AML BASE LLC
PROFESSIONAL FINANCIAL CRIME TRAINING | 94+ COUNTRIES

Alert & Case Generation Framework

Behavioural Detection, Alert Scoring & Case Investigation on NICE Actimize

A US Tier 1 Bank Reference Implementation

PLATFORM	NICE Actimize — SAM, WLF, CDD/ERCM, ActOne ECM, Policy Manager
COVERAGE	AML transaction monitoring, alert generation, alert scoring, case management, SAR
REGULATORY BASIS	BSA / 31 CFR Chapter X · FFIEC BSA/AML Examination Manual · FinCEN · OFAC · SR 11-7
CLASSIFICATION	Internal Use Only — Financial Intelligence Unit
VERSION	1.0 — Effective July 2026 — Annual review cycle

Prepared by AML Base LLC for professional training and reference purposes.
Illustrative institution: 'Meridian National Bank, N.A.' — a hypothetical US Tier 1 bank.

Document Control

Attribute	Detail
Document title	Alert & Case Generation Framework — NICE Actimize Transaction Monitoring
Illustrative institution	Meridian National Bank, N.A. (MNB) — hypothetical US Tier 1 bank, ~$450bn assets, OCC-supervised
Document owner	BSA Officer / Head of Financial Intelligence Unit (FIU)
Contributing functions	FIU Operations · AML Technology · Model Risk Management · Compliance Testing · Internal Audit (observer)
Approval authority	BSA/AML Steering Committee, ratified by the Board BSA & AML Committee
Review cycle	Annual, or upon: material system upgrade, new product/LoB launch, adverse regulatory finding, or material change in the Enterprise-Wide AML Risk Assessment (EWRA)
Classification	Internal Use Only. Contains detection thresholds — unauthorised disclosure could enable evasion. Distribution is need-to-know and logged.
Retention	Current version plus all superseded versions retained for a minimum of five (5) years per 31 CFR 1010.430
Version	1.0 — Effective July 2026

Acronyms and Abbreviations

Term	Meaning	Term	Meaning
ADM	Actimize Data Model	LoB	Line of Business
AIS	Actimize Integrated Server	MI	Monetary Instrument (cheque, money order, cashier's cheque, traveller's cheque)
ATL / BTL	Above-the-Line / Below-the-Line (threshold testing)	ML / TF	Money Laundering / Terrorist Financing
BSA	Bank Secrecy Act	MRM	Model Risk Management (SR 11-7 / OCC 2011-12)
CDD / EDD	Customer / Enhanced Due Diligence	MSB	Money Services Business
CIP	Customer Identification Program	NFA	No Further Action (alert disposition)
CRR	Customer Risk Rating	OFAC	Office of Foreign Assets Control
CTR	Currency Transaction Report (FinCEN Form 112)	PEP	Politically Exposed Person
DCN	Document Control Number (BSA E-Filing receipt)	RCM	Risk Case Manager (legacy Actimize case module; superseded by ActOne)
DQ	Data Quality	RFI	Request for Information (to the business / branch / RM)
ECM	Enterprise Case Management (ActOne)	SAM	Suspicious Activity Monitoring (Actimize AML detection engine)
ER	Entity Resolution	SAR	Suspicious Activity Report (FinCEN Form 111)
EWRA	Enterprise-Wide AML Risk Assessment	SIU	Special Investigations Unit
FFIEC	Federal Financial Institutions Examination Council	SLA	Service Level Agreement
FIU	Financial Intelligence Unit	TBML	Trade-Based Money Laundering
FinCEN	Financial Crimes Enforcement Network	TM / TMS	Transaction Monitoring / Transaction Monitoring System

Term	Meaning	Term	Meaning
HIDTA / HIFCA	High Intensity Drug Trafficking / Financial Crime Area	VASP	Virtual Asset Service Provider
IFM-X	Actimize Integrated Fraud Management (next-gen)	WLF	Watch List Filtering (Actimize sanctions screening engine)
KRI	Key Risk Indicator		

Table of Contents

1. Scope, Objectives and Regulatory Basis

1.1 Purpose of this Framework

This framework documents, end to end, how suspicious activity alerts are generated, scored, triaged, investigated and converted into Suspicious Activity Reports at Meridian National Bank (MNB), a hypothetical US Tier 1 institution running the **NICE Actimize** financial crime suite. It is written to the standard an examiner from the OCC or a FinCEN reviewer would expect to see on request: a document that explains not only *what* the bank monitors, but *how the engine physically arrives at an alert* — the feed, the data model, the profile, the binding, the match, the score, the queue, and the decision.

