

Beyond Observability: Proactive SRE for AI-Driven Personalisation

What telco churn models and grocery promo engines
taught me about personalising \$50M portfolios



Yury Lysak
Head of Strategic Projects @ Lionsoul Global

Conf42 SRE 2026
March 19, 2026

About Me: Personalisation Across Three Industries



Veon (Telco) · N = 50 Million

- NBA engines, churn prediction, ARPU uplift
- Real-time decisioning - sub-second offer selection

- Scale: millions of subscribers, daily interactions



McKinsey (Consulting) · N = 5 Million

- Personalised promotions, loyalty optimisation
- Basket analysis, cross-sell engines for top retailers

- CIS & Southeast Asia - grocery, consumer electronics



Lionsoul Global (WealthTech) · N = 50

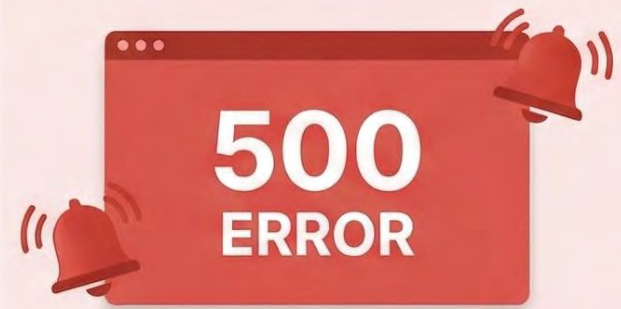
- HNWI digital-asset wealth management
- BTC Alpha Fund, Market Neutral, ETH Yield, DeFi strategies
- Personalised allocation & risk management per client

Personalisation: N=50M > N=5M > N=50

"Personalisation is the highest-ROI commercial lever in any industry - but the approaches must adapt radically when your client base is 50, not 50 million."

The Silent Killer: Why AI Personalisation Fails Differently

Loud Failures



- Traditional system failures:
 - 500 errors, timeouts, crashes
 - Alerts fire immediately
 - Everyone knows it's broken
 - MTTR measured in hours

Silent Failures



Revenue



Model drift: LTV drops 2% per month. No alerts fire.

- AI personalisation failures:
 - Model drift, feature decay, stale embeddings
 - No alerts - all dashboards green
 - Revenue erodes 2% per month
 - Nobody notices for a quarter
 - “Slowly boiling frog”

Black Swan Events



October 2025 ADL Cascade

\$19B liquidated · 9× largest ever · \$3.21B in 60 seconds
Models trained on normal conditions > dangerous recommendations during cascade

