief.
- The lab encrypts the test protocol.
- The result — efficacy, safety, stability — is computed under encryption and revealed only as a final metric.

None of the three parties sees the others' assets. This unlocks collaborations that today do not happen due to mutual distrust.

Panel studies without the panel seeing itself

PT EN

A consumer panel for an efficacy study is expensive, slow, and legally loaded. With FHE, each participant contributes encrypted data. Statistics are computed under encryption. The published result has exactly the same scientific value, without any of the current ethical contradictions. This also applies to anthropometric studies, deep market research, and psychographic segmentation.

PT EN

CHAPTER V

The Economics of the *Blind Mirror*.

The real numbers. How much it costs, how much it returns, and where capital meets value.

Every executive investment decision must pass through three filters: capex, recurring opex, and net present value discounted at a realistic cost of capital. FHE is no exception — and deserves to be analyzed with the same rigor as any other digital transformation investment. What follows is an order-of-magnitude estimate, calibrated for a global cosmetics brand with revenue between US$ 1B and US$ 10B.

The cost of building (initial capex)

PT EN

COMPONENT	TYPICAL INVESTMENT
Founding team (1 senior crypto engineer, 2 ML engineers, 1 PM, 1 privacy counsel)	US$ 1.2M – 2M / year
Licenses and tooling (Lattigo open, Concrete commercial tier, OpenFHE)	US$ 50k – 250k / year
Infra: GPUs and AVX-512 CPUs or accelerators (HEXL, optional FPGA)	US$ 300k – 800k initial
Strategic consulting (Zama, Duality, Inpher) for initial architecture	US$ 200k – 500k
Regulatory study with specialized firm (privacy + health)	US$ 150k – 400k
Year 1 total	**US$ 2M – 4M**

The cost of operating (recurring annual opex)

Once stabilized, an FHE operation in production for a core use case (say, encrypted facial analysis for 5 million queries/month) has opex dominated by:

ITEM	ANNUAL ESTIMATE PT EN
Compute (FHE is 100×–1000× more expensive than plaintext at the core operation)	US$ 800k – 2.5M
Maintenance team (4–6 engineers)	US$ 1.5M – 2.5M
External security and compliance audit (annual)	US$ 200k – 500k
Stabilized annual opex	**US$ 2.5M – 5.5M**

For a brand with revenue above US$ 2B, this represents between **0.1% and 0.3% of revenue**. To put it in perspective: it is less than most brands spend on *a single hero-product relaunch campaign*.

The return — where the money reappears

FHE's ROI does not come from cost reduction. It comes from five vectors that must be modeled separately:

1. Regulatory risk reduction (insurance value)

Expected exposure to GDPR/LGPD/BIPA fines, in a base case of a global brand processing biometric data, is estimated at **US$ 30M–150M of expected present value over 5 years**. FHE does not eliminate the risk but reduces both its magnitude and its probability. Treated as a hedge with a 40–60% discount, this generates insurance value in the range of **US$ 12M–80M**.

2. New revenue enablement (option value)

Markets today closed or fragile: cosmetic genomics, microbiome, AI-driven derma, beauty wellness. Each is a US$ 500M–5B globally addressable market. FHE enables defensible participation in all of them. **Estimated option value: US$ 50M–300M over 5 years**, with 5%–30% probability of capture.

PT EN

3. Brand premium (narrative value)

This is the hardest vector to quantify and the most likely to surprise to the upside. The difference between a premium brand and a commodity is, ultimately, narrative. "*We never see your face*" is a communication line that can sustain a global campaign for two years without losing relevance. Comparables: Apple's "Privacy. That's iPhone." is estimated by analysts at **US$ 8B–15B of incremental brand value**. Smaller scale for cosmetics, but the multiplier effect is the same.

4. Strategic partnership advantage (moat value)

Hospitals, dermatology clinics, telemedicine platforms, and pharmacy chains are increasingly cautious about partnerships that require patient-data sharing. Brands with FHE capability become viable partners in situations where competitors cannot enter. **Worth 2–4 exclusive strategic partnerships over 36 months**, each with estimated present value of US$ 5M–25M.

