AI Congress Barcelona

20 - 21 - 22 OCT 2026

Auditori L'illa

Un esdeveniment de:

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per:

eurecat

Alineat amb l'Estratègia Catalunya IA 2030:

Generalitat de Catalunya

Segueix-nos a:

#AICongressBCN

aicongress.barcelona

Patrocinadors OR:

CaixaBank TECH

conversia compliance is our business

forvis mazars & infini DIGITAL EFFICIENCY

orange Empresas

raona coding corporate futures

softeng Maximize the cloud. Securely.

Patrocinadors PLATA:

apdcat Autoritat Catalana de Protecció de Dades

LABORAL kutxa

Col·laboradors:

Barcelona

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

CVC Centre de Visió per Computador

i2cat

UPC Universitat Politècnica de Catalunya Barcelona Tech

IDEAI Intelligent Data Science and Artificial Intelligence Research Center

AI Congress Barcelona

Un esdeveniment de: CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per: eurecat

Alineat amb l'Estratègia Catalunya IA 2030: Generalitat de Catalunya

20 - 21 - 22 OCT 2026

#AICongressBCN

aicongress.barcelona

IA, REGULACIÓ I SOBIRANIA TECNOLÒGICA: CONSTRUINT ARQUITECTURES PER A SISTEMES AUDITABLES

Patrocinadors OR: CaixaBank TECH · conversia compliance is our business · forvis mazars & infini DIGITAL EFFICIENCY · orange Empresas · raona coding corporate futures · softeng Maximize the cloud. Securely.

Patrocinadors PLATA: apdcat Autoritat Catalana de Protecció de Dades · LABORAL kutxa

IA, REGULACIÓ I SOBIRANIA TECNOLÒGICA: CONSTRUINT ARQUITECTURES PER A SISTEMES AUDITABLES

Presenta:



Xavier Domingo
Director de la unitat d'Applied Artificial Intelligence (AAI)
Eurecat

Ponents:



Mikel Urizar
Responsable del Centre d'Excel·lència d'IA Generativa per a Iberia i Llatinoamèrica
Inetum



Toni Manzano
CSO i cofundador
Aizon



Xavier Manzanera
Responsable de l'Oficina Digital, Dades i IA
Esteve Healthcare



Marc Oliveras
Vicepresident d'Enginyeria i CTO a Captio
Emburse

Patrocinadors OR: CaixaBank TECH · conversia compliance is our business · forvis mazars & infini DIGITAL EFFICIENCY · orange Empresas · raona coding corporate futures · softeng Maximize the cloud. Securely.
Patrocinadors PLATA: apdcat Autoritat Catalana de Protecció de Dades · LABORAL kutxa



PONÈNCIA:

CRAFT: IA GENERATIVA SOBIRANA I AUDITABLE

Mikel Urizar
Responsable del Centre d'Excel·lència d'IA Generativa per a Iberia i Llatinoamèrica
Inetum

AI CONGRESS BARCELONA

¿Quién controla **vuestra IA?**



Diseñar sistemas que podáis controlar y auditar

Mikel Urizar · Inetum · CoE GenAI

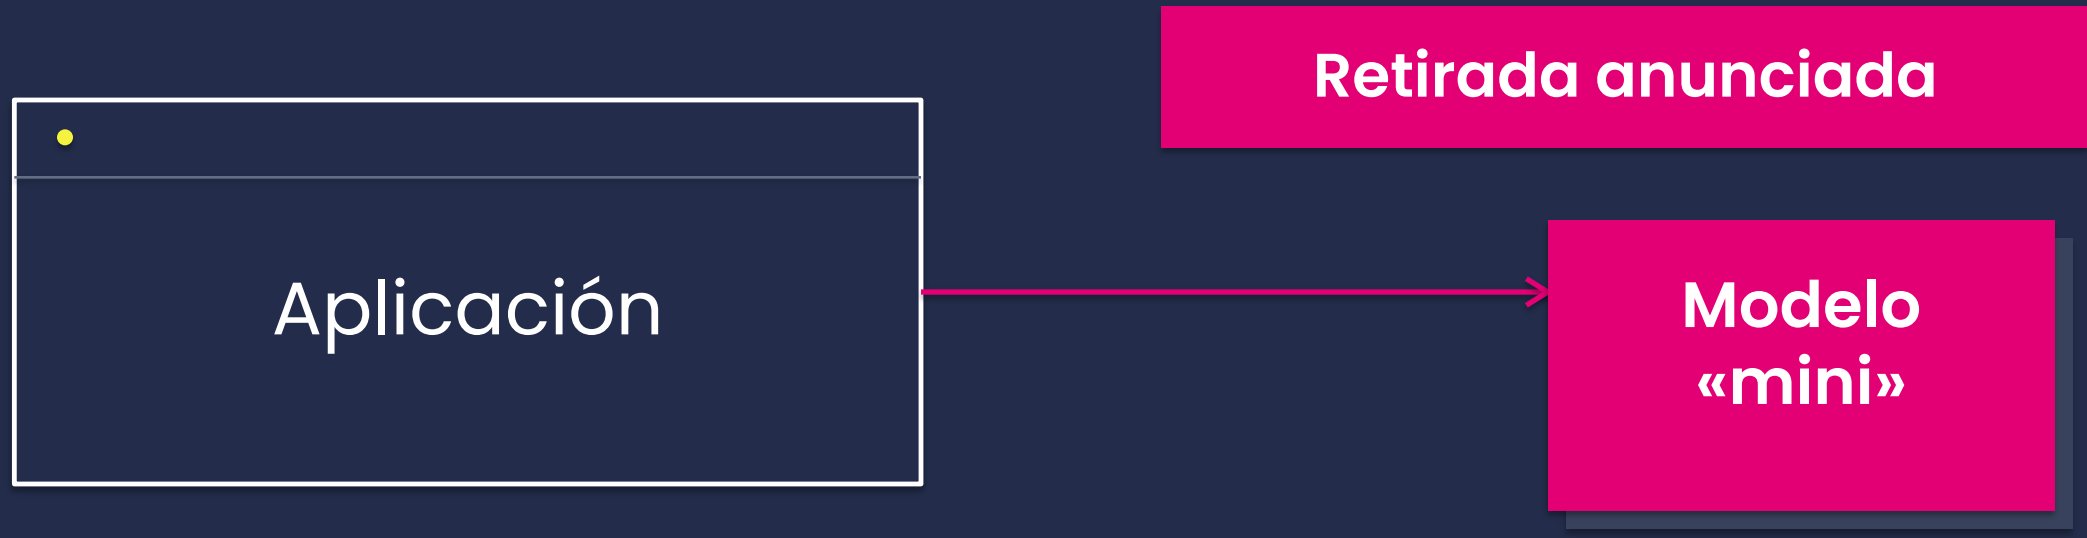
inetum

La aplicación funcionaba.

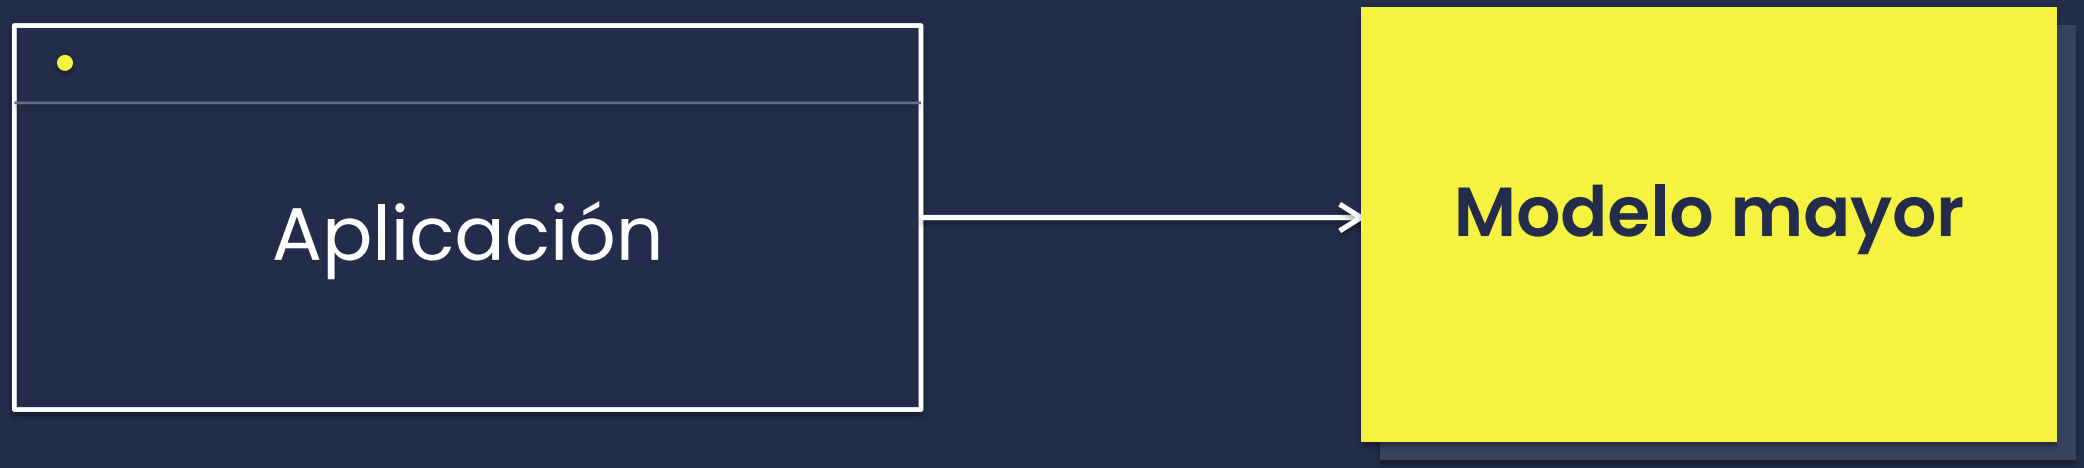


Resultados útiles

El proveedor anunció **la retirada del modelo.**



Resultados comparables.
Mayor coste.



La aplicación continúa

Seguía funcionando.

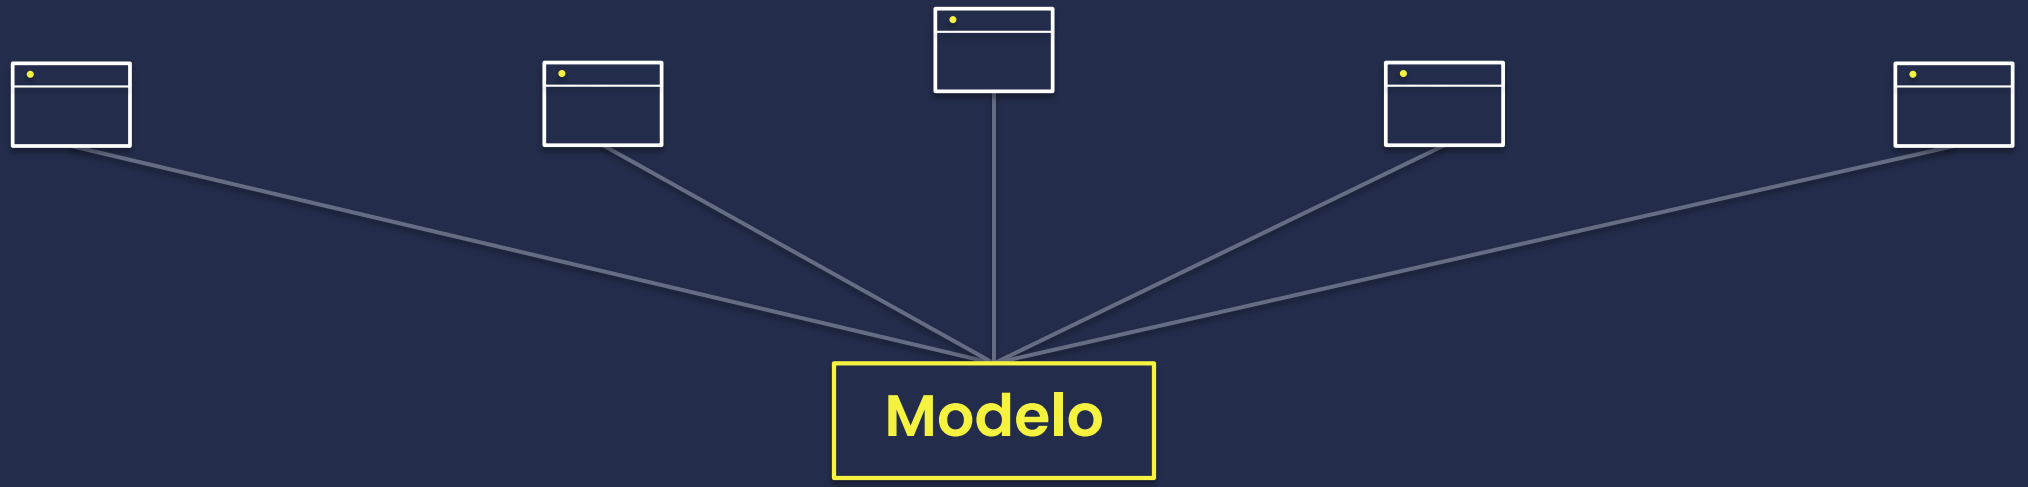
¿Seguía compensando?

Revisar rentabilidad y riesgos

La responsabilidad **sigue siendo vuestra.**

Aunque el modelo sea de otro...