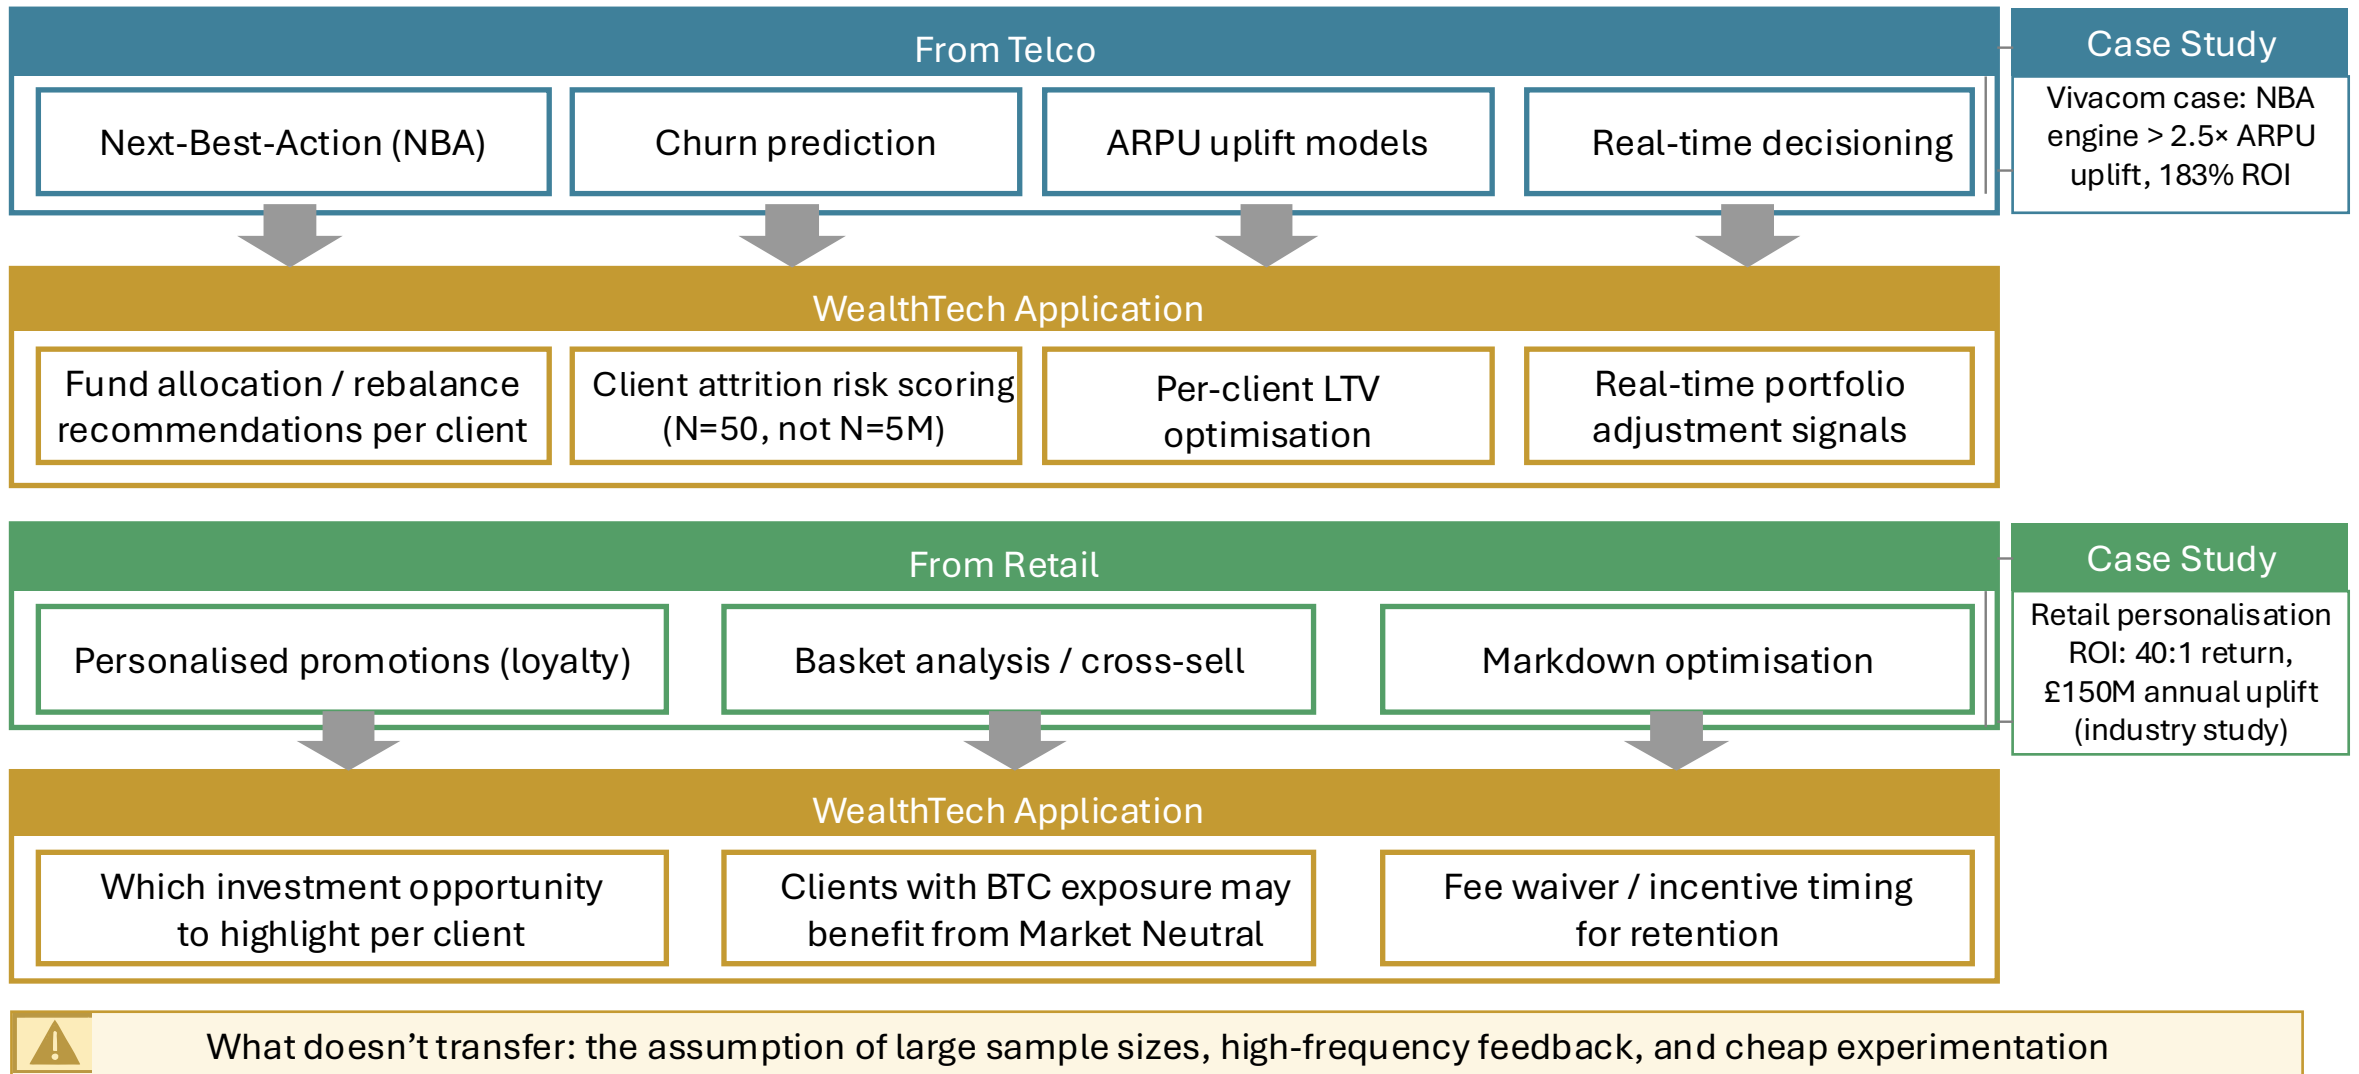


February 2026 BTC Crash

-44% drop · -6.05 σ move · \$2T market cap wiped
Personalisation models assuming normal liquidity > “buy the dip” into a void

Observability tells you the system is ‘up’ - not that it’s still making good decisions.

Personalisation Models That Transfer: Telco & Retail > WealthTech



Small-Data Personalisation: When N=50, Not N=50 Million


What Breaks

✗ **A/B testing** - can't split 50 clients into significant cohorts

✗ **Frequentist p-values** - meaningless at small N

✗ **Large training datasets** - doesn't exist

✗ **Fast feedback loops** - quarterly reviews, not daily transactions



What Works

✓ **Bayesian approaches**
Prior distributions from telco/retail + sequential updating with each client interaction