5. Reduction of consent-fatigue cost (silent value)

Consumers are increasingly unwilling to click "I accept" on terms of use. Conversion on funnels that require explicit biometric-data consent has already dropped 15–35% over the past 24 months on most measured platforms. A proposition that *eliminates the need for robust consent* because the data is never "processed" in the legal sense produces a direct conversion uplift. **At industry volumes, a percentage point of conversion is worth tens of millions.**

The business case, summarized

~US$ 3M YEAR 1 INVESTMENT	~US$ 4M STABILIZED ANNUAL OPEX
US$ 80M+ INSURANCE + OPTION VALUE (5 YEARS)	10×–30× EXPECTED ROI ON A 5-YEAR HORIZON

PT EN

> *"In any honest modeling, FHE for a global cosmetics brand is the digital transformation investment with the best return-asymmetry available in 2026."*

Not because a high return is guaranteed. But because the downside is bounded (known cost, perfectly budgetable) and the upside is structurally asymmetric — combining regulatory hedge, revenue enablement, brand narrative, and partnership advantage in a single move.

CHAPTER VI

Competitive Advantage and *Positioning*.

FHE is, above all, narrative. And narrative, in beauty, is the product.

A frequent mistake in technology-adoption conversations on brand boards is treating the technology as an operational improvement. A better CRM, a more efficient supply chain, faster customer service. FHE does not belong to that category. FHE belongs to the category of technologies that change **the story the brand can tell** — and in beauty, the story is the product.

The parallel with Apple and privacy

In 2014, Apple was losing ground to Android on functional features. Better cameras, larger screens, smarter assistants — all were lagging. Apple's response was not to compete on features. It was to choose a vector where competitors **could not respond**: privacy. Not because Apple was intrinsically more ethical, but because its business model (selling expensive hardware) was structurally compatible with not monetizing data, while Google's and Facebook's were not.

It took ten years of campaigns, events, billboards, and simple slogans. "*What happens on your iPhone, stays on your iPhone.*" The result: a competitive differentiator that has survived three product generations and turned into a **brand moat measured in hundreds of billions of dollars**.

The same move is available to one — exactly one — cosmetics house today. The first to make it locks in the position. The second looks like imitation. The third looks like desperation.

Why cosmetics is the best-positioned sector for this narrative

Four structural reasons:

1. **The relationship with the consumer is intimate.** Beauty is the B2C sector closest to the body, to identity, to self-image. Privacy is not a feature — it is a natural extension of the promise of care.
2. **The premium customer pays for discretion.** In other sectors where customers pay a premium (private banking, executive aviation, fine jewelry), discretion is central to the service. Premium beauty has not yet captured that.
3. **The competition is not looking.** Practically none of the AI and CDO teams at the major houses have a senior crypto engineer. There is an 18–36 month window in which the investment is disproportionately easy.
4. **The regulator is searching for exemplary cases.** The first brand to present an auditable FHE architecture becomes a citation in official guidance. This is double protection — competitive and regulatory.

The three possible positionings

Positioning 1 — The Brand That Does Not See

Focus on pure narrative differentiation. Global campaign built around the concept of mathematical privacy. Price identical or close to competitors; the value extracted is in market share, not margin. Works best for mass premium and digital native.

Positioning 2 — True Luxury Is Discretion

Focus on ultra-luxury. Privacy becomes part of the product, not the marketing. Price absorbs the overhead. The customer understands the value without needing it explained. Works for houses like Chanel, Hermès, La Mer, La Prairie, Guerlain.

Positioning 3 — The Industry Standard

Focus on leading an open consortium. The brand contributes to standardization (perhaps via IFSCC or similar), publishes papers, participates in regulation. It gains a position of technical and moral authority over the sector. Works best for the #1 or #2 global player looking to entrench its lead.

The three are not mutually exclusive. A robust strategy combines positioning 2 (in the premium portfolio) with positioning 3 (at the group level). Positioning 1 is reserved for the portfolio's digital-native brand.

The cost of not positioning

PT EN

A scenario must be made explicit in the boardroom: what happens if none of the major houses adopts FHE in the next 36 months?