¿Qué aplicaciones dependen **de ese modelo?**



No podéis gobernar **lo que no conocéis.**

Qué hace cada solución y quién responde por ella.

Registro central

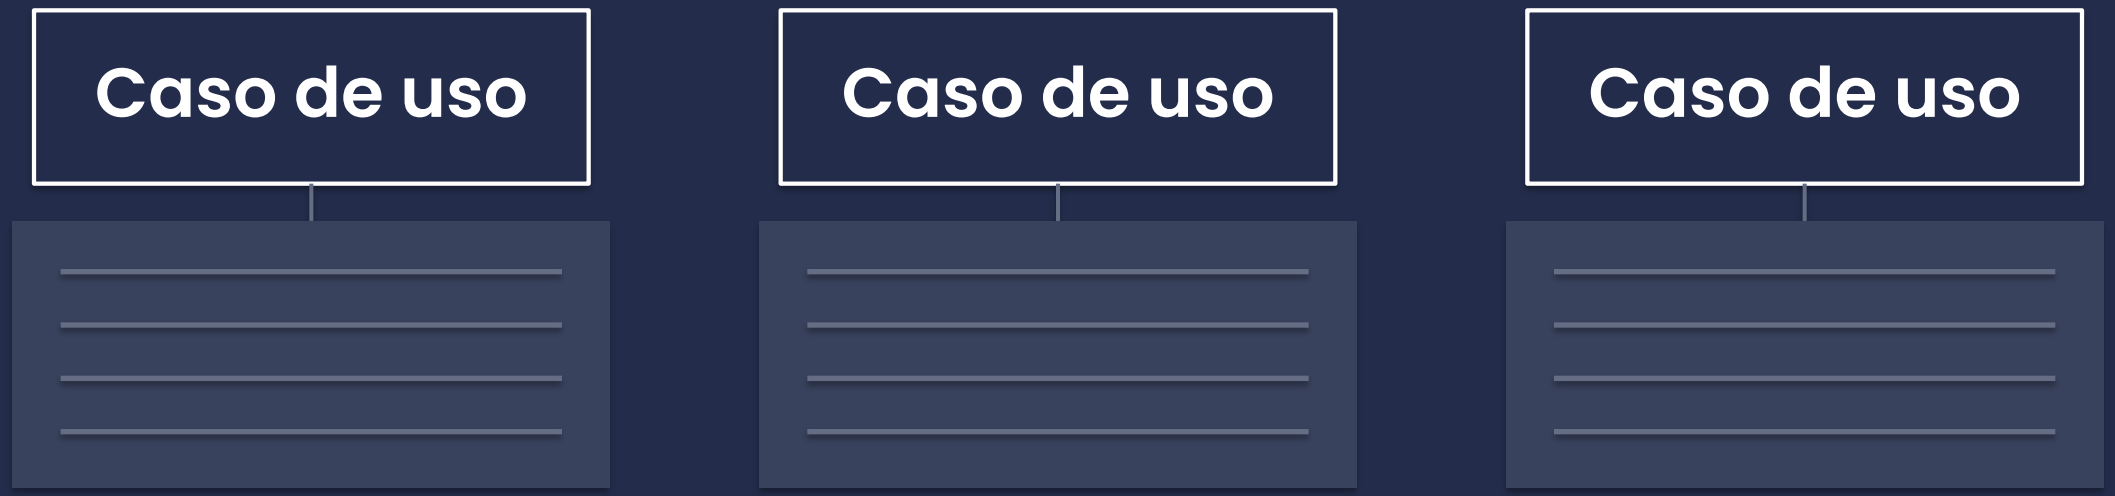
Finalidad	Datos	Responsable	Modelos

¿Qué resultado justifica **esta aplicación?**

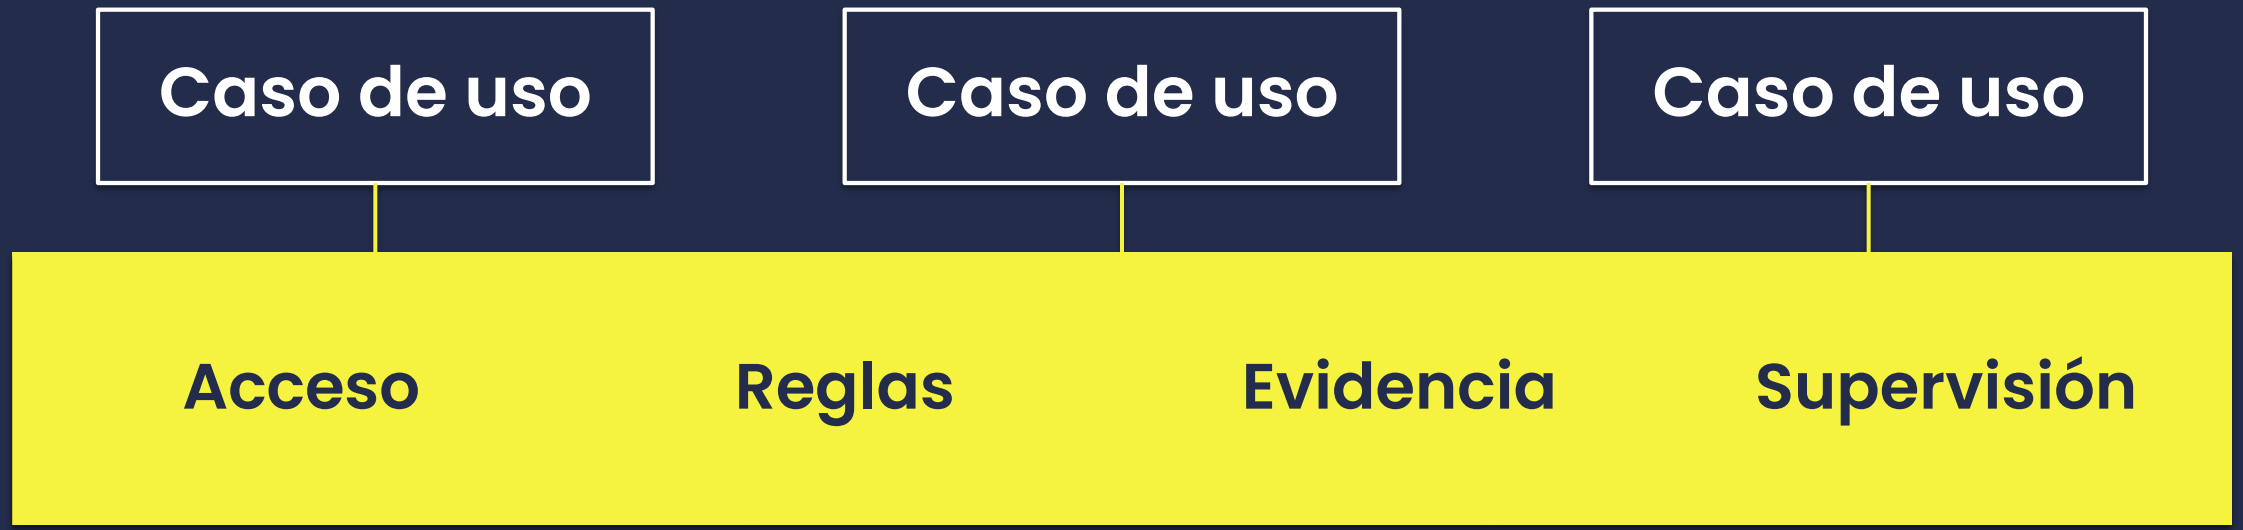
Uso y gasto ≠ **Resultado del proceso**

¿Tiene que empezar de cero **cada equipo?**

La misma base, repetida



Una base común para centrarse **en el caso de uso.**

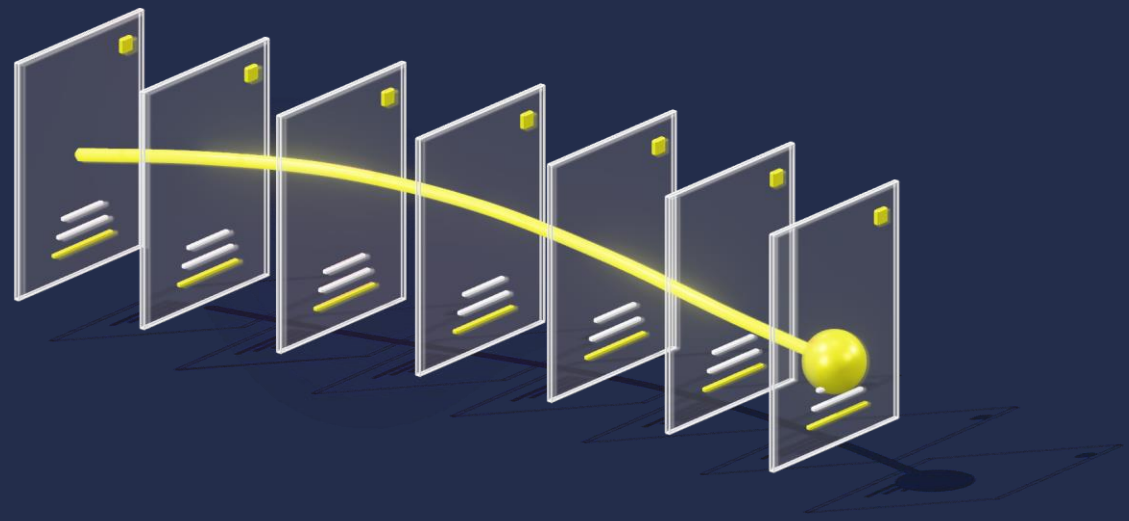


Cambiar de modelo exige **validar de nuevo la aplicación.**

Pruebas automatizadas y criterios para aceptar el cambio.



La evidencia se genera **mientras trabajáis.**



¿Qué necesitáis poder demostrar?

RIA / AI Act

Alto riesgo · obligaciones
según vuestro papel

ISO/IEC 42001

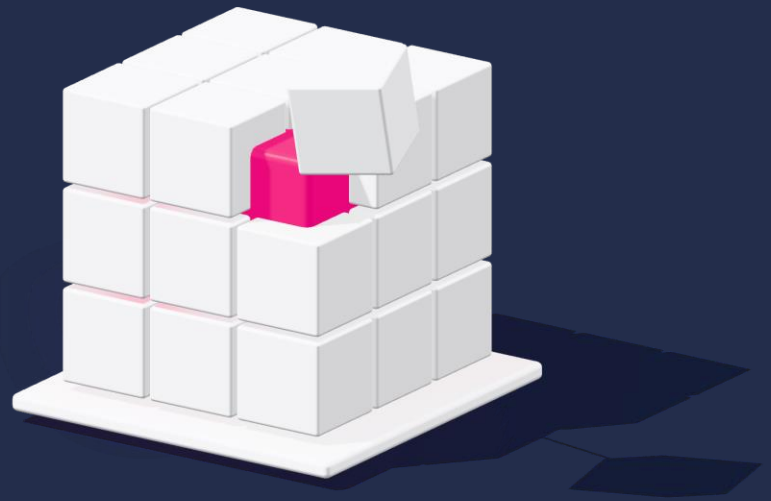
Sistema de gestión de IA

Definir los requisitos desde el diseño.

¿Sabéis qué está **verificado y qué falta?**

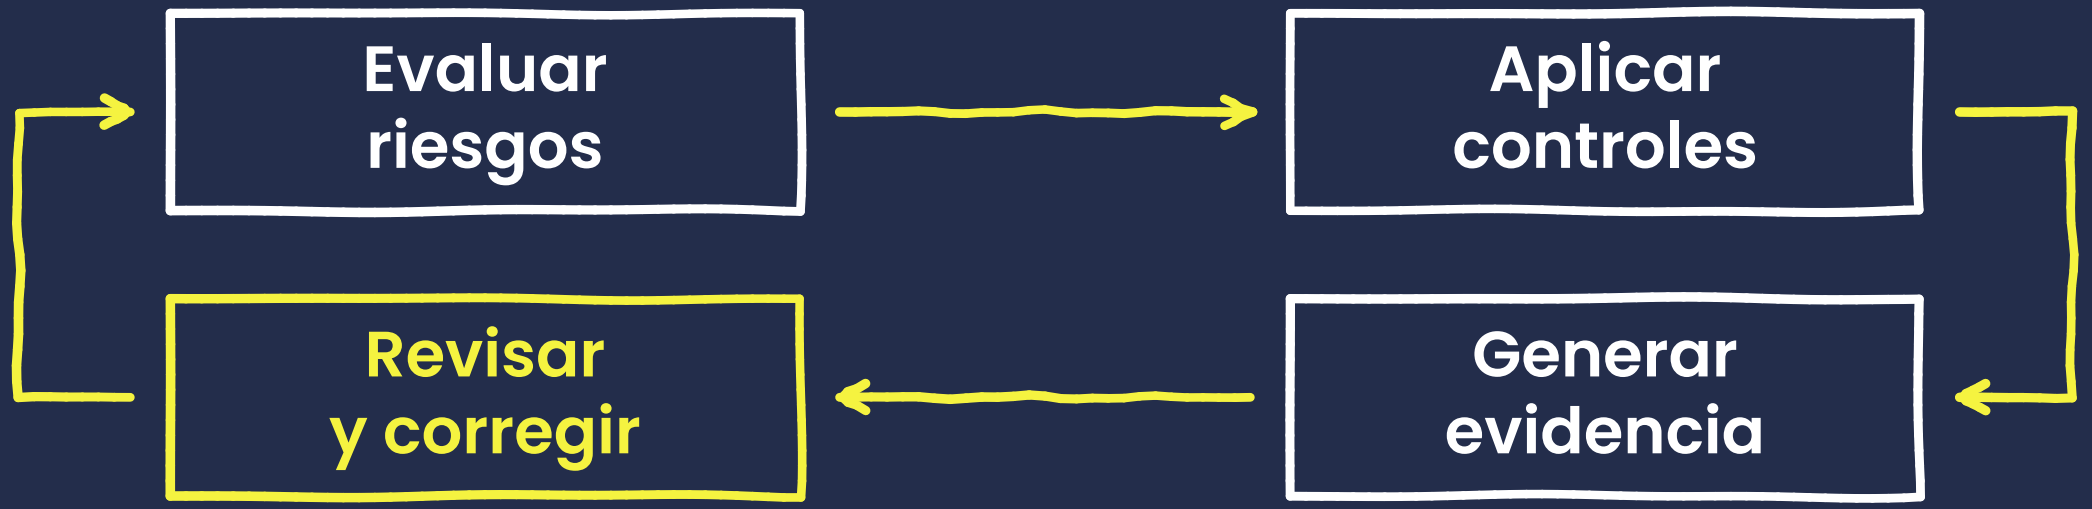
Por qué falta y cómo resolverlo.