Three audiences are served simultaneously. **Investigators** need to know what an alert means and what evidence closes it. **Technologists** need the data contracts, the batch dependencies and the detection semantics. **Second and third line reviewers** need to see that every parameter in the system is owned, evidenced, approved and periodically re-validated. A transaction monitoring programme fails examination far more often for the third reason than for the first two.

> **The examiner's question this document answers**
> "Show me how an alert is produced. Walk me from the wire that hit the Fedwire gateway at 14:03 on Tuesday to the analyst who dispositioned it on Friday — every table it passed through, every threshold it was tested against, every human who touched it, and who approved that threshold."

1.2 Scope of Coverage

The framework covers all products, channels and legal entities within MNB's US BSA perimeter:

Dimension	In Scope	Out of Scope (covered elsewhere)
Lines of business	Retail Banking · Business Banking · Commercial & Corporate Banking · Private Bank / Wealth · Correspondent Banking · Trade Finance · Treasury Services · Digital Bank	Broker-dealer affiliate (FINRA regime, separate SAR-SF programme); Insurance affiliate
Products	DDA / savings / CD · loans & lines · commercial cards · prepaid · trade instruments (LC, guarantees, documentary collections) · nostro/vostro accounts · custody	Proprietary trading book (market abuse surveillance under a separate Actimize markets module)
Channels	Branch/teller · ATM · online & mobile · wire (Fedwire, CHIPS, SWIFT) · ACH · RTP/FedNow · Zelle · cheque/lockbox · card rails · third-party payment processors	Internal GL-only movements with no customer nexus (monitored by Finance controls)
Detection domains	AML suspicious activity monitoring (SAM) · CDD anticipated-activity variance · sanctions cross-referral from WLF · fraud cross-referral from IFM-X	Real-time sanctions interdiction mechanics (documented in the Sanctions Screening Standard; the referral interface only is described in Chapter 11)
Jurisdictions	All US-booked business plus US-dollar clearing activity for foreign branches	Locally-booked non-US activity monitored on local instances under local addenda

1.3 Regulatory Foundation

Every design decision in this document traces to an obligation. The table below is the control mapping the FIU maintains and presents at examination; the right-hand column names the chapter where the obligation is operationalised.

Authority	Obligation	Operationalised in
31 U.S.C. 5318(h); 31 CFR 1020.210	Maintain an AML/CFT programme with internal controls, a designated BSA Officer, independent testing, training, and ongoing risk-based customer due diligence.	Ch. 12, Ch. 13

Authority	Obligation	Operationalised in
31 CFR 1020.320	File a SAR for known or suspected violations aggregating $5,000 or more where the bank can identify a suspect ($25,000 or more where it cannot), within 30 calendar days of initial detection.	Ch. 9.7, Fig. 9
31 CFR 1010.311 / 1010.313	File a CTR for cash transactions exceeding $10,000 in a business day by or on behalf of the same person.	Ch. 6.2 (CTR scenarios), Annex A
31 CFR 1010.230 (CDD Rule)	Identify and verify beneficial owners of legal entity customers; understand the nature and purpose of the relationship to develop a customer risk profile and conduct ongoing monitoring.	Ch. 5.3, Ch. 6.6
USA PATRIOT Act §312 / §313 / §319	Enhanced due diligence for correspondent and private banking accounts; prohibition on shell-bank correspondents.	Ch. 6.5 (correspondent scenarios)
USA PATRIOT Act §314(a) / §314(b)	Respond to FinCEN information requests; permit voluntary information sharing between institutions under safe harbour.	Ch. 9.8
FFIEC BSA/AML Examination Manual	Monitoring systems must be commensurate with the bank's risk profile; the bank must be able to explain filtering criteria, thresholds and their reasonableness, and evidence periodic review.	Ch. 6, Ch. 12
SR 11-7 / OCC Bulletin 2011-12	Detection scenarios, segmentation logic and any AI alert-scoring model are *models*; they require development documentation, independent validation, ongoing monitoring and an inventory.	Ch. 12
12 CFR Part 30, Appendix D (OCC Heightened Standards)	Three lines of defence, front-line accountability for risk, independent risk management, credible challenge from Internal Audit.	Ch. 13, Fig. 11
31 CFR Part 501 (OFAC Reporting)	Report blocked and rejected transactions within 10 business days; annual blocked property report.	Ch. 11
FinCEN Advisories (e.g. human trafficking, elder exploitation, ransomware, fentanyl, real estate)	Incorporate published red-flag typologies into the monitoring programme and reference the advisory key term in the SAR narrative.	Ch. 6.7, Annex D
Anti-Money Laundering Act of 2020	Programmes must be effective, risk-based and reasonably designed; encourages innovation, feedback loops and prioritisation of national AML/CFT priorities.	Ch. 8 (AI prioritisation), Ch. 12

