

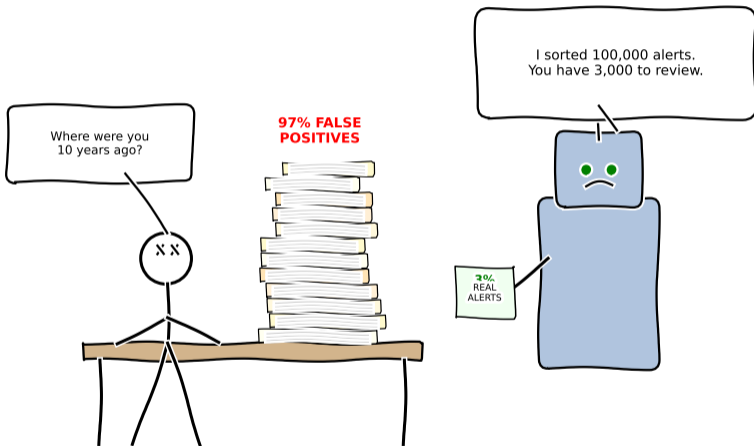
Lesson 7.2: RegTech — Automating Compliance at Scale

Module 7: The Compliance Problem

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

RegTech to the Rescue



The compliance officer was last seen smiling in 2014.

After this lesson you will be able to:

- 1 **Explain** why manual compliance does not scale for modern financial institutions
- 2 **Compare** rule-based and ML-based transaction monitoring approaches
- 3 **Design** an alert triage pipeline that reduces false positive rates
- 4 **Describe** the XBRL regulatory reporting framework and its automation benefits
- 5 **Analyze** how biometric and document-based identity verification systems work
- 6 **Evaluate** privacy-preserving techniques (differential privacy, federated learning, secure MPC) for compliance use cases
- 7 **Map** the RegTech vendor landscape by functional category

Bloom's taxonomy levels: **Understand (1), Apply (2–3), Analyze (4–5), Evaluate (6–7).**

Lesson 7.1 — The compliance landscape:

- Regulatory requirements: AML, KYC, CFT
- Cost of compliance for banks and FinTechs
- Consequences of non-compliance (fines, license loss)
- Manual processes and their limitations

Lesson 7.2 — Automating compliance:

- Transaction monitoring: rules vs. ML
- Alert optimization and false positive reduction
- Regulatory reporting automation (XBRL)
- Identity verification technology
- Privacy-preserving compliance techniques
- The RegTech vendor ecosystem

We understand the rules. Now: how do we automate compliance?

RegTech applies technology to regulatory challenges — the same way FinTech applies technology to financial services.

What Is RegTech?

Definition: RegTech (Regulatory Technology) refers to technology solutions that help financial institutions comply with regulations more efficiently, accurately, and cost-effectively than manual processes.

- The term was coined around 2015 by the UK Financial Conduct Authority (FCA)
- Global compliance spending by financial institutions exceeds \$200 billion annually (conceptual industry estimate)
- Manual compliance teams at large banks: 10,000–30,000 employees
- Key driver: regulations have grown in volume and complexity since the 2008 financial crisis
- RegTech is to compliance what FinTech is to banking: automation of costly manual processes

Core promise: Same or better compliance outcomes at a fraction of the cost and time.

The FCA Innovation Hub and the Bank of England were early champions of RegTech as a distinct sector.

RegTech: Five Functional Categories

Category	What It Does	Example Technology
Transaction Monitoring	Detect suspicious patterns in payment flows	ML anomaly detection
Regulatory Reporting	Automate submission of mandated reports	XBRL, natural language processing
Identity Verification	Verify customer identity (KYC)	Biometrics, document AI
Risk Management	Monitor and quantify compliance risk	Real-time dashboards
Regulatory Change	Track and interpret new regulations	NLP on legal text

This lesson covers the first three categories in depth, plus privacy-preserving techniques that cut across all five.

Each category addresses a different pain point in the compliance value chain. Most large banks use vendors from multiple categories.