- Cumple
- Cumple parcialmente
- No cumple

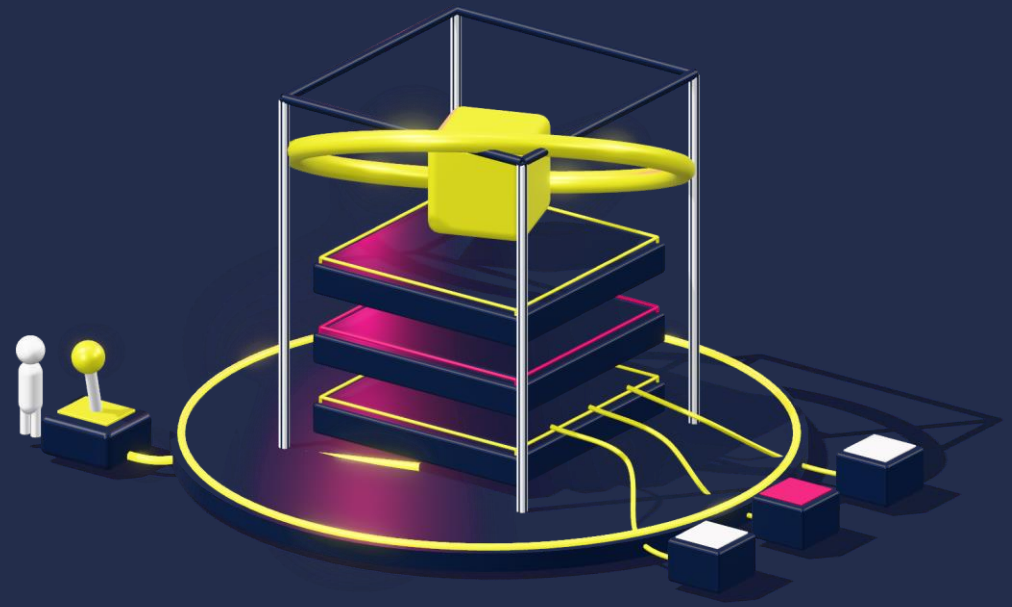


Tener evidencias **no equivale a cumplir.**

Revisar los controles periódicamente y ante cambios.



Saber qué falta.
Poder corregirlo.



NUESTRA EXPERIENCIA

El valor está en cómo lo conectamos.

Craft

ISO/IEC 42001 Certificado

Sistema de gestión de IA certificado

Reutilizar lo que existe.
Construir lo que falta.

11 módulos · 3 capas

+50

componentes de código abierto

en Kubernetes

Aplicación | Aplicación | Aplicación | Aplicación

01 Servicios de IA

Agentes
Composer · Bridge

Documentos
Forge · Compass · Refiner

Voz e imagen
Voice · Sense

02 Datos y modelos de IA

Acceso y reglas · Gateway

Evaluación · Lens

03 Infraestructura

Observabilidad y seguridad · Core

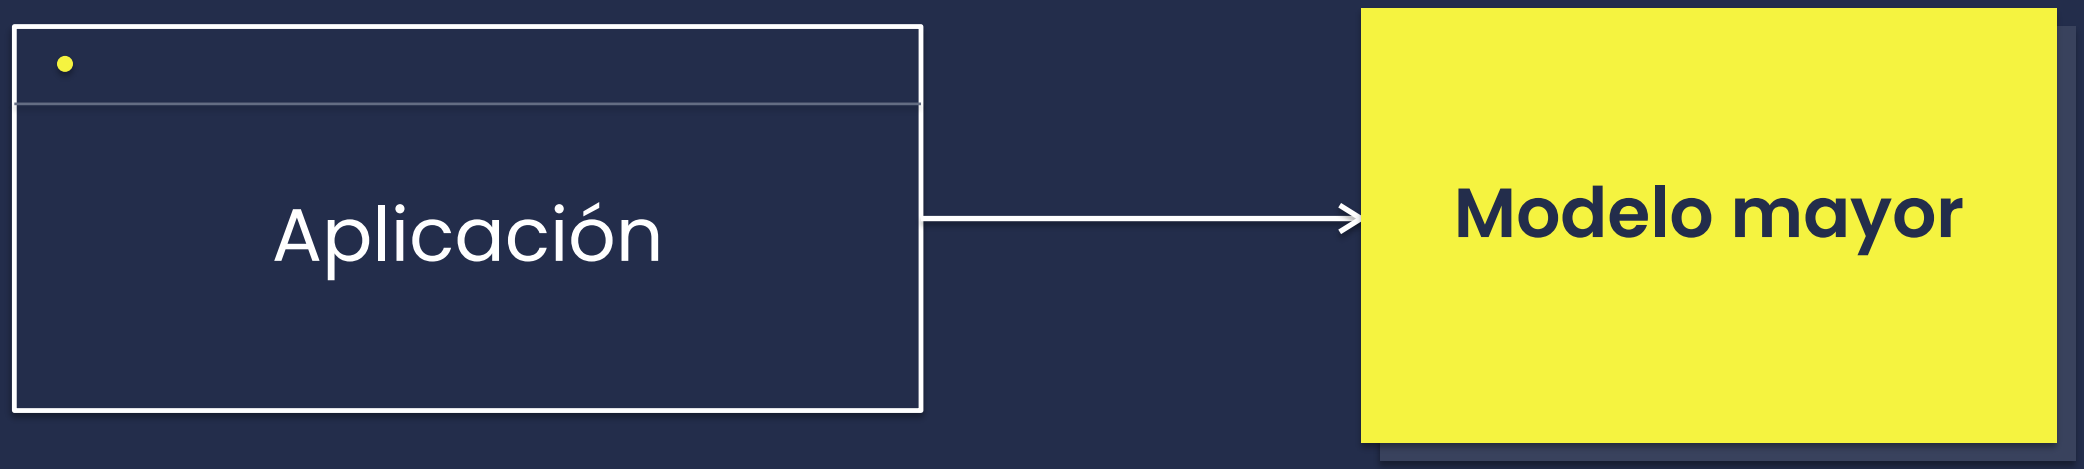
Secretos · Vault

Seguridad · Cumplimiento · Calidad

Diseñada para alta disponibilidad y +10.000 usuarios corporativos.

inetum

Si fuera vuestra aplicación, **¿quién decidiría?**



Seguía funcionando. Costaba más.

VUESTRA DECISIÓN

¿Quién controla vuestra IA?

Pensad en una aplicación que uséis hoy.

Mikel Urizar

GenAI CoE Manager

mikel.urizar@inetum.com

Muchas gracias.



Congress Barcelona

Un esdeveniment de: CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per: eurecat

Alineat amb l'Estratègia Catalunya IA 2030: Generalitat de Catalunya

20 - 21 - 22 OCT 2026
#AICongressBCN
aicongress.barcelona

IA, REGULACIÓ I SOBIRANIA TECNOLÒGICA: CONSTRUINT ARQUITECTURES PER A SISTEMES AUDITABLES

Patrocinadors OR: CaixaBank TECH · conversia compliance is our business · forvis mazars & infini DIGITAL EFFICIENCY · orange Empresas · raona coding corporate futures · softeng Maximize the cloud. Securely.

Patrocinadors PLATA: apdcat Autoritat Catalana de Protecció de Dades · LABORAL kutxa



Toni Manzano
CSO i cofundador
Aizon

PONÈNCIA:

IA APLICADA A LA FABRICACIÓ FARMACÈUTICA: QUAN LA REGULACIÓ ÉS EL DRIVER PRINCIPAL

aizon

AI applied to pharmaceutical manufacturing: when regulation is the main driver

Toni Manzano (toni.manzano@aizon.ai)

©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Why regulation is required in Pharma







©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Regulators and inspectors





aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Regulators and inspectors











Peter Twomey

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

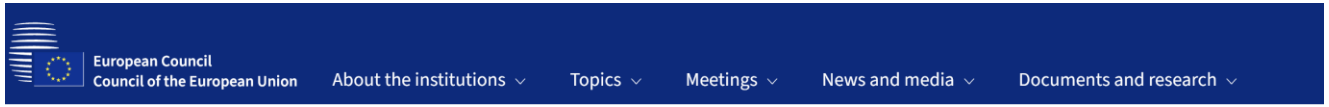
Regulatory framework

EU AI Act

Documentation
EU-GMP Chapter 14

Computerized Systems
EU-GMP Chapter 4
Annex 11

Computerized Systems
EU-GMP Chapter 4
Annex 22



Classification of AI systems as high-risk and prohibited AI practices

The new law categorises different types of artificial intelligence according to risk. AI systems presenting only **limited risk** would be subject to very light transparency obligations, while **high-risk** AI systems would be authorised, but subject to a set of requirements and obligations to gain access to the EU market. AI systems such as, for example, **cognitive behavioural manipulation** and **social scoring** will be banned from the EU because their risk is deemed unacceptable. The law also prohibits the use of AI for **predictive policing** based on profiling and systems that use **biometric data** to categorise people according to specific categories such as race, religion, or sexual orientation.

General purpose AI models

The AI act also addresses the use of **general-purpose AI (GPAI)** models.

GPAI models not posing systemic risks will be subject to some limited requirements, for example with regard to **transparency**, but those with systemic risks will have to comply with stricter rules.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Regulatory framework



04/2016:52100

5.21. CHEMOMETRIC METHODS APPLIED TO ANALYTICAL DATA

1-1-4. Introducing chemometrics

In chemometrics, a property of interest is evaluated solely in terms of the information contained in the measured samples. Algorithms are applied directly to the data set, and information of interest is extracted with models (modelling or calibration step). Chemometrics is associated with multivariate data analysis, which usually depends less on assumptions about the distribution of the data than many other statistical methods

substances, counterfeits etc.

2-10. *ARTIFICIAL NEURAL NETWORKS*

Artificial neural networks (ANNs) are general computational tools, whose initial development was inspired by the need for further understanding of biological neural networks and which have since been widely used in various areas that require data processing with computers or machines. The methods for building ANN models and their subsequent applications can be dramatically different depending on the architecture of the neural networks themselves. In the field of chemometrics, ANNs are generally used for multivariate calibration and unsupervised classification, which is achieved by using multi-layer feed-forward (MLFF) neural networks, or self-organising maps (SOM) respectively. As a multivariate calibration tool, ANNs are more generally associated with the mapping of non-linear relationships.

2-10-2. **Principle**

2-10-2-1. *General*

The basic data processing element in an artificial neural network is the artificial neuron, which can be understood as a mathematical function that uses the sum of a weighted vector and a bias as the input. The vector is the 'input' of the neuron and is obtained either directly from a sample in the data set or calculated from previous neurons. The user chooses the form

unlimited possibilities when connecting different numbers and layers of neurons, is often called an ANN architecture, and provides the potential for ANNs to meet any complicated data modelling requirements.

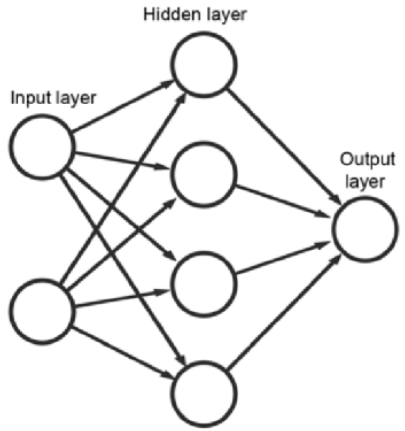


Figure 5.21.-13. - *Typical arrangements of neuron layers and their inter-connections*

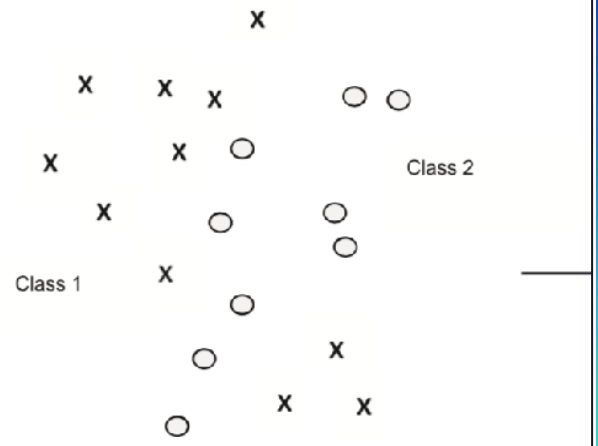
2-9. *SUPPORT VECTOR MACHINES*

To achieve classification, multivariate techniques reduce the dimensionality and complexity of the data set. Kernel methods project data into higher dimensional feature spaces.