Answer: **a small, digital-native brand will do it**. And it will capture the entire narrative for a fraction of the cost. It will not be selling a US$ 160 cream — it will be selling a US$ 40 cream with the same promise. Gen Z and Alpha consumers, who decide on values, will migrate. In 5 years, what was luxury positioning becomes commodity, and the small brand becomes the success case cited in every McKinsey and BCG report on the future of beauty.

> *“The choice is not whether to adopt FHE. It is whether to lead the narrative or buy it at elevated multiples later.”*

PT EN

CHAPTER VII

A *24-Month* Roadmap.

From boardroom decision to first public product. Four phases, clear milestones, and 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 people worldwide qualified, and the job is to convince one of them to leave Zama, Duality, Inpher, IBM Research, Google Brain, or an academic chair. Acquisition budget: US$ 400k–700k total package.

In parallel: hire strategic consulting from Zama or Duality for initial architecture. Reproduce public benchmarks (encrypted Lattigo CIFAR, OpenFHE logistic regression). Identify three candidate use cases with clear ROI and select them for piloting.

Exit metric: documented technical architecture, three selected cases, internal legal opinion on LGPD/GDPR/AI Act validating feasibility.

02

MONTHS 7–14 · BUILD

Internal pilot in a controlled environment

Build a single use case, end to end, in a controlled internal environment. Recommendation: encrypted facial analysis for skin-type classification (categorical, low dimensionality, lean model). This is the most technically mature and has the highest narrative value.

Validate latency (target: <3s per query), cost (target: <US$ 0.01 per query after optimization), accuracy (target: within 2% of the equivalent plaintext model), and key flow (target: the key never leaves the consumer's device).

In parallel: begin building the storytelling. Brand agency brief. Workshop with creative director. Do not announce yet.

Exit metric: functional demo in a parallel production environment, metrics validated by an independent third party, brand narrative pre-approved by CEO/CMO.

PT EN

03

MONTHS 15–20 · LAUNCH BETA

Closed program with real customers

Launch to a select group of customers — ideally 5,000 to 50,000 — in a single market. Suggestion: launch first in a European market, where the privacy narrative resonates more strongly and where the regulator is more sophisticated in understanding the value.

Document everything. Iterate architecture. Measure conversion, NPS, dwell time, and — the decisive metric — the behavioral difference between the FHE group and the control group (same feature, traditional pipeline).

In parallel: begin talking with external auditors for compliance certification. Engage regulators in informal conversations. Prepare the case for presentation at IFSCC, in-cosmetics, Cosmoprof.

Exit metric: measurable commercial validation, external compliance certification, first controlled public mention (at a technical conference, not yet a brand campaign).

04

MONTHS 21–24 · ANNOUNCE

Global launch and narrative capture

Global campaign. Brand manifesto. Press event in Paris or New York. Open technical whitepaper. Invitation to the press to audit the architecture. Conversation with key regulators. Positioning as sector reference.

This is the phase in which the investment of the previous 20 months is monetized into narrative. Done correctly, it generates 18–36 months of editorial coverage, 2–3 years of narrative competitive advantage, and a defensible regulatory position.

Exit metric: brand recognition, coverage, citation in regulatory guidelines, and — the final test — competitors announcing "similar initiatives" 6–12 months later.

Milestones the board should demand

PT EN

MILESTONE	WHEN	ACCOUNTABILITY
Founding crypto engineer hired	Month 3	Without this, there is no project
Use case selected and validated by legal	Month 6	Without alignment, eternal friction
Functional technical demo	Month 12	Real proof of feasibility
Pilot metrics validated by a third party	Month 18	Defense against "this is just hype"
Public launch with narrative	Month 24	ROI starts here

CHAPTER VIII

Risks, Mitigations, and *Pitfalls*.

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

Risk 1 · Failing to hire the founding talent

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

There are perhaps 200 people worldwide qualified to lead an FHE operation in production. Most work at a handful of companies (Zama, Duality, Inpher, IBM Research, Google) or in academia. Most have never considered working for cosmetics.

Mitigation: treat it as a C-level hire, not an engineer hire. Equity package, technical autonomy, clear mission. Consider acquiring a startup (acquihire) as a shortcut — Zama, Optalysys, Inpher already receive regular offers; the path to partial acquisition exists.