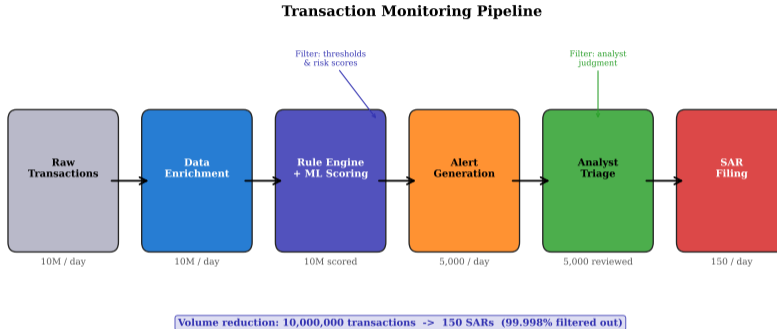
Definition: Transaction monitoring is the automated process of reviewing customer transactions to detect patterns indicative of money laundering, terrorist financing, fraud, or sanctions violations.

- Required by AML regulations in every major jurisdiction
- Processes millions of transactions per day at large banks
- Generates **alerts** when transactions match suspicious patterns
- Alerts are reviewed by human analysts who decide: *escalate* or *dismiss*
- Escalated cases become **Suspicious Activity Reports (SARs)** filed with regulators

The fundamental tension: too few alerts → missed laundering (regulatory risk); too many alerts → analyst overload (operational risk).

A typical large bank generates 50,000–200,000 alerts per month but files only 1–5% as SARs.

Transaction Monitoring Pipeline



Each stage filters out legitimate activity. The goal is high recall (catch all true positives) with manageable volume for human review.

Traditional approach: Hard-coded rules (“scenarios”) flag transactions that exceed predefined thresholds.

Common rule examples:

- Cash deposits above \$10,000 (Bank Secrecy Act threshold)
- Multiple deposits just below the reporting threshold (“structuring”)
- Rapid movement of funds through multiple accounts within 24 hours
- Transactions with sanctioned countries or entities
- Sudden increase in transaction volume relative to customer history

Advantages: Transparent, auditable, easy to explain to regulators.

Limitations:

- Criminals adapt to known thresholds → rules become stale
- Static thresholds produce excessive false positives (often >95%)
- Cannot detect novel patterns not anticipated by rule designers

Most legacy transaction monitoring systems use 100–500 pre-configured rules. Tuning these rules is a full-time job.

Modern approach: Machine learning models learn patterns from historical data to score transactions by suspicion level.

Supervised models:

- Trained on labeled SARs (positive) and cleared alerts (negative)
- Random Forest, Gradient Boosting, Neural Networks
- Output: probability score per transaction
- Strength: high precision on known patterns

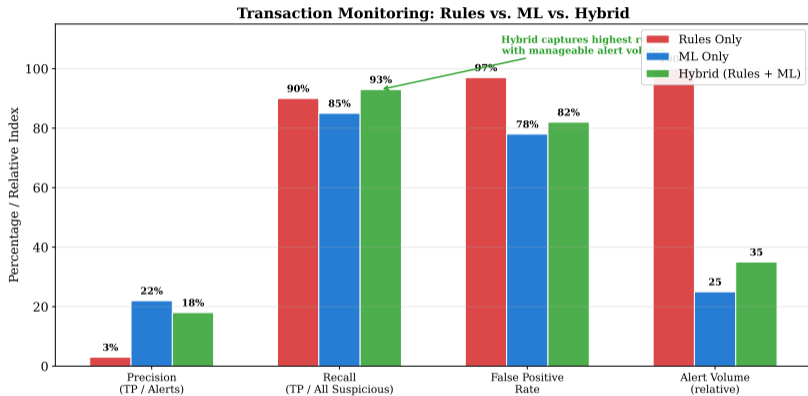
Unsupervised models:

- Detect anomalies without labeled training data
- Isolation Forest, Autoencoders, Clustering
- Output: anomaly score per transaction or customer
- Strength: can find novel, previously unseen patterns

Hybrid approach (most common): Rules for regulatory mandates (e.g., the \$10,000 threshold) + ML for everything else.

ML models can reduce false positives by 40–70% while maintaining the same true positive detection rate (conceptual benchmark).

Rule-Based vs. ML-Based Detection



ML models trade some interpretability for substantially better precision. Hybrid systems combine the strengths of both.

The False Positive Problem

The #1 operational challenge in transaction monitoring: false positives.