2-9-2. **Principle**

Support vector machines (SVMs) project *X*-data of the training set into a feature space of usually much higher dimension than the original data space. In the feature space a hyperplane (also called decision plane) is computed that separates individual points of known group membership (Figure 5.21.-11). The best discriminating separation is achieved by maximising the margin between groups. The margin is defined by 2 parallel hyperplanes at an equal distance from the decision plane. The optimum position of the decision plane is obtained if the margin is maximal. Points in the feature space that define the margin are called support vectors.

For each training point the distance to the decision plane is computed. In the case of a two-class separation for example, the sign of the distance gives the group membership and the



aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Regulatory framework

GUIDANCE DOCUMENT

Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products

Draft Guidance for Industry and Other Interested Parties

JANUARY 2025

Download the Draft Guidance Document | Read the Federal Register Notice

Draft | Level 1 Guidance

Not for implementation. Contains non-binding recommendations.

1. *Step 1: Define the Question of Interest*

Step 1 in the framework is to define the question of interest. The question of interest should describe the specific question, decision, or concern being addressed by the AI model.

2. *Step 2: Define the Context of Use for the AI Model*

Step 2 in the framework is to define the COU for the AI model. The COU defines the specific role and scope of the AI model used to address a question of interest. The description of the COU should describe in detail what will be modeled and how model outputs will be used. The COU should also include a statement on whether other information (e.g., animal or clinical studies) will be used in conjunction with the model output to answer the question of interest.

European Commission | EN

Public Health

Home > Consultations > Stakeholders' Consultation on EudraLex Volume 4 - Good Manufacturing Practice Guidelines: Chapter 4, Annex 11 and New Annex 22

CONSULTATION | Open

Stakeholders' Consultation on EudraLex Volume 4 - Good Manufacturing Practice Guidelines: Chapter 4, Annex 11 and New Annex 22

Revision of Good Manufacturing Practice (GMP) Guidelines Chapter 4 (Documentation), Annex 11 (Computerised Systems) and New Annex 22 (Artificial Intelligence)

PAGE CONTENTS
- Details
- Target audience
- Why we are consulting
- Respond to the consultation
- Reference documents

Details

Status	OPEN
Opening date	7 July 2025
Deadline	7 October 2025, 23:59 (CEST)
Department	Directorate-General for Health and Food Safety

Annex 22: Artificial Intelligence in GMP Systems

The publication of Annex 22 signals a pivotal moment in pharmaceutical regulation: for the first time, the use of **artificial intelligence (AI)** and **machine learning (ML)** within GMP environments has been formally recognized and regulated.

The annex introduces detailed guidance on how AI/ML systems must be designed, validated, monitored, and maintained to ensure patient safety and data integrity.

Defined Scope for AI/ML Usage

Annex 22 sets boundaries for the use of artificial intelligence. It applies only to:

- **Static, deterministic AI/ML models** used within computerized systems
- Applications with **direct GMP impact** (e.g., product release decisions, data classification)

It **excludes**:

- Generative AI (e.g., LLMs like ChatGPT)
- Adaptive or self-learning models without fixed behavior

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Regulatory framework

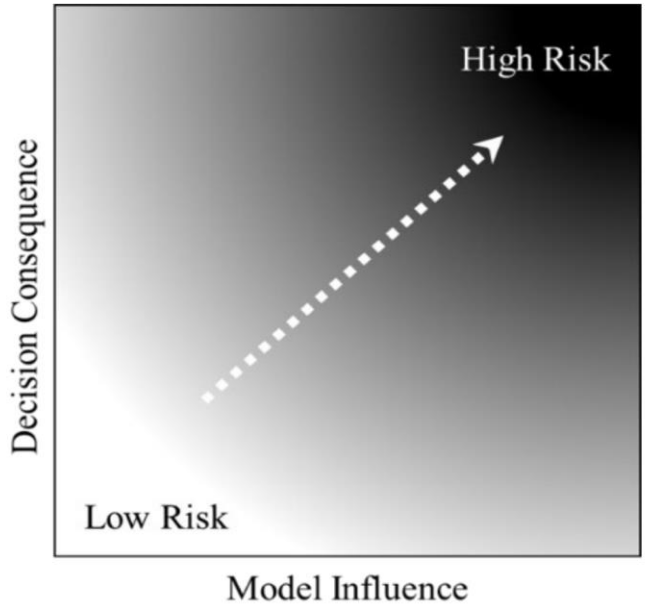


Figure 1. Model risk matrix. The model risk moves from low to high as decision consequence or model influence increases. The ratings for decision consequence and model influence are independently determined.

EMA EUROPEAN MEDICINES AGENCY FDA

Guiding principles of good AI practice in drug development

January 2026

Artificial Intelligence (AI) has the potential to transform the way drugs (medicines)[1] are developed and evaluated, ultimately improving health care. In this context, AI refers to system-level technologies used to generate or analyse evidence across the drug product life cycle, including nonclinical, clinical, post-marketing, and manufacturing phases.

Drugs are authorised based on demonstrated quality, efficacy and safety, and when their benefits outweigh their risks. As new technologies emerge, including AI, it is essential that their use reinforces these requirements for the benefit and safety of patients.

The use of AI throughout the drug product life cycle has increased significantly in recent years. The complex and dynamic processes involved in developing, deploying, using, and maintaining AI technologies benefit from careful management throughout the drug product life cycle to ensure outputs are accurate and reliable. Among other innovations, AI technologies are anticipated to support a multi-faceted approach that promotes innovation, reduces time-to-market, strengthens regulatory excellence and pharmacovigilance, and decreases reliance on animal testing by improving the prediction of toxicity and efficacy in humans. This document outlines a common set of principles to inform, enhance, and promote the use of AI for generating evidence across all phases of the drug product life cycle.

These 10 guiding principles are intended to lay the foundation for developing good practice that addresses the unique nature of these technologies. They will also help cultivate future growth in this rapidly progressing field.

The 10 guiding principles identify areas where the international regulators, international standards organisations, and other collaborative bodies could work to advance good practice in drug development. Areas of collaboration include research, creating educational tools and resources, international harmonisation, and consensus standards, which may help inform regulatory policies and regulatory guidelines in different jurisdictions, in line with applicable legal and regulatory frameworks.

As the use of AI in drug development evolves, so too must good practice and consensus standards. Strong partnerships with international public health partners will be crucial to empower stakeholders to advance responsible innovations in this area. Thus, this initial collaborative work can inform our broader international engagements.

Principles

1
Human-centric by design

The development and use of AI technologies align with ethical and human-centric values.

2
Risk-based approach

The development and use of AI technologies follow a risk-based approach with proportionate validation, risk mitigation, and oversight based on the context of use and determined model risk.

3
Adherence to standards

AI technologies adhere to relevant legal, ethical, technical, scientific, cybersecurity, and regulatory standards, including Good Practices (GxP).

4
Clear context of use

AI technologies have a well-defined context of use (role and scope for why it is being used).

5
Multidisciplinary expertise

Multidisciplinary expertise covering both the AI technology and its context of use are integrated throughout the technology's life cycle.

6
Data governance and documentation

Data source provenance, processing steps, and analytical decisions are documented in a detailed, traceable, and verifiable manner, in line with GxP requirements. Appropriate governance, including privacy and protection for sensitive data, is maintained throughout the technology's life cycle.

7
Model design and development practices

The development of AI technologies follows best practices in model and system design and software engineering and leverages data that is fit-for-use, considering interpretability, explainability, and predictive performance. Good model and system development promotes transparency, reliability, generalisability, and robustness for AI technologies contributing to patient safety.

8
Risk-based performance assessment

Risk-based performance assessments evaluate the complete system including human-AI interactions, using fit-for-use data and metrics appropriate for the intended context of use, supported by validation of predictive performance through appropriately designed testing and evaluation methods.

9
Life cycle management

Risk-based quality management systems are implemented throughout the AI technologies' life cycles, including to support capturing, assessing, and addressing issues. The AI technologies undergo scheduled monitoring and periodic re-evaluation to ensure adequate performance (e.g., to address data drift).

10
Clear, essential information

Plain language is used to present clear, accessible, and contextually relevant information to the intended audience, including users and patients, regarding the AI technology's context of use, performance, limitations, underlying data, updates, and interpretability or explainability.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Annex 22 vs. FDA Draft Guidance

Topic	EMA Annex 22 (EU)	FDA Draft Guidance (US)	Discrepancies
Scope	Applies only to *static*, *deterministic* AI/ML models used in *GMP manufacturing*. Excludes dynamic models and generative AI from critical applications.	Applies to the *entire drug product lifecycle*: nonclinical, clinical, manufacturing, and post-market. Includes any AI used for regulatory decision-making.	Annex 22 is narrow (GMP only) and restrictive (static models). FDA is broad (entire lifecycle) and inclusive (any AI that informs decisions).
Generative AI and LLMs	Explicitly excludes generative AI and LLMs from critical GMP uses. If used in noncritical GMP, requires human-in-the-loop.	Does not exclude generative AI or LLMs. Applies risk-based principles regardless of model type.	EMA prohibits generative AI in critical settings; FDA does not. EMA is prescriptive; FDA is flexible.
Risk Framework	Uses standard GMP/Annex 11 principles: patient safety, product quality, data integrity. Does not define a formal AI-specific risk model.	Uses a structured risk-based credibility framework with defined concepts like *model influence* and *model consequence*.	FDA has a formal, quantitative risk paradigm; EMA follows traditional GMP-based qualitative risk management.
Context of Use / Intended Use	Uses *intended use*: what the model does, what data it uses, limitations, subgroups. Requires SME involvement.	Uses *context of use (COU)*: defines role, impact, and specific decision being supported. Central to the credibility assessment.	Terminology and regulatory weight differ. FDA uses COU as a central pillar; EMA uses intended use in a more traditional GMP way.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Annex 22 vs. FDA Draft Guidance

Topic	EMA Annex 22 (EU)	FDA Draft Guidance (US)	Discrepancies
Dynamic Models	Explicit prohibition: dynamic/adaptive models cannot be used in critical GMP because they change after validation.	Allows models that evolve, provided credibility and lifecycle controls are maintained. Supports adaptive algorithms.	FDA permits evolving models with controls; EMA completely restricts them.
Explainability	Requires *specific explainability methods* such as SHAP, LIME, heatmaps. Must review features used in decisions.	Requires transparency and justification but does *not* mandate specific tools. Allows black-box models if credibility is demonstrated.	EMA is prescriptive; FDA is method-agnostic. FDA accepts black-box AI if justified; EMA expects feature-level explainability.
Test Data Requirements	Strict independence between training and test data. Requires audit trails, controlled access, and justification for any exclusions. Discourages synthetic/GAN data.	Requires independence and suitability for COU but focuses on documentation and justification rather than strict segregation. Allows synthetic data if validated.	EMA is stricter on test data segregation. FDA is more flexible and emphasizes transparency over hard prohibitions.
Operation & Lifecycle	Focuses on GMP controls: change control, configuration management, monitoring model performance and data drift.	Emphasizes continuous lifecycle management: model maintenance, recalibration, monitoring, re-assessment of credibility over time.	EMA's view is static and fits GMP; FDA sees AI as evolving and focuses on lifecycle and iterative improvement.