1.4 Design Principles

- **Risk-based, not volume-based.** The objective of tuning is not fewer alerts; it is a higher proportion of alerts that describe genuinely unusual behaviour. Volume reduction that suppresses true positives is a programme failure even if it improves every operational metric.
- **Every parameter has an owner and a rationale.** No threshold exists in production because it shipped as a vendor default. Each is evidenced, approved and dated.
- **Detect on behaviour, not on transactions.** The unit of interest is a focal entity exhibiting a pattern over a lookback window — not an individual payment.
- **Nothing is silently discarded.** Suppressed and auto-closed alerts are generated, persisted, actioned with a reason code and retained. An alert the system never wrote down is an alert the bank cannot prove it considered.
- **The data layer is a first-class control.** Detection quality is bounded above by feed completeness. Data quality is monitored as a compliance control, not as an IT service metric.
- **Explainability over sophistication.** Any machine-learning component must be able to state, in plain language, why it ranked an alert where it did — because an investigator will have to defend that ranking to an examiner.

1.5 Document Users

Role	How this document is used
L1 Triage Analyst	Understands what each scenario is designed to catch (Ch. 6) and the evidence standard for closure (Ch. 9.4).
L2 / L3 Investigator	Case construction, EDD standards, SAR decisioning and narrative quality (Ch. 9).
FIU Team Lead / Queue Manager	Routing, capacity, SLA management, escalation matrix (Ch. 9.2, Ch. 13).
BSA Officer	Owns the framework; approves thresholds, suppressions and SAR filings.
AML Technology / Actimize Admin	Architecture, feed contracts, ADM schema, batch operations, Policy Manager change control (Ch. 3, 4, 7).
Model Risk Management (2LoD)	Validation scope: segmentation, scenario logic, thresholds, AI scoring (Ch. 12).
Compliance Testing / QA	Sampling methodology and decision-quality standards (Ch. 12.6).
Internal Audit (3LoD)	End-to-end lifecycle audit programme (Ch. 13).
Regulators / External Audit	Primary written evidence of the design and governance of the monitoring system.

2. Glossary — Actimize Detection Terminology

Precision here matters. Investigators, technologists and examiners routinely use the words *alert*, *match* and *case* interchangeably, and the resulting confusion shows up in MIS, in SAR timeliness calculations and in audit findings. The definitions below are the ones used throughout this document and configured in the platform.