Scale of the problem:

- Rule-based systems: 95–99% false positive rate
- Each alert requires 30–60 minutes of analyst review
- Cost per alert investigation: \$20–\$50 (conceptual estimate)
- A bank with 100,000 monthly alerts and 97% false positives wastes resources on 97,000 non-suspicious cases

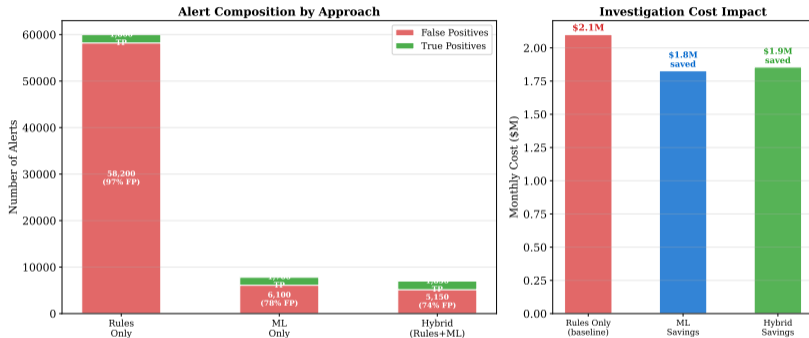
Key insight: Reducing false positives is not just about efficiency — it directly improves the detection of real financial crime.

Consequences:

- Analyst fatigue → true positives are missed
- Compliance teams are understaffed relative to volume
- Regulatory expectation: every alert must be reviewed
- Banks cannot simply raise thresholds without regulatory approval

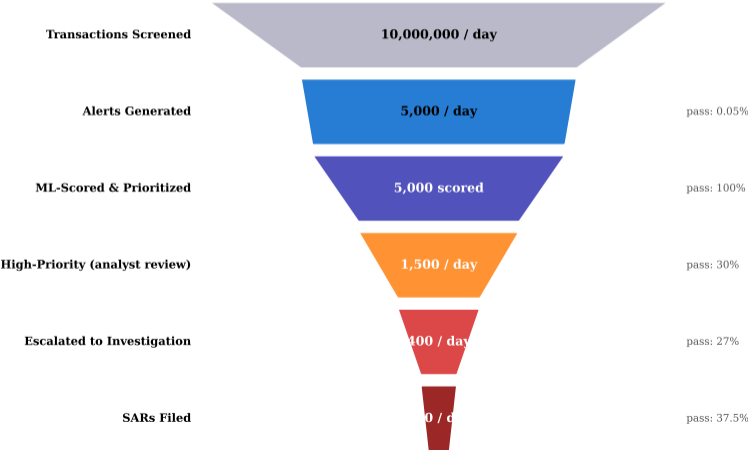
“Alert fatigue” is a recognized risk factor. Regulators increasingly accept ML-based prioritization to address it.

False Positive Reduction: Rules vs. ML vs. Hybrid



Even modest false positive reduction translates to millions of dollars in annual savings for large institutions.

Alert Triage Funnel: Transaction Monitoring



How to reduce false positives without missing true positives:

- 1 **Customer risk segmentation:** Apply different thresholds to high-risk vs. low-risk customer segments
- 2 **Contextual scoring:** Enrich alerts with customer history, peer group behavior, and external data
- 3 **Network analysis:** Examine transaction counterparties, not just individual transactions
- 4 **Above-the-line / below-the-line:** ML scores prioritize the alert queue; low-scoring alerts get expedited review
- 5 **Feedback loops:** Analyst decisions on past alerts retrain the ML model continuously
- 6 **Scenario tuning:** Periodically review rule thresholds using historical true/false positive data

Regulatory requirement: Any optimization must be documented and justified. Regulators expect evidence that true positive rates are maintained.

“Model risk management” (SR 11-7 in the US) applies to ML models used in compliance just as it does to credit models.

Definition: Regulatory reporting is the mandatory submission of structured financial data to supervisory authorities on a periodic or event-driven basis.

- Banks submit hundreds of distinct reports per jurisdiction per year
- Report types: capital adequacy (Basel III), liquidity (LCR), transaction reporting (MiFID II / EMIR), tax (FATCA / CRS)
- Manual reporting involves data extraction, transformation, validation, and submission
- Error rates in manual reporting: 5–15% of submissions require restatement (conceptual estimate)
- A single reporting error can trigger regulatory investigation and fines

Cost: Large banks spend \$50–100 million annually on regulatory reporting infrastructure.

The EU alone introduced over 50,000 pages of financial regulation between 2009 and 2020 (approximate industry estimate).

Definition: **XBRL** (eXtensible Business Reporting Language) is an open, XML-based standard for tagging financial data so it can be read, validated, and analyzed by machines.