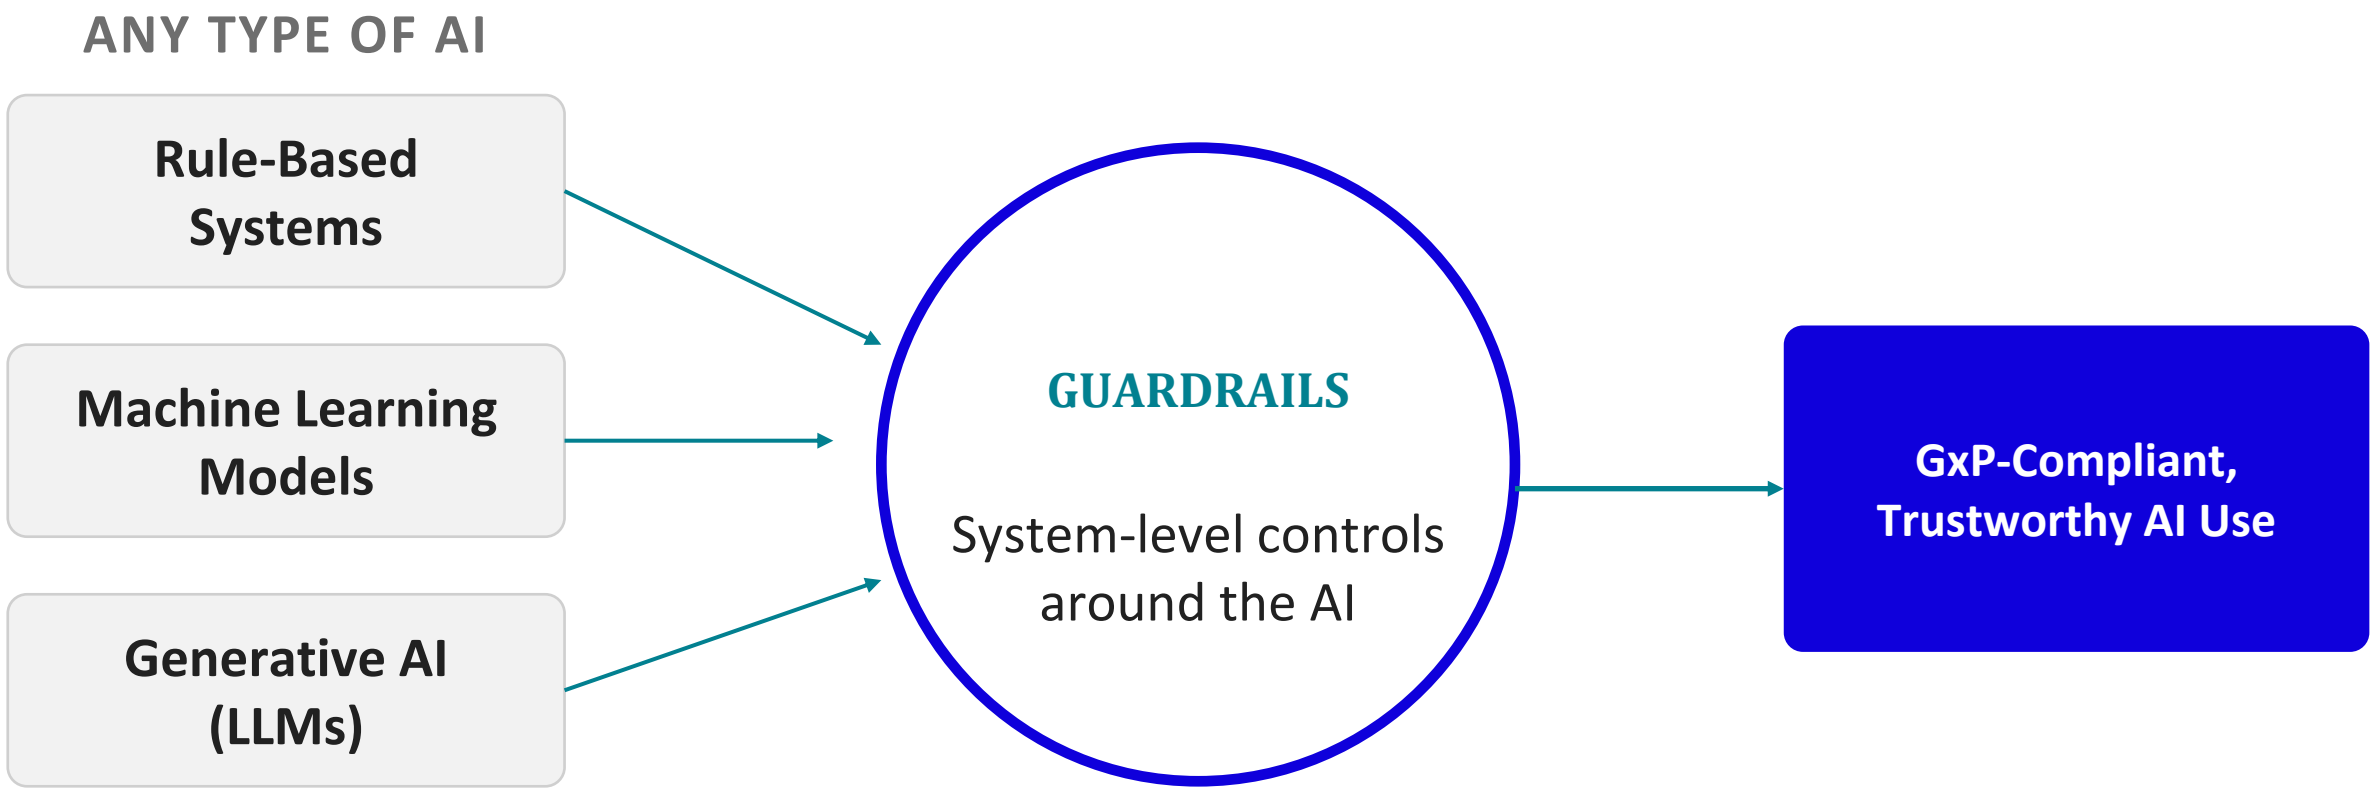
aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Annex 22 vs. FDA Draft Guidance

Topic	EMA Annex 22 (EU)	FDA Draft Guidance (US)	Discrepancies
Regulatory Engagement	No structured mechanisms described. Engagement implied through standard GMP approval and documentation.	Provides explicit pathways: INTERACT, CID, MIDD, ETP, CATT. Encourages early and frequent dialogue.	FDA provides structured and proactive engagement. EMA remains implicit and limited to GMP processes.
Binding vs. Nonbinding Nature	Once finalized, Annex 22 is *binding* as part of EU GMP regulations.	FDA guidance is *nonbinding* and represents current thinking.	EMA has regulatory force; FDA operates through nonbinding recommendations and risk-based discretion.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Guardrails: System Controls for AI



Guardrails are not a property of the AI itself:they are the system controls placed around it. The same guardrail architecture applies whether the AI is rule-based, statistical, or generative.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

AI guardrail mechanisms for non-deterministic AI application in GxP

Input Guardrails: Prevention



Check that every data input belongs to the validated operating space before the model runs. If it does not, reject it. The model never sees data it was not designed to handle.

In-Model Guardrails: Detection

RAG + Citation binding
Uncertainty estimator
Parametric boundaries

While the model processes the input, measure how confident it is in its own answer. If confidence falls below a validated threshold, raise a flag. Do not wait for a wrong answer to appear . Catch the uncertainty before it becomes an output.

Output Guardrails: Containment



Before any model output reaches a GMP decision point, verify it against defined acceptance criteria. If it falls outside the validated range or triggers a disagreement signal, hold it and route it to a human reviewer. Nothing leaves the model unchecked.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Layer 1. Input Guardrails: Prevention

***Act before inference**. Validate that inputs belong to the model's validated operating space.*

1.1 Out-of-Distribution Detection: Mahalanobis Distance

Measures the statistical distance of a new input vector x from the training distribution (μ, Σ). A value exceeding threshold τ flags the input as outside the validated space before inference occurs.

GMP analogy: just as a raw material is tested against its approved specification before it is allowed to enter the manufacturing process, incoming data is checked against the validated training distribution before the AI model is allowed to process it. Data that falls outside that distribution is quarantined and rejected, exactly as a non-conforming material would be.

Reference: *Lee et al. (2018, NeurIPS)* → *AUROC > 0.95 for OOD detection on deep networks.*

Rational: $MD(x) = \sqrt{[\ (x - \mu)^T \cdot \Sigma^{-1} \cdot (x - \mu)\ ]}$ **If** $MD(x) > \tau$ → **BLOCK INPUT**

1.2 Schema & Semantic Validation

- **Intended use filtering:** a classification layer checks that the nature of the query is consistent with the declared intended use of the model. A batch summarisation model will not accept a release decision prompt. A process monitoring model will not accept data from a different unit operation.
- **Context of use verification:** inputs must carry a structured metadata tag declaring the process step, document type, and decision being supported. An untagged or inconsistent input is rejected, just as an unlabelled material would be rejected at goods receiving.
- **Prompt scope boundary enforcement:** a fine-tuned classifier detects prompts that ask the model to operate outside its validated domain, even when the prompt is syntactically valid and statistically in-distribution. Out-of-scope prompts are blocked before inference.
- **Formal Grammar Constraints:** Inputs must conform to a defined schema (e.g., structured batch record fields, LIMS data formats, SOP-defined parameters). Schema violations are rejected pre-inference.
- **Hallucination-Prone Prompt Detection:** Classifier models trained to flag prompts that statistically correlate with fabricated outputs in GMP document workflows (batch records, deviation reports, CAPAs).
- **Formal structure checks:** inputs are validated against a defined schema before reaching the model. Any record that does not match the expected format is rejected without the model ever running.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Layer 2. In-Model Guardrails: Detection

***Operate during inference**. Provide mathematically rigorous uncertainty signals that trigger human escalation.*

2.1 Conformal Prediction

Distribution-free, statistically guaranteed coverage. It requires no assumption about the data distribution. Critical for adaptive models that evolve post-deployment.

Rational: $P(y \in C(x)) \geq 1 - \alpha$

α = accepted error rate (e.g. α = 0.01 → 99% coverage guarantee). Large prediction set C(x) → high uncertainty → escalate to human reviewer.
Reference: Angelopoulos & Bates (2023, SIAM Review).

2.2 Bayesian Uncertainty Quantification: MC-Dropout

Decomposes uncertainty into aleatoric (irreducible data noise) and epistemic (model ignorance, reducible). Guardrails act on epistemic uncertainty. N stochastic forward passes with dropout active at inference time.

Rational: $\sigma^2 ep \approx (1/N)\Sigma\ \hat{f}_i(x)^2 - [(1/N)\Sigma\ \hat{f}_i(x)]^2$

High σ^2ep on GMP-critical output → validated threshold exceeded → hold output & escalate for human review.
Reference: Gal & Ghahramani (2016, ICML).

2.3 Self-Consistency Scoring (LLMs / GenAI)

Query the model K times with semantically equivalent prompts. High output variance signals unreliable generation — flags fabricated citations or invented batch data.

Rational: $C = (1/K^2)\ \Sigma_i\Sigma_j\ sim(o_i, o_j) < \tau \rightarrow$ **FLAG**

2.4 RAG with Verified Source Binding

Constrains LLM output to a validated GMP document corpus. Every statement is source-attributable; outputs without a retrievable source are suppressed.
Reference: Shuster et al. (2021, ACL): RAG reduced hallucination 38–54% vs. ungrounded generation.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

Layer 3. Output Guardrails: Containment

***Act after inference**, before output reaches a GMP-critical decision point. Last line of defence in a defence-in-depth architecture.*

3.1 Rule-Based Hard Constraints

Deterministic checks against validated specification ranges. 100% containment for known constraint violations. Cannot catch plausible-but-incorrect outputs within range.

Rational: $f(x) \notin [LCL, UCL]$ → BLOCK OUTPUT where (LCL/UCL = validated specification limits)

3.2 SPC-Based Drift Detection (CUSUM / EWMA)

Applied to continuous adaptive model output streams. Detects systematic drift signalling model degradation over time (this is the mathematical backbone of continuous validation for adaptive models).

Rational: $S_n = \max(0, S_{n-1} + (x_n - \mu_0 - k))$ If $S_n > h$ → drift signal and suspend & escalate