Risk 2 · Compute cost not dropping as fast as expected

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

The FHE cost-reduction curve depends on algorithmic and hardware advances. If the pace slows, the use case may be viable only for low-volume premium products, not for mass market.

Mitigation: start the portfolio in the ultra-luxury segment, where the math already closes. Scale to mass only when the economics allow. In any scenario, the premium case sustains itself.

Risk 3 · Regulator interprets the feature as insufficient

Probability: low. **Impact:** high.

There is a small chance that ANPD, CNIL, or ICO decides that even with FHE, the act of encrypting and sending biometric data to the server already constitutes "processing" in the legal sense. This would be a minority interpretation, but possible.

Mitigation: engage the regulator **before** launch, in consultative mode. Present the architecture. Seek a prior opinion (formal mechanism in the EU; informal but possible in Brazil). Publish the architecture

openly for academic review. This is a legal decision with enormous upside: a regulator who opines favorably becomes an ally.

Risk 4 · Competitor announces first

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

L'Oréal, Estée Lauder, Shiseido, Unilever, P&G — any of them may have a similar project underway. Unlikely, but possible. If it happens, the second mover loses most of the narrative prize.

Mitigation: speed. Every month of delay is another month of risk exposure. Consider partnering with an FHE vendor as an accelerator (Zama Concrete, Duality Confidential AI) instead of building everything from scratch. Higher cost, but significantly shorter time-to-narrative.

Risk 5 · Narrative does not resonate with the consumer

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

There is a possibility that the consumer simply does not care — that privacy fatigue has caused everyone to stop paying attention to these messages.

Mitigation: test the narrative before the campaign. Focus groups, A/B testing on digital copy, sentiment measurement. Position the feature as an additional quality of the product, not an isolated manifesto. In the worst case, the feature still generates regulatory and operational value.

Risk 6 · Vulnerability discovered in the FHE scheme

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

Modern FHE schemes (CKKS, BFV, BGV, TFHE) are based on well-studied lattice problems — the same problems on which NIST standardized the next generation of post-quantum cryptography (ML-KEM, ML-DSA). Mathematical confidence is high. But unexpected theoretical advances are always possible.

Mitigation: use conservative parameters. Track the literature. Maintain a migration plan across schemes (the architecture should be decoupled from any specific scheme). Do not use FHE as the sole defense layer — combine with TLS, AES, and key segregation.

Pitfall 1 · Treating it as an IT project

PT EN

The most common mistake in adopting narrative technology is placing it under the CIO/CTO instead of the CMO/CDO/CEO. This results in perfect technical delivery and a dead narrative. **FHE should report to the CEO or to a C-level sponsor with a broad mandate.**

Pitfall 2 · Communicating too early

Announcing before the product is functional generates two worst-case outcomes: (a) competitors adapt their strategy and neutralize; (b) the press tests the feature and discovers it is not ready. Communicate only after independent validation.

Pitfall 3 · Forgetting the key

FHE protects during computation. But consumer-side key management is where most implementations fail. If the key is in the brand's cloud, FHE has lost its point. If it is on the phone without backup, losing the phone means losing history. **Key management design is half the project.**

PT EN

A letter for the next decade of *beauty*.

For the CEOs, board members, and creative directors of the houses that can still choose to lead.

The industry you lead was built on an old promise: that caring for beauty is, at bottom, caring for the person. That selling a cream, a perfume, a lipstick is selling a form of relationship with oneself. That the brand stands on the consumer's side — not against her, not above her, not over her, but *with* her.

This promise endured for a century. It survived world wars, cultural revolutions, the rise of feminism, the decline of conspicuous consumption, and the rise of digital. It survived because it was — and largely still is — true. The women and men who choose to buy from a house have a real affective relationship with that house. They trust it.

But over the past five years, without anyone decreeing it, the relationship changed in nature. The consumer is no longer merely the one who hands over money and receives product. She became the one who hands over biometric data, behavioral data, emotional data, and in some cases genetic data — and receives product, yes, but also receives an exposure she did not request, does not fully understand, and that no honest terms-of-use document could explain in time.