How XBRL works:

- Each financial fact (e.g., “Total Assets”) is tagged with a machine-readable label from a taxonomy
- Taxonomies define the allowed tags, their relationships, and validation rules
- Reports are submitted as structured XBRL documents, not PDFs or spreadsheets

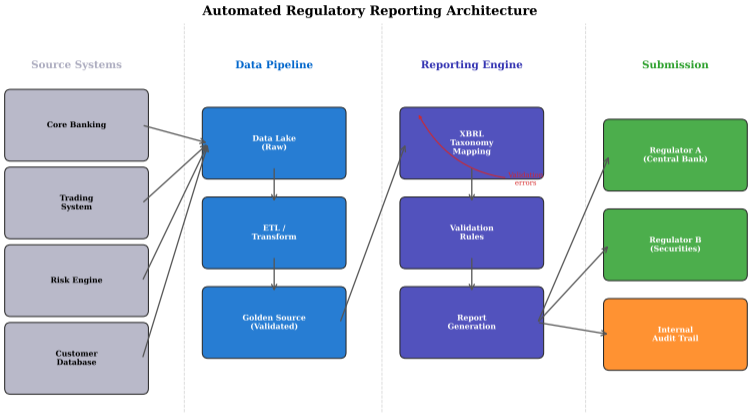
Adoption: Mandated by the US SEC (since 2009), European Banking Authority (EBA), Bank of England, and 60+ regulators worldwide.

Benefits:

- Automated validation at submission time
- Machine-comparable data across institutions
- Reduced manual data entry errors
- Faster regulatory analysis and response

XBRL eliminates “PDF to spreadsheet” manual transcription, which is the largest source of reporting errors.

Automated Regulatory Reporting Architecture



Automation replaces manual extract-transform-load (ETL) with continuous, validated data pipelines.

Definition: Digital identity verification is the process of confirming a person's identity remotely using technology rather than in-person document inspection.

Three pillars of modern KYC verification:

- 1 **Document verification:** AI-powered extraction and authentication of identity documents (passport, national ID, driving license)
- 2 **Biometric verification:** Matching a live selfie or fingerprint to the photo on the verified document
- 3 **Liveness detection:** Confirming the person is physically present (not a photo, video replay, or deepfake)

Method	Speed	Fraud Resistance
Manual document review	1–3 days	Low–Medium
Automated document + biometric	30–90 seconds	Medium–High
Document + biometric + liveness	60–120 seconds	High

Automated KYC can reduce onboarding time from days to minutes while improving fraud detection accuracy.

Digital Identity Verification Flow



Each step reduces the probability of a fraudulent identity passing through. Liveness detection is critical against deepfakes.

Problem: Fraudsters attempt to bypass biometric checks using photos, pre-recorded videos, or AI-generated deepfakes.

Two types of liveness detection:

Active liveness:

- User performs a challenge: blink, turn head, smile, read a number
- Verifies real-time responsiveness
- Higher friction (takes 5–10 seconds)
- Harder to spoof with static images

Passive liveness:

- Analyzes a single selfie for micro-textures, depth cues, reflection patterns
- No user action required (frictionless)
- Uses convolutional neural networks (CNNs)
- Vulnerable to advanced deepfakes without additional signals

Current best practice: Combine passive liveness (for low-risk onboarding) with active liveness (for high-risk transactions or step-up verification).

The ISO/IEC 30107-3 standard defines Presentation Attack Detection (PAD) levels for biometric liveness testing.

Tension: Effective compliance requires analyzing customer data. Privacy regulations (GDPR, CCPA) restrict how that data is collected, stored, and shared.

Compliance needs data for:

- Transaction monitoring across accounts
- Cross-institutional information sharing (e.g., typology networks)
- Customer risk profiling
- Regulatory reporting of individual transactions

Privacy regulations restrict:

- Data minimization (collect only what is necessary)
- Purpose limitation (use data only for stated purposes)
- Cross-border data transfers
- Profiling without explicit consent

Solution direction: Privacy-preserving computation techniques allow compliance analysis without exposing raw personal data.

The GDPR (Article 6(1)(c)) provides a legal basis for processing personal data when required by law, but the principle of data minimization still applies.

Definition: Differential privacy (DP) is a mathematical framework that adds calibrated noise to data or query results so that the output does not reveal whether any individual's data was included.