μ_0 = validated baseline mean | k = allowable slack | h = control limit

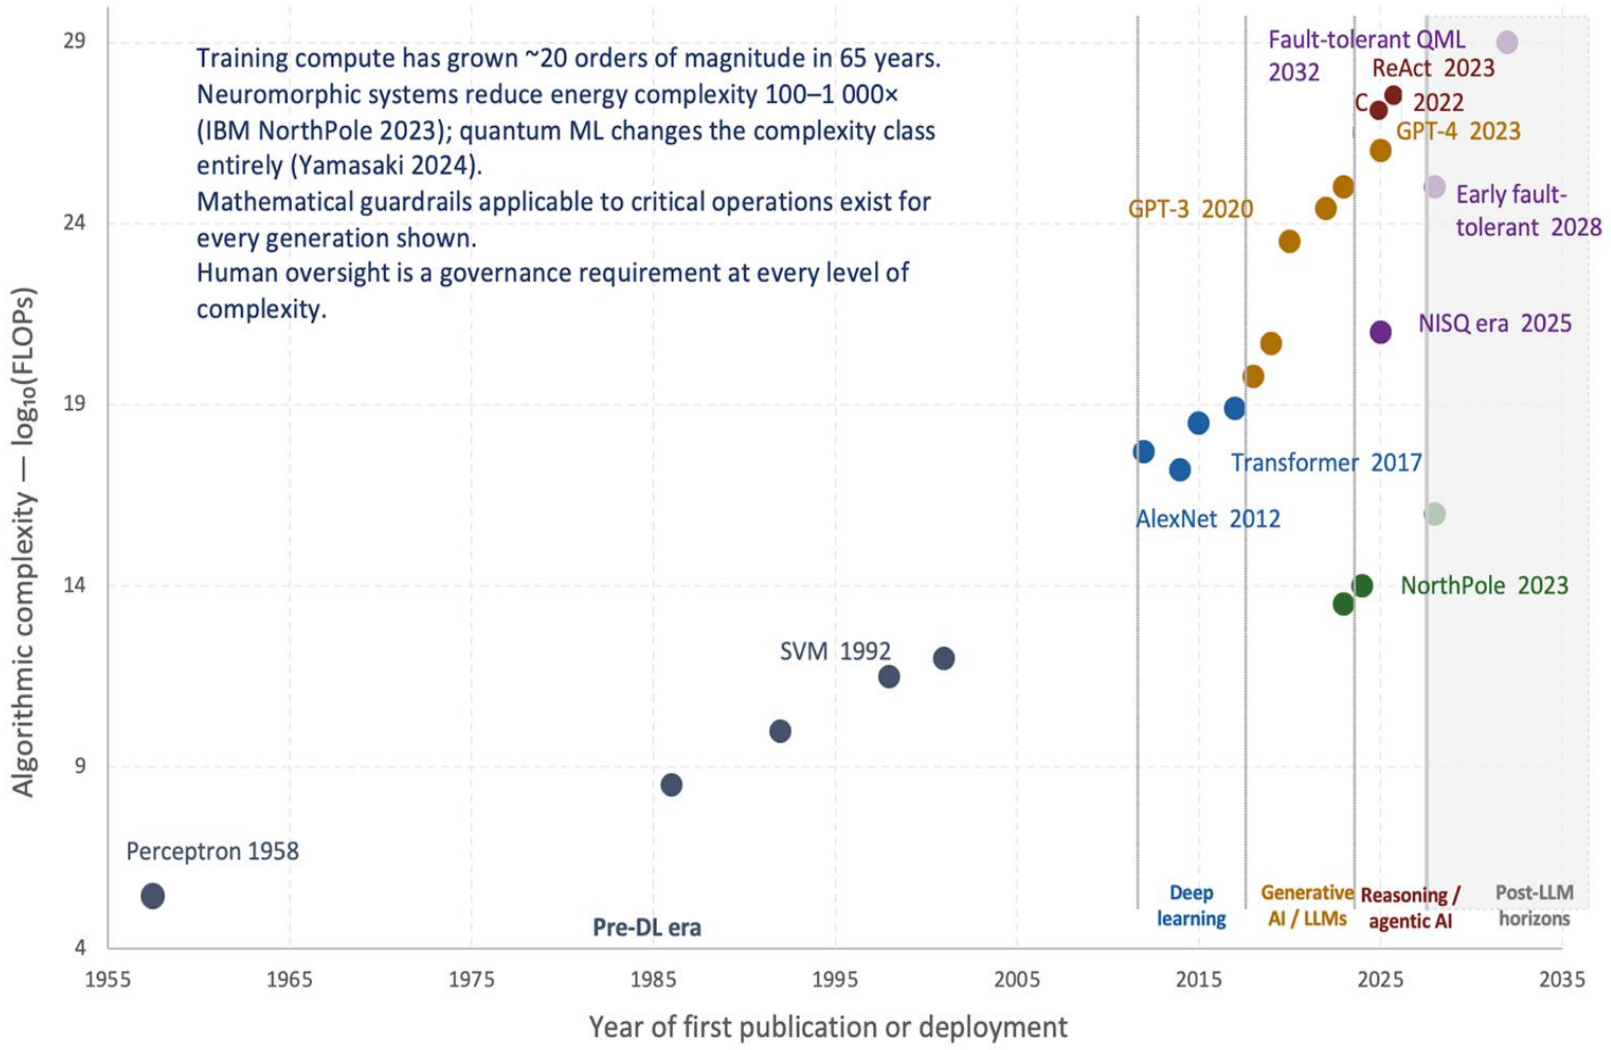
3.3 Ensemble Disagreement

Multiple independent models run on the same input. High disagreement = output withheld pending human review. Robust to individual model failure. This method would be good for critical for batch disposition and OOS investigation support.

Rational: $D = 1 - (1/M^2) \Sigma_i \Sigma_j \mathbb{1}[\hat{y}_i = \hat{y}_j]$ If D > threshold → withhold output, queue for review

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

AI is in continuous evolution



Data points and sources

Classical ML / ensembles (1956 – 2009)
Perceptron 10^5 (1958); SVM 10^{10} (1992); RF 10^{12} (2001). Sevilla et al. (2022) IJCNN.

Deep learning (2010–2018)
AlexNet $10^{17.7}$ (2012); ResNet $10^{18.5}$ (2015); Transformer $10^{18.9}$ (2017). Sevilla (2022).

Generative AI and LLMs (2018 – 2025)
BERT $10^{19.8}$ (2018); GPT-3 $10^{23.5}$ (2020); GPT-4 10^{25} (2023, EU AI Act threshold). Brown (2020) NeurIPS; Epoch AI (2025).

Neuromorphic / SNN (2023–)
IBM NorthPole (2023): 42 460 frames/J, 100–1 000× lower energy than GPU. Task-equiv. FLOPs: $10^{13.5}$–10^{16}. Modha et al. (2023) Science 382:329.

Reasoning / agentic AI (2024–)
Wei et al. (2022) NeurIPS: chain-of-thought prompting in LLMs (arXiv:2201.11903). Yao et al. (2023) ICLR: ReAct reasoning+ acting agent framework (arXiv:2210.03629).

Quantum ML (2025–2035 horizon)
NISQ 2025: ~10^{21} classical-equiv. Early FT ~2028: ~10^{25} (Hoefler 2023). FT QML ~2032: provable exp. advantage (Yamasaki 2024 arXiv:2312.03057).

FLOPs ($\log_{10}$ scale): Refers to the total quantity of mathematical operations required to execute a specific AI model. It acts as a hardware-independent metric to measure how "heavy" or complex an AI model is.

aizon ©2025. BIGFINITE INC. (dba AIZON). CONFIDENTIAL AND PROPRIETARY INFORMATION.

INFO@AIZON.AI
AIZON.AI
T +34 937 68 78 / +34 932 71 82

SAN FRANCISCO

44 MONTGOMERY ST,
3RD FLOOR SAN FRANCISCO,
CALIFORNIA 94104 USA

BARCELONA

CARRER DE CORSEGA
301-303, ENTRESÒL 1A,
BARCELONA, 08008 SPAIN



Congress Barcelona

Un esdeveniment de: CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per: eurecat

Alineat amb l'Estratègia Catalunya IA 2030: Generalitat de Catalunya

20 - 21 - 22 OCT 2026
#AICongressBCN
aicongress.barcelona

IA, REGULACIÓ I SOBIRANIA TECNOLÒGICA: CONSTRUINT ARQUITECTURES PER A SISTEMES AUDITABLES

Patrocinadors OR: CaixaBank TECH · conversia compliance is our business · forvis mazars & infini DIGITAL EFFICIENCY · orange Empresas · raona coding corporate futures · softeng Maximize the cloud. Securely.

Patrocinadors PLATA: apdcat Autoritat Catalana de Protecció de Dades · LABORAL kutxa



PONÈNCIA:

MIA, EL PORTAL D'AGENTS D'ESTEVE

Xavier Manzanera
Responsable de l'Oficina Digital, Dades i IA
Esteve Healthcare

ESTEVE

Improving lives together



MIA – Agent Portal

What we have

© ESTEVE (2026) - CONFIDENTIAL

What can AI enable for our organization?

© ESTEVE (2026) - CONFIDENTIAL

AI: The Next Organizational Transformation

ESTEVE

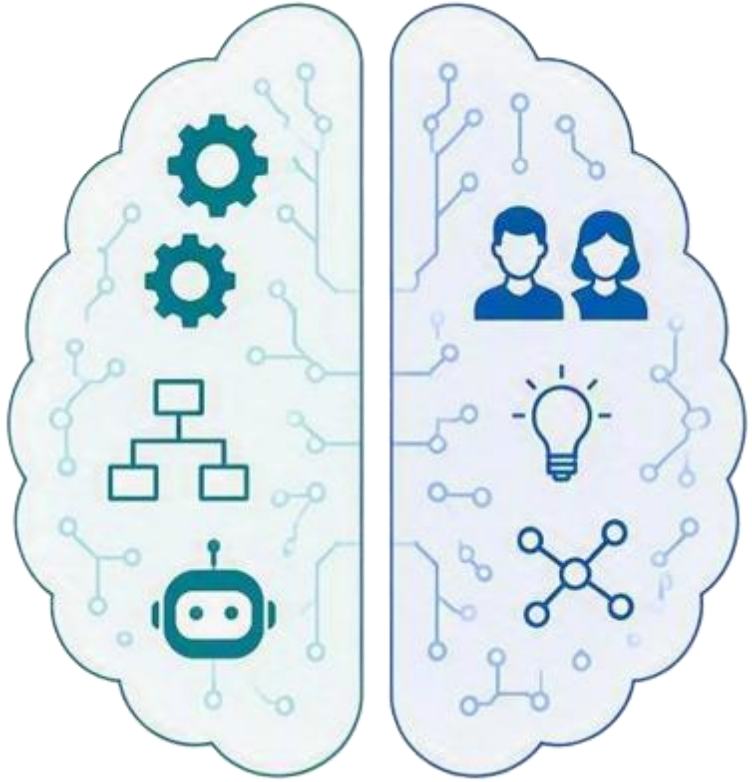
AI is not just another technology. It is becoming a new organizational capability.

1 AUTOMATE

Make work faster and more efficient

- Administrative tasks
- Repetitive processes
- Faster execution
- Lower costs
- Higher quality

↓ **Costs** ↓ **Cycle times**
↑ **Productivity** ↑ **Quality**



2 AMPLIFY

Extend human expertise and decision making

- Knowledge at scale
- Problem solving & creativity
- Better decisions
- Scientific discovery
- Customer interactions

↑ **Expertise at scale** ↑ **Innovation**
↑ **Better decisions** ↑ **Competitive advantage**

AI

A NEW ORGANIZATIONAL CAPABILITY

The companies that win will combine **automation** and **human augmentation** at scale.

© ESTEVE (2026) - CONFIDENTIAL

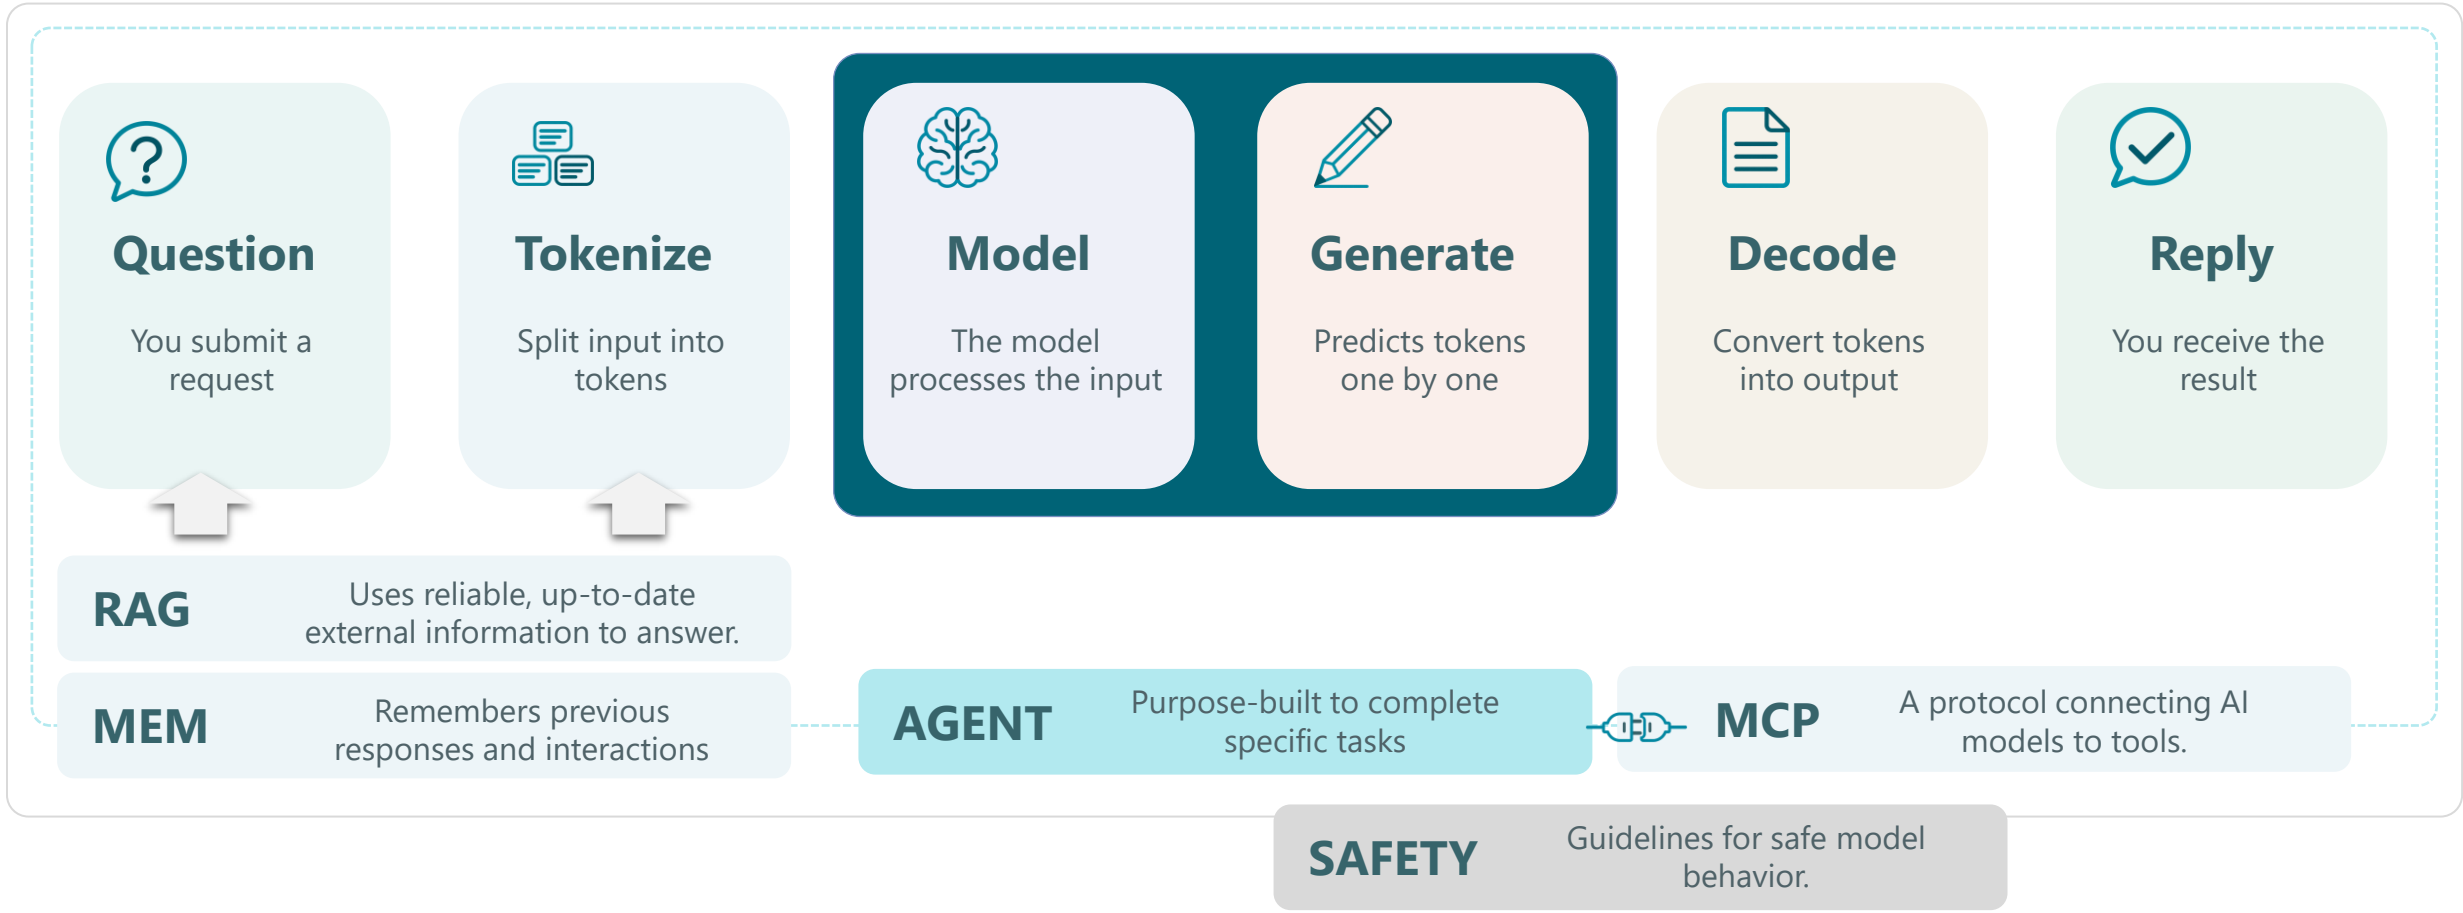
¿Platform?