This change was not caused by bad faith. It was caused by accretion. Each feature added a little more collection. Each campaign asked for a little more personalization. Each CRM had a few more fields. At some point, without decree, without manifesto, the relationship crossed an ethical frontier that no individual executive would consciously cross.

It is possible to walk back. More than that: it is strategically preferable to walk back. Not because regulation demands it — though it does. Not because the regulator is watching — though it is. But because **the original relationship with the consumer was more valuable**. It was more durable. It was the foundation on which the entire brand narrative that still sustains the sector was built.

FHE — Fully Homomorphic Encryption — is the first technology in decades that makes it possible to walk back **without losing the features**. It is possible to keep offering facial analysis, personalized recommendation, hair diagnosis, bespoke fragrance, customized supplementation. It is possible to keep

doing everything the industry learned to do with data over the past ten years. **It is possible to do all of that without ever seeing the data.**

On first reading this sentence sounds paradoxical. On second reading it sounds too technical. On third reading it sounds like the most obvious argument the industry has ever had at its disposal. How is it possible to recommend a cream without seeing the skin? 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 a brand saying, with mathematical truth and not a management promise: "*we take care of you without invading you*". It is the possibility of offering care without extracting surveillance. It is the possibility of giving the consumer back the dignity of not having to choose between personalization and privacy.

This possibility is open today for exactly one brand. The first to understand what lies in its hands. The first to hire the right engineer, call the right counsel, present the case to the right board. The first to publish the whitepaper, run the campaign, defend the position in an interview, and sustain the narrative for two straight years without giving in to the temptation to dilute the message.

In three years, this position will be taken. The question that must be asked, in every boardroom that still has the privilege of choosing, is simple:

> *“Will we be the ones who spoke first, or the ones who had to explain later why it took so long?”*

There is a window. It is short. It is real. It is historically rare. Entire industries wait decades for windows like this — and most of the houses inside them lose them through excess operational prudence. Cosmetics already lost one such window to digital (and most houses lost it to native D2C brands). It lived through another with sustainability (and some houses used it to redefine themselves, others only greenwashed). This is the next. It may be the last of this decade capable of redefining the core narrative of the industry.

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.

PT EN

— End of Volume I

PT EN

APPENDIX A

Executive *Glossary*.

The terms you will hear from the CTO. In plain language.

FHE — FULLY HOMOMORPHIC ENCRYPTION

Cryptography that allows computations to be performed on encrypted data without decrypting it. The result, when decrypted, is identical to what would have been 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 both classical and quantum computers.

CKKS

FHE scheme for real numbers with approximate arithmetic. Best suited for machine learning, statistical analysis, image processing.

BFV / BGV

FHE schemes for integers with exact arithmetic. Used when the result must be identical to the plaintext result — encrypted databases, scoring, counts.

TFHE

Boolean, bitwise FHE scheme. Slower per operation but the most flexible — allows arbitrary programs to run under encryption.

BOOTSTRAPPING

The operation that "reduces the noise" accumulated during an FHE computation, enabling computations of unlimited depth. It is the most expensive operation in FHE, and most optimization effort focuses here.

PSI — PRIVATE SET INTERSECTION

A protocol derived from FHE/MPC that allows two parties to discover the intersection of their sets without revealing the rest. Useful for "who is a common customer" without disclosing the databases.

PT EN

TEE — TRUSTED EXECUTION ENVIRONMENT

Alternative technology based on hardware (Intel SGX, AMD SEV-SNP). Trusts the manufacturer. FHE is mathematically stronger; TEE is currently faster.

MPC — MULTI-PARTY COMPUTATION

Another alternative technology based on collaborative protocols between multiple parties. It works, but requires all parties online; FHE allows asynchronous operation.

LATTIGO

FHE library in Go, maintained by Tune Insight (EPFL spin-off). High maturity, clean code, well suited for integration into commercial products.

OPENFHE

Open-source C++ library, successor to PALISADE. Maintained by Duality Technologies. The most complete in the variety of supported schemes.

CONCRETE

FHE framework from Zama (Paris). Focused on TFHE and on ease of use for non-cryptographer developers. Offers a commercial edition.

PT EN

APPENDIX B

Vendors and *Talent*.