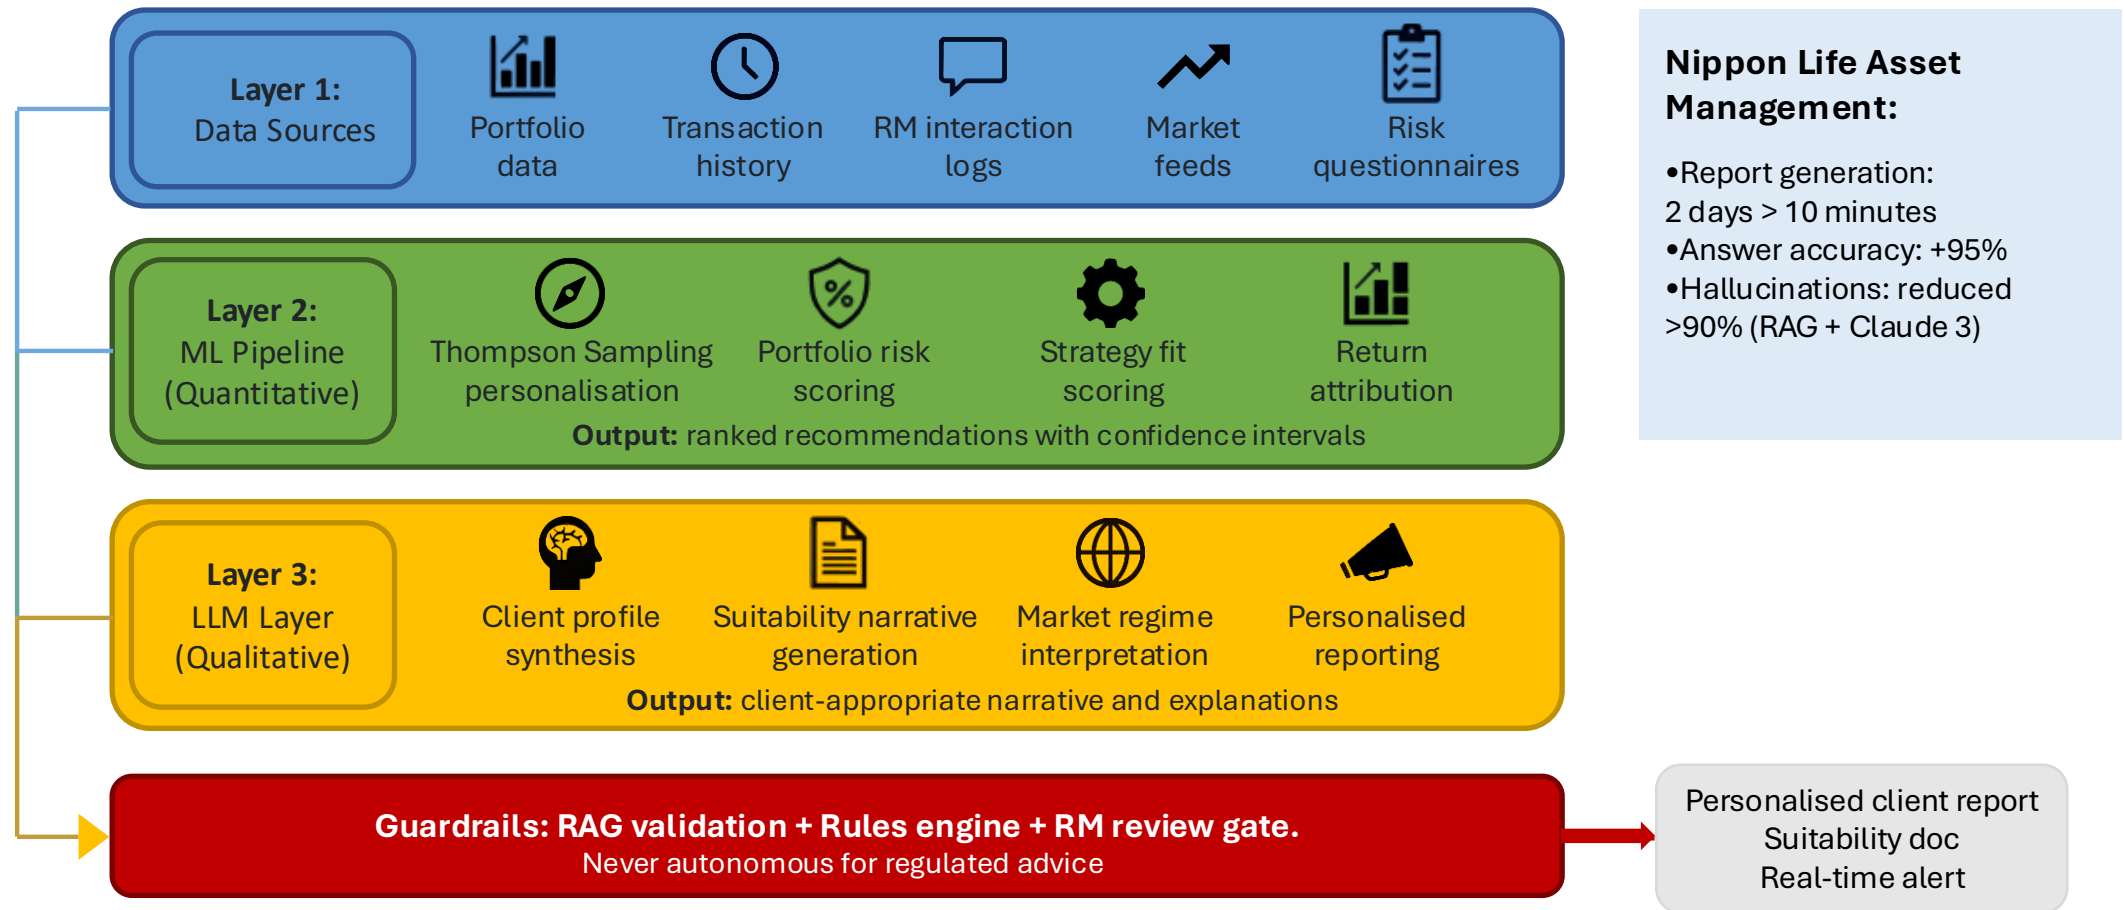
✓ **Thompson Sampling**
Explore vs. exploit - works with tiny