- **Core idea:** The result of an analysis should be approximately the same whether or not any single individual's record is in the dataset
- **Privacy parameter ϵ (epsilon):** Controls the privacy–utility trade-off. Lower ϵ = stronger privacy but noisier results
- **Mechanism:** Add random noise (Laplace or Gaussian) calibrated to the query's sensitivity


Compliance application: A bank can share aggregate transaction statistics with regulators or industry utilities without exposing individual customer transactions.

Example: Reporting the average transaction amount per customer segment \pm noise, rather than individual transaction records.

Differential privacy was pioneered by Cynthia Dwork (2006) and is now used by Apple, Google, and the US Census Bureau.

Differential Privacy: How Noise Protects Individuals

Differential Privacy: The Privacy-Utility Trade-off



The noise is calibrated so that aggregate statistics remain useful while individual records cannot be reverse-engineered.

Definition: Federated learning (FL) is a machine learning technique where a model is trained across multiple institutions without sharing raw data. Only model updates (gradients) are exchanged.

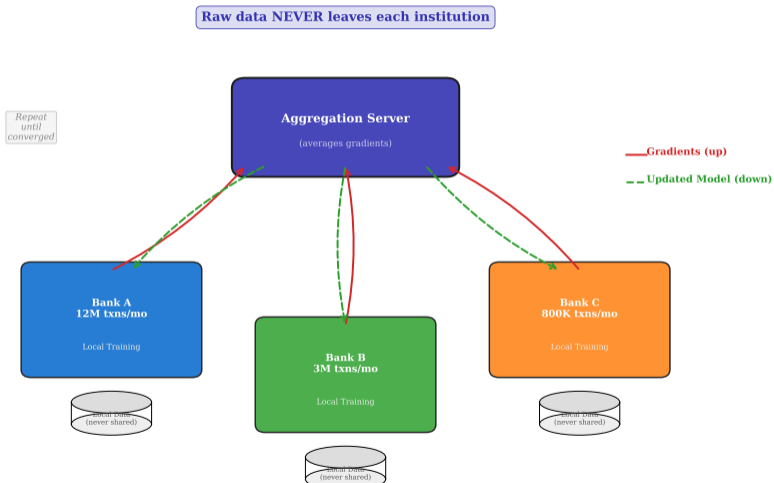
- 1 A central server distributes a global model to participating banks
- 2 Each bank trains the model on its local transaction data
- 3 Banks send only the model weight updates (not data) back to the server
- 4 The server aggregates updates and distributes an improved global model
- 5 Repeat until the model converges

Compliance application: Multiple banks collaboratively train an AML detection model that benefits from cross-institutional patterns — without any bank seeing another bank's customer data.

Key benefit: Catches typologies that span multiple institutions (e.g., layering across different banks) without data pooling.

Federated learning was introduced by Google (McMahan et al., 2017) for mobile keyboard prediction and has since been adapted for finance.

Federated Learning for Cross-Institutional AML Detection



Raw data never leaves each institution. The aggregation server sees only model parameters, not customer information.

Definition: **Secure MPC** is a cryptographic protocol that enables multiple parties to jointly compute a function over their combined inputs without revealing those inputs to each other.

How it works (simplified):

- Each party “secret-shares” its input
- Computation is performed on encrypted shares
- Only the final result is revealed
- No party learns anything beyond the output

Compliance applications:

- Sanctions screening across banks without sharing customer lists
- Joint fraud detection on shared transaction graph
- Privacy-preserving identity deduplication
- Aggregate risk reporting to regulators

Trade-off: MPC provides the strongest privacy guarantees but is computationally expensive. Practical for batch processing, challenging for real-time use.

Secure MPC was first theorized by Andrew Yao (1982). Modern implementations are fast enough for many financial applications.

Property	Differential Privacy	Privacy	Federated Learning	Secure MPC	Homomorphic Encryption
Data leaves institution?	Noisy aggregates only		No (gradients only)	No (shares only)	No (ciphertext only)
Computational cost	Low		Medium	High	Very High
Accuracy impact	Noise reduces precision		Minimal	None	None
Best use case	Aggregate reporting		Collaborative ML	Joint computation	Cloud computation

Practical reality: Most RegTech deployments combine techniques — for example, federated learning with differential privacy applied to the gradient updates.

Homomorphic encryption (computing on encrypted data) is included for completeness but remains too slow for most real-time compliance workloads.