Where the commercial capability sits and where the people are.

Commercial vendors

VENDOR	HEADQUARTERS	FOCUS
Zama	Paris	TFHE, Concrete framework, focus on developer experience
Duality Technologies	US / Israel	OpenFHE, focus on healthcare and finance, heavy consulting
Inpher	Switzerland / US	Hybrid FHE + MPC, focus on finance and healthcare
Tune Insight	Switzerland (EPFL)	Lattigo, focus on federated medical research
Optalysys	United Kingdom	Optical FHE acceleration; specific hardware
Cornami	US	FHE hardware accelerators
Fhenix	Israel	FHE on blockchain; less relevant for cosmetics

Academic centers where the talent lives

- **EPFL (Lausanne)** — laboratory of Christian Cachin / Jean-Pierre Hubaux, origin of Lattigo

- **IBM Research** — historical origin of FHE (Craig Gentry, 2009), still producing papers
- **Microsoft Research** — SEAL library, Kim Laine
- **NJIT / Duality** — Yuriy Polyakov, Kurt Rohloff
- **Université Paris-Saclay / ENS** — Léo Ducas, several central contributions
- **KU Leuven (COSIC)** — Frederik Vercauteren
- **Stanford / MIT / CMU** — several active researchers, especially in ML+FHE

Open-source libraries for the technical team to evaluate

- **Lattigo** (Go) — github.com/tuneinsight/lattigo
- **OpenFHE** (C++) — github.com/openfheorg/openfhe-development
- **Concrete** (Rust + Python) — github.com/zama-ai/concrete
- **SEAL** (C++) — github.com/microsoft/SEAL
- **HElib** (C++) — github.com/homenc/HElib

PT EN

APPENDIX C

30 Questions for the *CTO*.

The list you should bring to your next meeting with the technical team.

Strategy and capability

1. Do we have anyone on the team with a background in lattice cryptography? If not, what is the acquisition plan?
2. What is our current exposure to biometric data processing, by volume and by jurisdiction?
3. Which product features today depend on processing sensitive personal data on the server?
4. Do we have an up-to-date inventory of which data leaves the consumer's device and where it is processed?
5. What is our current legal opinion on the sustainability of that architecture over the next 36 months?

Technical architecture

6. Which FHE scheme makes most sense for our first use case? CKKS, BFV, TFHE, or hybrid?
7. What is the maximum multiplicative depth of our current inference model?
8. Can we run our skin-classification model under FHE with acceptable latency today?
9. What would be the estimated computational overhead, in orders of magnitude, for our case?
10. How do we plan to manage keys on the consumer side? Backup? Recovery?
11. What happens if the consumer changes phones?
12. Can we perform longitudinal comparisons (before/after photos) under encryption?
13. Where are the unavoidable decryption points, and how will we justify them?

Cost and infrastructure

14. What is the estimated computational cost per query, in FHE versus plaintext?
15. Which hardware accelerators are we evaluating? HEXL? FPGA? Specific GPU?
16. What is the path to reduce cost by 10× over the next 18 months?

17. Will we build internally or use a vendor (Zama, Duality, Inpher)?
18. What are the projected capex and opex for the next 24 months?

PT EN

Security and compliance

19. Do the security parameters we plan to use meet the 128-bit level? Post-quantum?
20. Who will perform the independent security audit?
21. How do we intend to demonstrate compliance with LGPD Art. 11 and GDPR special category?
22. Can we publish the architecture openly for public review?
23. What is our migration plan if a scheme is compromised?

Product and narrative

24. Which product feature will we launch first with FHE?
25. How will we communicate the feature to the consumer without jargon?
26. How will we explain it to non-technical press?
27. What defensible marketing claim could we make?
28. Do we have an independent third party willing to validate our claims publicly?

Risks

29. What is the worst case if a competitor announces before us?
30. What is the worst case if a regulator interprets our architecture as insufficient?

PT EN

~

The Mirror That Cannot See

A strategic eBook for the executive leadership of the global cosmetics industry.

Volume I · 2026 Edition · Confidential distribution.

Set in Iowan Old Style and SF Pro.

Built as a self-contained HTML document.

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

— end —