✓ **RM as Human Sensor**
Structured feedback: "Client seemed uncomfortable with crypto exposure"
Each RM observation = thousands of

✓ **Transfer Learning**
Pre-train on telco/retail data (N=millions),

LLMs are the great equaliser: with 50 clients, you can analyse each client's full history - something impossible at telco scale.

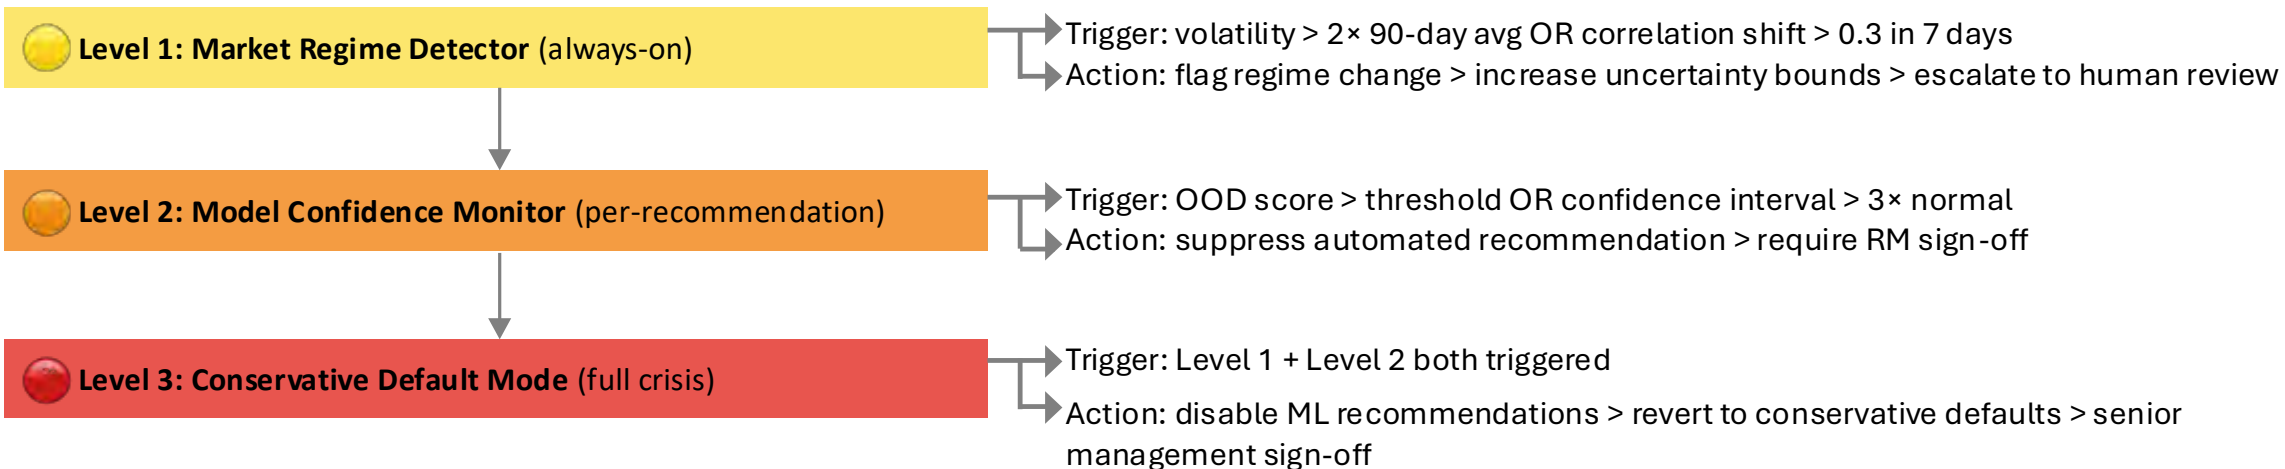
LLMs as the Personalisation Engine for Small-Data WealthTech



ML tells you **WHAT** to recommend.