Term	Definition
Focus / Focal Type	The type of business entity an alert is centred on. MNB monitors on four focus types: **CU** (Customer — aggregates all accounts under one Party ID), **AC** (Account), **CB** (Correspondent Bank), and **HH** (Household, for retail relationship views). Choosing the wrong focus is the single most common scenario-design error: structuring belongs on CU (a structurer splits across accounts), while dormancy reactivation belongs on AC.
Focal Entity	The specific instance of the focus type that triggered the behaviour — Party ID 8842391, Account 0019-448220. This is the subject of the alert and, if escalated, the subject of the SAR.
Scenario / Model	A codified behaviour of interest. In Actimize SAM a scenario is implemented as a *model* containing one or more *patterns*; the model is parameterised by a threshold set and executed inside a *plan*.
Pattern	The deterministic detection logic inside a scenario — the aggregation, the comparison, the sequencing. One scenario can carry several patterns (e.g. Structuring carries a 'multiple cash below threshold' pattern and a 'cash plus monetary instrument' pattern).
Binding / Binding Set	The grouping key over which the pattern aggregates — typically (focal entity, business date window), sometimes extended by counterparty, corridor or channel. One binding set that satisfies the pattern produces one match.
Match	A single detected occurrence of a pattern, carrying the *matched records* (the specific transactions that evidence it) and a match score. Matches are the atoms; alerts are the molecules.
Alert	A unit of work presented to a human. One or more matches for the same scenario and the same focal entity on the same business date are consolidated into a single *multi-match alert*. An alert is always about behaviour, never about a single transaction.
Case	An investigation object in ActOne. A case may link many alerts across many scenarios and several related parties, and it is the container in which EDD, RFI responses, flow-of-funds analysis and the SAR decision live.
Threshold	A tunable parameter inside a scenario (amount, count, percentage, day-count). A **threshold set** is the complete vector of parameters bound to one *segment* at run time.
Segment	A population of focal entities that share monitoring treatment — defined by customer risk rating × line of business × entity type × product family. Segments are why a $9,500 cash deposit alerts for a retail salaried customer and does not alert for a supermarket chain.
Frequency	How often the detection job for a scenario runs (daily, weekly, month-start, month-end). Frequency also governs de-duplication: within one frequency period the same focal entity will not be re-alerted by the same scenario for the same behaviour.
Lookback Period	The span of history each run examines. Lookback is routinely longer than frequency — a scenario may run daily but evaluate a rolling 30 days. Lookback ≥ frequency is the normal configuration; lookback < frequency creates blind windows and is prohibited.
Profile / Aggregate	Pre-computed rolling statistics per focal entity (30/60/90/180/365-day sums, counts, means, standard deviations, peer-relative z-scores). Detection reads profiles, not raw transactions, which is what makes the nightly run tractable at MNB's volumes.
Peer Group	A statistically or business-defined cohort against which a focal entity's behaviour is compared. A cash-intensive restaurant is only anomalous relative to other cash-intensive restaurants.
Match Score / Alert Score	A 0–100 measure of relative risk used to prioritise work. The score is a workload management device; it is not a probability of criminality and must never be cited in a SAR narrative.
Suppression	An approved, time-bound rule that continues to *generate* alerts for a (scenario × focal entity) pair but auto-closes them with the action 'Suppressed' before an analyst sees them. The alert record persists; the human effort does not.
Trusted Pair	An approved, time-bound exclusion for a *pair* of parties (A ↔ B) whose relationship has been investigated and understood — e.g. a customer and its own payroll processor. Scoped by direction, channel and an amount ceiling.

Term	Definition
Auto-Close Rule	A rules-based post-processing step that closes alerts meeting objectively low-risk criteria (e.g. score below floor AND focal entity Low CRR AND no prior alert in 12 months AND no watch-list nexus). Auto-close rates are reported to the Steering Committee and independently sampled by QA.
Prior	The count of previous matches for the focal entity: *All Prior* (across all scenarios) and *Same-Scenario Prior*. Both are inputs to the match score; a repeat is more interesting than a first.
Exempt Entity	An entity excluded from AML detection entirely — used almost exclusively for the bank's own internal suspense and settlement accounts. Exempting a *customer* requires BSA Officer approval and is heavily disfavoured; CTR exemption (31 CFR 1020.315) is a distinct concept and does not exempt a customer from SAR monitoring.
Initial Detection	The date on which the bank knows, or has reason to suspect, facts constituting a basis for filing. This is **not** automatically the alert-generation date, but MNB's conservative policy starts the 30-day SAR clock at alert generation unless the FIU documents a later, justified detection date.

3. NICE Actimize Platform Architecture

MNB runs the Actimize financial crime suite as an integrated stack rather than as four unrelated point solutions. The value of the integration is contextual: a sanctions near-match in WLF, a declined card authorisation in IFM-X and a cash structuring pattern in SAM converge on the same resolved Party ID and land in the same ActOne case.