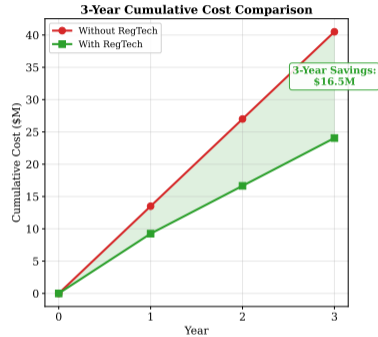
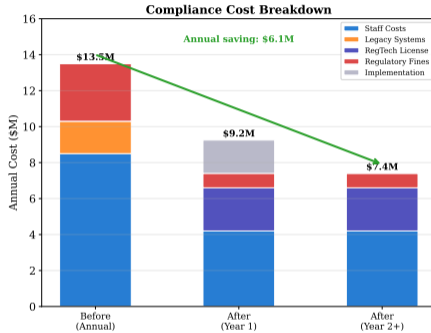
RegTech Vendor Landscape by Category

Total RegTech Market: ~\$15B | Growth: ~22% CAGR



The RegTech market is estimated at \$10–15 billion globally and growing at 20–25% annually (conceptual industry estimates).

The Business Case: Compliance Automation ROI



ROI is driven primarily by headcount reduction in alert review and reporting, plus avoided regulatory fines.

Technical challenges:

- Legacy core banking systems resist integration
- Data quality: inconsistent formats, missing fields, duplicate records
- Model validation: regulators require explainability for ML-based decisions
- Real-time processing at scale (millions of transactions per day)
- Cross-border data residency requirements

Organizational challenges:

- Compliance culture favors conservatism over innovation
- Regulator skepticism toward “black box” ML models
- Vendor lock-in concerns with proprietary platforms
- Change management: retraining compliance analysts
- Budget competition with revenue-generating projects

Key insight: Technology is rarely the bottleneck. Organizational readiness and regulatory acceptance are the harder problems.

Successful RegTech adoption requires close collaboration between compliance officers, data scientists, and regulators.

Question: Will regulators accept ML-based compliance decisions?

Current state by jurisdiction:

- **US (FinCEN / OCC):** Supportive in principle. Joint statement (2018) encourages “innovative approaches” to AML compliance. Requires model risk management (SR 11-7).
- **EU (EBA):** Published guidelines on ML use in AML/CFT (2023). Emphasizes explainability, human oversight, and ongoing validation.
- **UK (FCA):** “TechSprints” actively test RegTech solutions. Open to ML if governance is adequate.
- **Singapore (MAS):** Actively promotes RegTech via sandbox programs and industry grants.

Common requirements across jurisdictions:

- ① Explainability — regulators must understand why an alert was (or was not) generated
- ② Human-in-the-loop — final SAR filing decisions must involve a human
- ③ Ongoing monitoring — model performance must be tracked and documented

No regulator has banned ML in compliance. The trend is toward cautious acceptance with strong governance requirements.

- 1 **RegTech** automates compliance across transaction monitoring, regulatory reporting, identity verification, risk management, and regulatory change tracking
- 2 **Transaction monitoring** generates massive alert volumes; ML reduces false positives by 40–70% vs. rule-based systems alone
- 3 **Alert triage** combines ML scoring with human review to ensure both efficiency and regulatory acceptance
- 4 **XBRL** standardizes regulatory reporting into machine-readable formats, eliminating manual transcription errors
- 5 **Biometric verification** with liveness detection enables fast, fraud-resistant KYC that replaces in-person document checks
- 6 **Privacy-preserving techniques** (differential privacy, federated learning, secure MPC) resolve the tension between compliance data needs and privacy regulations
- 7 The RegTech market is large and growing, but **organizational change and regulatory acceptance** remain the primary adoption barriers

RegTech is not optional — it is how compliance will work at scale as transaction volumes and regulatory complexity continue to grow.

This lesson:

- Defined RegTech and its five functional categories
- Compared rule-based and ML-based transaction monitoring
- Analyzed the false positive problem and alert optimization strategies
- Explained XBRL-based regulatory reporting automation
- Examined biometric identity verification with liveness detection
- Evaluated privacy-preserving techniques for cross-institutional compliance
- Surveyed the RegTech vendor landscape and implementation challenges

Next lesson — Lesson 7.3:

- Deep dive into AML case studies and typology detection
- Network analysis for uncovering layered laundering schemes
- Real-world SAR filing workflows and regulatory feedback loops

With the technology toolkit established, we next apply it to real compliance scenarios and measure outcomes.