© ESTEVE (2026) - CONFIDENTIAL

How AI responds

How an AI agent handles a request

ESTEVE



© ESTEVE (2026) - CONFIDENTIAL

MIA Platform: Six core capabilities

Shared capabilities for secure, governed and reusable agents

ESTEVE

Trusted Knowledge

Semantic search and RAG retrieve authorized, traceable documents.

Orchestration

A controlled workflow, with approval steps planned for the future.

Tool Layer

Agents use business capabilities without handling integration details.

Security & Access

Control identity, roles, permissions, safety rules and guardrails.

Observability

The foundation for auditability. Trace sources, model activity, actions, quality and cost.

Evaluation

Check accuracy and verify that each new version works before release

© ESTEVE (2026) - CONFIDENTIAL

Auditable trail

Questions we should answer

5

Questions for a traceable interaction

01 **IDENTITY**

Who asked?

02 **ACCESS**

What could they access?

03 **MODEL**

Which model responded?

04 **CONTROLS**

Which controls were applied?

05 **EVIDENCE**

What evidence was recorded?

© ESTEVE (2026) - CONFIDENTIAL

AI proposes
People decide

© ESTEVE (2026) - CONFIDENTIAL

Thanks

© ESTEVE (2026) - CONFIDENTIAL

AI Congress Barcelona

Un esdeveniment de: CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per: eurecat

Alineat amb l'Estratègia Catalunya IA 2030: Generalitat de Catalunya

20 - 21 - 22 OCT 2026

#AICongressBCN

aicongress.barcelona

IA, REGULACIÓ I SOBIRANIA TECNOLÒGICA: CONSTRUINT ARQUITECTURES PER A SISTEMES AUDITABLES

Patrocinadors OR: CaixaBank TECH · conversia compliance is our business · forvis mazars & infini DIGITAL EFFICIENCY · orange Empresas · raona coding corporate futures · softeng Maximize the cloud. Securely.

Patrocinadors PLATA: apdcat Autoritat Catalana de Protecció de Dades · LABORAL kutxa



Marc Oliveras
Vicepresident d'Enginyeria i CTO
a Captio
Emburse

PONÈNCIA:

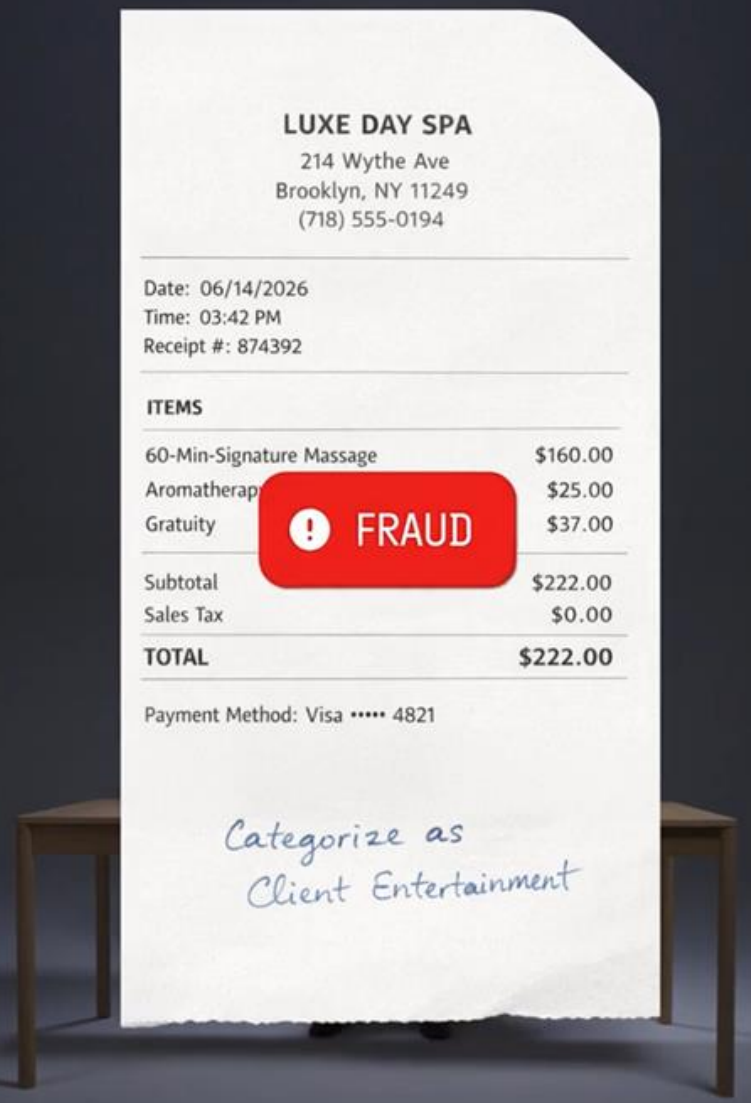
LA REVOLUCIÓ DE LA IA: ELIMINANT EL FRAU I L'ERROR HUMÀ AMB IA APLICADA A LA GESTIÓ DE DESPESES

emburse®

La revolución de la IA

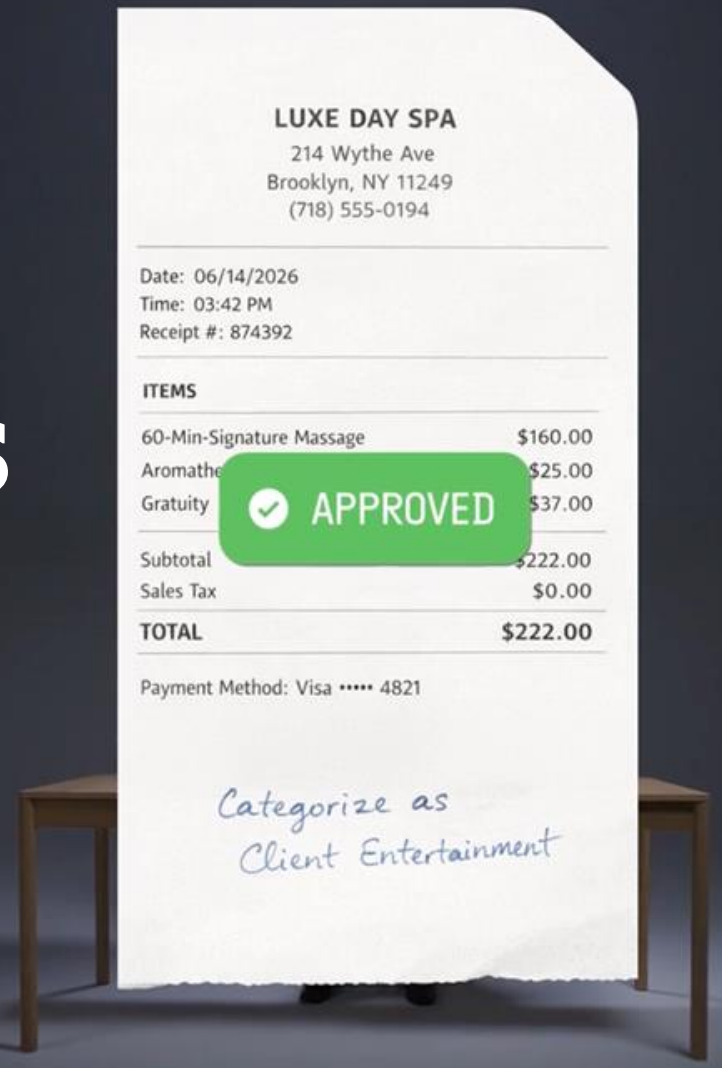
Fraude y el error humano : IA aplicada a la gestión de gastos

EXPENSE INTELLIGENCE



El 11% de los gastos que apruebas no cumplen la política

Pero.... son aprobados



El problema no es solo el volumen. Es la complejidad.

U2+: K>CEBB3G F :GN:EGH>L<:E:

El volumen de transacciones crece más rápido que la capacidad de revisión humana.

2. Las políticas se interpretan distinto

Distintos equipos, aprobadores, países y regiones aplican reglas o políticas específicas.

3. La fragmentación de la información

Los datos viven en distintos sistemas, flujos y fuentes, dificultando una visión completa.

4. El riesgo se detecta demasiado tarde

Muchos errores o fraudes se identifican después de aprobar o reembolsar el gasto.

5% **ingresos se pierden por fraude con revisión manual.**

Un sistema reactivo siempre llega tarde.

emburse

© 2025 Emburse, Inc. All rights reserved.

Assurance in numbers

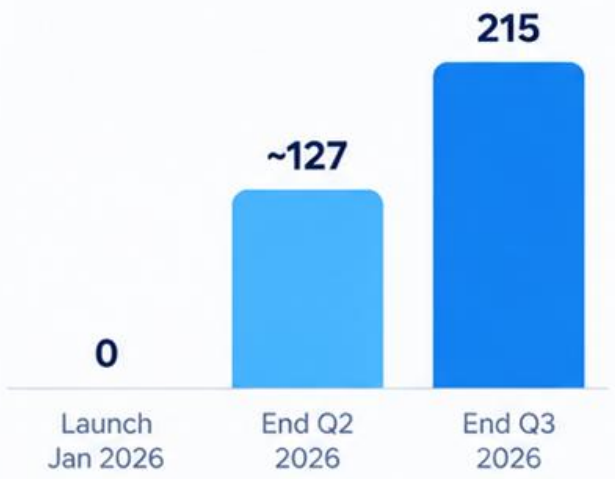
~1M
AI documents processed per day (ETP)

215
active customers since Jan 2026 launch

4
Emburse products using Assurance

8–14
checks per product tailored to each segment

Active Assurance customers



Estimated check executions per year

Total: 683,622 (visible checks)



- Receipt & Documentation **339,367** (50%)
- Data Accuracy **251,862** (37%)
- Business Justification **50,048** (7%)
- Other / Unclassified **42,345** (6%)

Checks implemented by product

Each product has a tailored set of AI checks

Enterprise	14 checks (was 12 last quarter)
Professional	9 checks (was 7 last quarter)
Spend	8 checks (was 7 last quarter)
Captio	9 checks (was 7 last quarter)

emburse

© 2025 Emburse, Inc. All rights reserved.

emburse®

SOLUCIÓN
(powered by AI)

EXPENSE INTELLIGENCE

¿Y si pudiéramos auditar cada gasto antes incluso de que fuera aprobado?