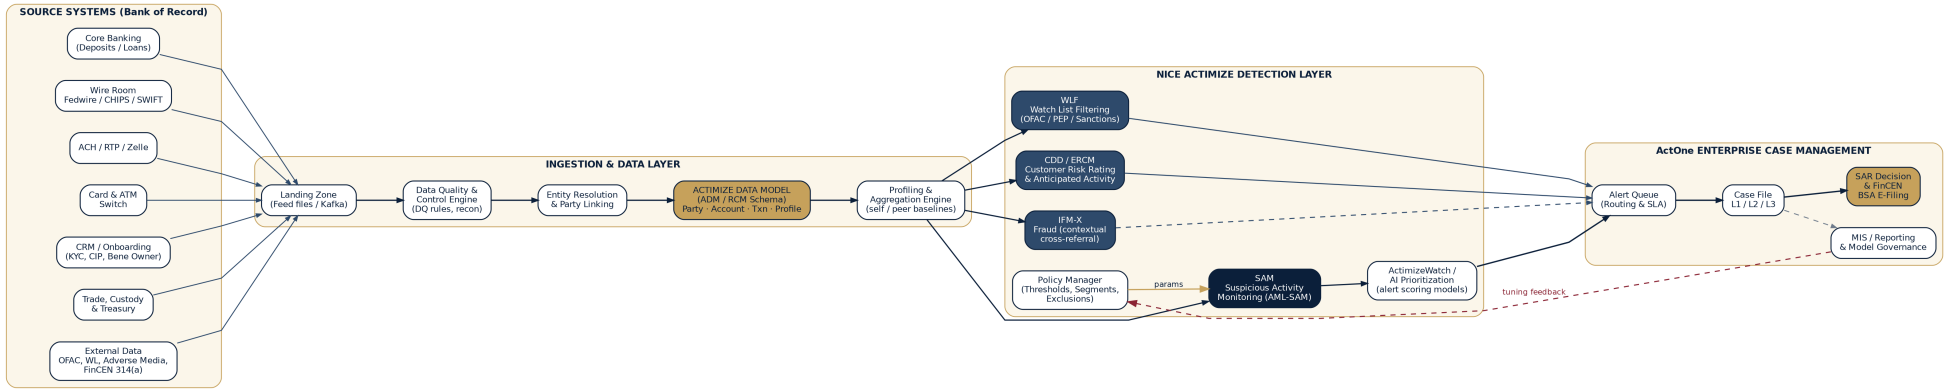

Figure 1 — End-to-end platform architecture: source systems, ingestion and data layer, Actimize detection layer, and ActOne enterprise case management, with the tuning feedback loop.

3.1 Component Inventory

Component	Function	Deployment at MNB
AIS — Actimize Integrated Server	The execution runtime. Hosts detection models, the scheduler, the analytics engine and the services layer. Orchestrates plans, jobs and model execution against the ADM.	Clustered, 6 application nodes; Oracle Exadata backing store; RHEL
SAM — Suspicious Activity Monitoring	The AML behavioural detection engine. Executes scenario models over profiles and transactions and emits matches and alerts. This is the subject of Chapters 6 and 7.	24 scenarios in production (Annex A); nightly and intraday plans
Policy Manager	The business-facing control plane for thresholds, segments, exclusion lists, suppression rules and trusted pairs. Every change is versioned, dual-controlled and audit-logged — this is the artefact examiners ask for when they ask 'who changed this threshold and when?'	Restricted to the AML Model Governance team; four-eyes enforced
WLF — Watch List Filtering	Real-time and batch sanctions/PEP/adverse-media screening of parties and payment messages against OFAC SDN and Consolidated lists, UN, EU, UK OFSI and internal lists.	Real-time inline on the payment rails; nightly full-book party rescreen
CDD / ERCM	Customer risk rating engine and anticipated-activity capture. Produces the CRR that segments every SAM scenario, and the expected-activity envelope that scenario SAM-BD-02 tests against.	Event-driven CRR recalculation plus periodic review cycles
IFM-X — Integrated Fraud Management	Real-time fraud detection. Interfaces with SAM through a bidirectional referral: confirmed fraud raises an AML alert (fraud proceeds are laundered proceeds), and AML findings feed fraud contextual features.	Real-time on card, digital and payment channels
ActOne — Enterprise Case Management	The investigator workspace. Alert queues, routing, SLA clocks, evidence management, flow-of-funds visualisation, narrative authoring, SAR generation and BSA E-Filing integration.	~380 named FIU users; queue model in Chapter 9.2
ActimizeWatch / AI Prioritisation	Supervised models trained on historical alert dispositions that rank alerts by likelihood of escalation. Advisory only at MNB — the model re-orders the queue, it never closes an alert.	Champion/challenger; quarterly retraining; SR 11-7 governed
Visual Modeler / Model Builder	The environment in which custom scenarios and derived attributes are authored, unit-tested against historical data, and promoted through DEV → SIT → UAT → PROD.	Custom scenarios: SAM-TB-01, SAM-VA-01, SAM-HT-01 (Ch. 6)