LLM tells you **HOW** to explain it in terms that resonate with each individual client.

Black Swan Resilience: When Markets Break Your Models



BTC-S&P 500 Rolling 90-Day Correlation

BTC Volatility vs. VIX Correlation: 0.88 - highest ever



Diversification assumption broken.

Silence is better than a confident wrong answer. In a black swan, "I don't know" is more valuable than a hallucinated recommendation.

Proactive Signals: What to Monitor and How



Tier 1 - Always-On (Daily)

- **Market regime:** crypto volatility (24h/7d), BTC-equity correlation, ADL metrics



- **Data quality:** completeness, price feed freshness, RM feedback rate




- **Model confidence:** avg prediction score, % below confidence threshold



Tier 2 - Weekly Checks

- **PSI (Population Stability Index)** per input feature - threshold: >0.25 = retrain



- **JS divergence** for recommendation output distribution



- **RM override rate:** if $>30\%$, model is misaligned



- **Out-of-distribution (OOD)** detection score



Tier 3 - Monthly Business Metrics


- **ARPU uplift** vs. non-personalised benchmark



- **Strategy uptake rate** by recommendation source




- **Client satisfaction** signals (RM qualitative + NPS)




Every alert must answer: “How much revenue is at risk, for which clients, right now?”

Compliance Automation: Audit Trails for AI Decisions



1 AI Recommendation Generated

Logged: inputs, features, model version, confidence, SHAP values



2 Compliance Validation (automated)


Rules engine checks: MiFID suitability ✓ Concentration limits ✓ Risk tolerance ✓ Client restrictions ✓ . If any check fails > recommendation blocked, compliance team alerted



3 LLM Explanation Generated



LLM-Generated Explanation: This recommendation of Market Neutral Fund was driven primarily by the client's stated preference for capital preservation (weight: 0.42) and current overexposure to crypto (weight: 0.31)...



4 RM Review & Approval

RM reviews AI recommendation + LLM explanation > approves / modifies / rejects



5 Immutable Audit Log

Full decision trail stored: recommendation + inputs + explanation + RM action + timestamp.
Retained 5+ years (MiFID II requirement)



Manual regulatory audit response: **4–6 weeks**



AI-assisted audit response: **2–4 hours**



Key Message Regulators will ask: “Why was this recommended to this client at this time?”
Your system must answer in minutes, not weeks.

Real Impact: Business Metrics That Matter



20–25%

LTV Uplift

From proactive personalisation + drift detection preventing silent degradation



30%

Higher Client Retention

Deloitte: firms investing >10% of budget in AI personalisation (vs. industry avg)



40–70%

MTTR Reduction

Model issues: from weeks (quarterly discovery) → hours (proactive alerting)



2–4 hours

Audit Response Time

Regulatory inquiry → complete decision trail (vs. 4–6 weeks manual process)



>10:1

ROI Ratio

\$150–300K engineering investment vs. \$750K+ per-event revenue risk protected

Investment Cost

\$150–300K

Revenue Risk Protected per Event

\$750K+

The ROI Story

- One mis-recommendation to 10 clients during a black swan:
 - Potential AUM outflow = **\$50M** → Fee revenue lost = **\$750K/year**
- One MiFID II compliance breach:
 - Fines up to €5M or 10% of global turnover + reputational damage
- Proactive SRE investment: \$150–300K/year
- Revenue risk mitigated: multiple of investment

Biggest ROI came from catching issues early - not from building better models.

Key Takeaways




Personalisation is the highest-ROI commercial lever

The models from telco and retail transfer to WealthTech - but must adapt for small data



LLMs are the great equaliser

Enterprise-grade personalisation with boutique-sized client bases - qualitative reasoning over unstructured data



AI personalisation fails silently

Build model health signals (drift, uplift decay, regime mismatch) alongside traditional golden signals



Black swans will happen - build circuit breakers

Fall back to conservative defaults when conditions exceed model training bounds (Oct 2025 ADL+ Feb 2026 BTC crash = your proof points)



Compliance is a feature, not a burden

Automated audit trails + LLM explanations reduce risk AND build client trust



Cross-industry experience is your superpower

Telco NBA engines, retail promo optimisation, churn models - directly applicable frameworks

Yury Lysak · Head of Strategic Projects @ Lionsoul Global | LinkedIn | yury.lysak@live.com