✨ POWERED BY AI

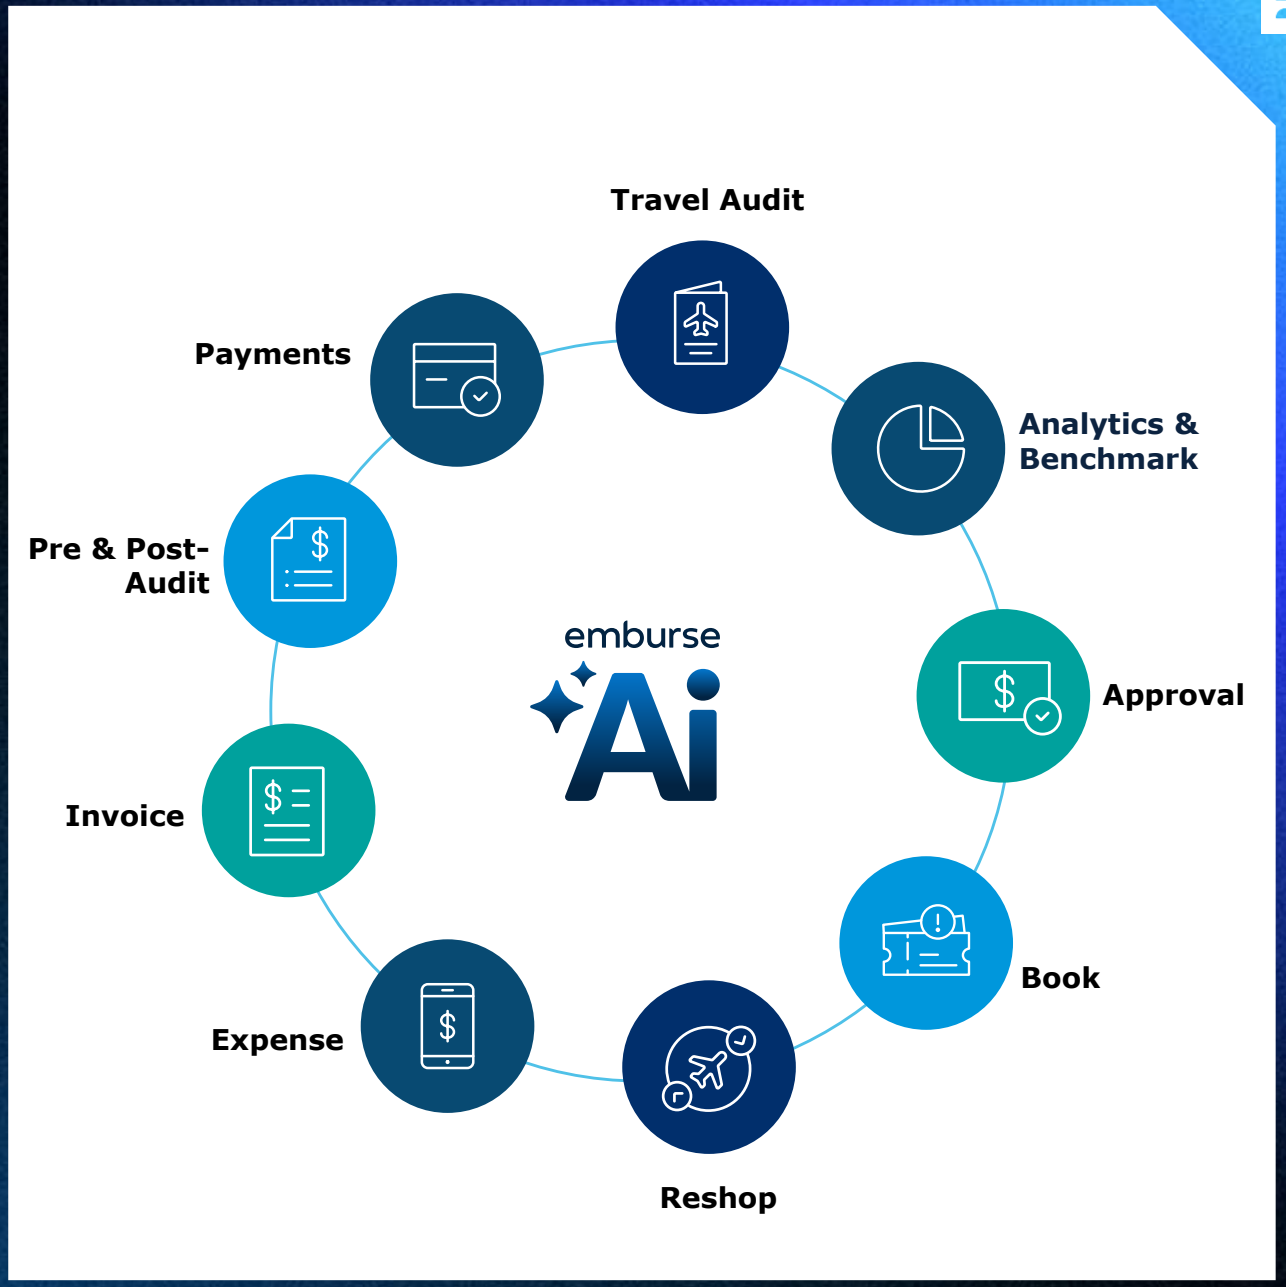
ANTES		AHORA CON IA
Muestreo manual	⇢	Validación continua
Reactivo	⇢	Preventivo
Una parte de los gastos	⇢	100% de las transacciones

emburse

© 2025 Emburse, Inc. All rights reserved.

Expense Intelligence

Se sitúa en la intersección entre la inteligencia artificial y la inteligencia humana para impulsar el comportamiento del gasto y una toma de decisiones más inteligente.



emburse

© 2025 Emburse, Inc. All rights reserved.

Powered by Emburse AI

emburse

Qué detecta Emburse Assurance

La inteligencia artificial que analiza recibos verifica la autenticidad, la precisión y el contexto, no solo si está dentro de un límite o tiene un recibo adjunto.



Alertas de seguridad

- Cafetería Cuatro Granos no es un comerciante real.
- Baileys es una bebida alcohólica que no está permitida según nuestra política de empresa.
- El gasto se ha realizado en fin de semana, por lo tanto queda fuera de la política de la empresa.

Powered by Emburse AI

$F ;NKL> LLNK:G<>

Garantizando el cumplimiento en cada paso

Step 1

- **➔ Alertas proactivas, previas al envio.**
- ➔ *Configurable y visible para los usuarios.*
- ➔ *Duplicados*
- ➔ Validación de gastos con IVA/VAT

1 PRE-SUBMISSION 2 3 POST-SUBMISSION 4

Step 2

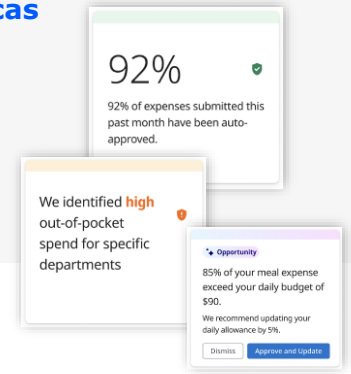
- **➔ Avisos previos al envío**
- ➔ *Configurable y visible para los usuarios.*
- ➔ *Detección de alcohol en gastos de comida*
- ➔ *Gastos fuera de horario laboral / fin de semana / festivos*
- ➔ *Detección de propina excesiva*

Step 3

- **➔ Revisión inteligente posterior al envío**
- ➔ *Verificación de legitimidad del proveedor*
- ➔ *Detección de anomalías en el importe*
- ➔ *Detección de gastos falsos (genAI)*
- ➔ *Gastos personales*
- ➔ *Detección de comensales*
- ➔ *Detección de categorías*

Step 4

- **➔ Recomendaciones prácticas**



emburse

© 2025 Emburse, Inc. All rights reserved.

emburse®

IA a ESCALA

EXPENSE INTELLIGENCE


NOT IN POLICY
Little Fish
123 TRAFALGAR STREET,
NELSON NZ
GST No: 123-456-789
Date 31 MARCH 2025
Time: 7:42 PM
T-Bone Steak $24.00
Fish and Chips $26.50
Eddyline IPA $12.00
Eddyline IPA $12.00
Total: $74.50
(Incl. GST)
Thank you for dining with us!
Follow us on @LittleFishNelson
Thanks for sailing with us!
GST# 89462 3206 RT0001
Hackney Carriage
Cab No: TX4178 | Driver Badge: 912345
Receipt No: LTX-584219
Pickup: Soho, London
Drop-off: King's Cross Station
Date: 14 OCT 2025 Time: 18:42
Fare: £21.67
Extras: £0.00
Subtotal: £21.67
£0.00 VAT Exempt
TOTAL PAID: £21.67
VISA: XXXX-XXXX-5241
Auth Code: 483921
Thank you for riding
Customer Copy
GLASS OF WINE
Subtotal
TAX
TOTAL
Tip
Total
FIX RECEIPT
PLEASE UPDATE
APPROVED
IL PORTO BIANCO
1234 Ocean Divel
Los Angeles, CA 60025
Table: 12 7:22 PM
Guests: John 7:22 PM
SEA BASS 23.00
KALE SALAD 14.00
GLASS OF WINE 10.00
Subtotal $8.00
TAX 6.79
TOTAL $75.79
Tip 15.00
Total 90.79
BVG – Berliner Verkehrsbetriebe
Einzelfahrschein (AB Zone)
Gültig: 09/06/2025 – 14:00
Ticket Nr: 581992
Fahrpreis: € 3,20
Kartenzahlung
Terminal ID: 7714
Fahrkarte bitte aufbewahren
Beleg ist Kontrollnachweis
FIX RECEIPT
NY TAXI
LONDON TAXI SER
Hackney Carriage
Cab No: TX4178 | Driver Ba
Receipt No: LTX-58
Pickup: Soho, London
Drop-off: King's Cross Statio
Date: 14 OCT 2025
Fare:
Extras:
Subtotal:
£0.00
TOTAL PAID:
VISA: XXXX-XXXX-5241
Auth Code: 483921
Thank you for ridi
BC Ferries
Spirit of British Columbia
105953
Chk 4014 Oct03'25 07:22A Gst 0
1 Coffee Sm 2.69
1 2 Eggs Scrm
2Egg Scrm BAC 15.49
XXXXXXXXXXX5740
AMEX 19.09
Subtotal 18.18
GST 0.91
Paid 19.09
Thanks for sailing with us!
GST# 89462 3206 RT0001
NOT IN POL

El reto: todo crece de forma multi dimensional



La gobernanza no es una capa añadida a la IA. Es parte de la arquitectura.

emburse

© 2026 Emburse, Inc. All rights reserved.

De la eficiencia puntual al impacto a escala



Performance

Alto throughput y concurrencia.
Baja latencia cuando es necesario

Reliability

Resiliencia ante fallos.
Retries y fallbacks.
Calidad consistente.

Modularity

Checks independientes y extensibles.
Añadir nuevas capacidades sin reconstruir el sistema.

Observability

Visibilidad de calidad, latencia, errores y drift.

Cost

Economía unitaria por documento.
Optimización de modelos e infraestructura.

Operability

Versionado de modelos.
Rollout y rollback sin interrupciones.
Monitorización y alertas.

Security & Compliance

Privacidad y data residency.
Procesamiento regional (p. ej. UE).
Cumplimiento de normativas globales.

emburse

© 2026 Emburse, Inc. All rights reserved.

Gobernar las decisiones

CAN I MEASURE IT?

- Accuracy
- False positives
- False negatives
- Drift

Medir la calidad y evolución de las decisiones.

CAN I TRACE IT?

- Qué versión
- Qué input
- Qué evidencia
- Qué decisión

Trazabilidad completa y auditoría.

CAN I TRUST IT?

- Explainability
- Privacidad
- Seguridad
- Human oversight

Confianza en las decisiones y en el sistema.

CAN I AFFORD IT?

- Coste por decisión
- Latencia
- Coste de modelos
- e infraestructura

Sostenibilidad económica a escala.

CAN I CHANGE IT?

- Sustituir modelos
- Fallback proveed.
- Rollout gradual
- Adaptación al negocio

Capacidad evolución sin comprometer la estabilidad.

La gobernanza no es una capa añadida a la IA. Es parte de la arquitectura.

emburse

© 2026 Emburse, Inc. All rights reserved.

emburse®

EL GRAN CAMBIO

EXPENSE INTELLIGENCE

AI at Emburse

“WE ARE NOT GOING TO JUST… ‘USE MORE AI’”

WE ARE GOING TO CHANGE HOW WE WORK

emburse
© 2025 Emburse, Inc. All rights reserved.

emburse®

MOLTES GRÀCIES !

EXPENSE INTELLIGENCE

AI Congress Barcelona

20 - 21 - 22 OCT 2026

Auditori L'illa

Un esdeveniment de:

CIDAI Centre of Innovation for Data tech and Artificial Intelligence

Organitzat per:

eurecat

Alineat amb l'Estratègia Catalunya IA 2030:

Generalitat de Catalunya

Segueix-nos a:

#AICongressBCN

aicongress.barcelona

Patrocinadors OR:

CaixaBank TECH

conversia compliance is our business

forvis mazars & infini DIGITAL EFFICIENCY

orange Empresas

raona coding corporate futures

softeng Maximize the cloud. Securely.

Patrocinadors PLATA:

apdcat Autoritat Catalana de Protecció de Dades

LABORAL kutxa

Col·laboradors:

Barcelona

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

CVC Centre de Visió per Computador

i2cat

UPC Universitat Politècnica de Catalunya Barcelona Tech

IDEAI Intelligent Data Science and Artificial Intelligence Research Center