3.2 Environment Topology and Change Control

Environment	Data	Purpose	Promotion gate
DEV	Synthetic + masked	Scenario authoring, unit tests	Peer code review
SIT	Masked 3-month copy	Integration, feed contract testing	Technical test evidence
UAT / Model Lab	Full production copy (12–24 months), access-restricted	**Threshold tuning, ATL/BTL sampling, impact simulation, scenario back-testing.** This is the single most important non-production environment in the AML estate — no threshold reaches production without a UAT simulation of its volume impact.	MRM validation + AML Model Committee
PROD	Live	Detection and case management	Dual-control release; CAB approval
DR	Replicated (RPO 15 min, RTO 4 h)	Business continuity	Semi-annual failover test with a live detection run executed in DR

Detection-affecting change is compliance change. Any modification to a scenario's logic, a threshold, a segment definition, an exclusion list, an auto-close rule or a scoring weight is treated as a change to the AML programme, not as an IT release. It requires: a documented rationale, a UAT impact simulation, MRM review where the change is material, AML Model Committee approval, BSA Officer sign-off, and a 90-day post-implementation monitoring report. Emergency changes (e.g. a new OFAC designation requiring immediate

coverage) may be applied under BSA Officer authority with retrospective committee ratification within 5 business days.

4. Data Architecture, Feeds and Schemas

The governing constraint
A transaction monitoring system cannot detect what it was never sent, and cannot correctly aggregate what it cannot resolve to a party. In enforcement actions against large US banks, the root cause is far more often *incomplete or mis-mapped data feeds* than *badly designed scenarios*. Data completeness is therefore treated here as a Tier 1 compliance control with board-level reporting, not as an IT hygiene metric.

4.1 Source System Inventory and Feed Contracts

Feed ID	Source	Content	Mode / Frequency	Format	Arrival SLA (ET)
F-CUST-01	CRM / Onboarding	Party master, CIP, beneficial ownership, NAICS, PEP flags	Batch delta, daily	Fixed-width DIS	00:20
F-ACCT-01	Core Banking	Account master, product, status, open/close, anticipated activity	Batch delta, daily	Fixed-width DIS	00:20
F-TXN-01	Core Banking	All ledger postings (deposits, withdrawals, transfers, fees)	Batch full-day, daily	Pipe-delimited	00:35
F-CASH-01	Teller / Vault / ATM	Cash-in / cash-out, conductor, branch, teller ID	Batch full-day, daily	Pipe-delimited	00:35
F-WIRE-01	Wire Room (Fedwire/CHIPS/SWIFT)	MT103/MT202, pacs.008/009, full originator & beneficiary chain	Near-real-time stream + daily reconciliation file	ISO 20022 XML / Kafka	Continuous; recon 00:45
F-ACH-01	ACH Operator	NACHA entries, SEC codes, company IDs, addenda	Batch, 4× daily	NACHA	01:00
F-RTP-01	RTP / FedNow / Zelle	Instant payment credit transfers	Near-real-time stream	ISO 20022 / JSON	Continuous
F-MI-01	Item Processing	Cheques, cashier's cheques, money orders, traveller's cheques	Batch, daily	Pipe-delimited	01:30
F-CARD-01	Card Switch	ATM withdrawals, POS, cash advance	Batch, daily	ISO 8583 derived	01:00
F-TRD-01	Trade Finance	LCs, guarantees, documentary collections, goods description, HS codes, ports	Batch, daily	XML	02:00
F-WL-01	Sanctions Data Vendor	OFAC SDN/Consolidated, UN, EU, OFSI, PEP, adverse media	Push on publication (intra-day)	XML	≤ 24 h from OFAC publication
F-314A-01	FinCEN Secure Portal	314(a) subject lists	Bi-weekly	CSV	Within 14 days of receipt
F-REF-01	Reference Data Hub	Country risk, FATF lists, HIDTA/HIFCA ZIPs, NAICS risk, correspondent risk	Weekly / on change	CSV	Sun 22:00

Each feed is governed by a signed **Data Interface Specification (DIS)**: field list, datatypes, mandatory/optional designation, code sets, null-handling, and the control totals the file must carry. The DIS is a two-way contract — the source system owner attests annually that the feed remains complete and correctly mapped, and that attestation is evidence the FIU produces at examination.

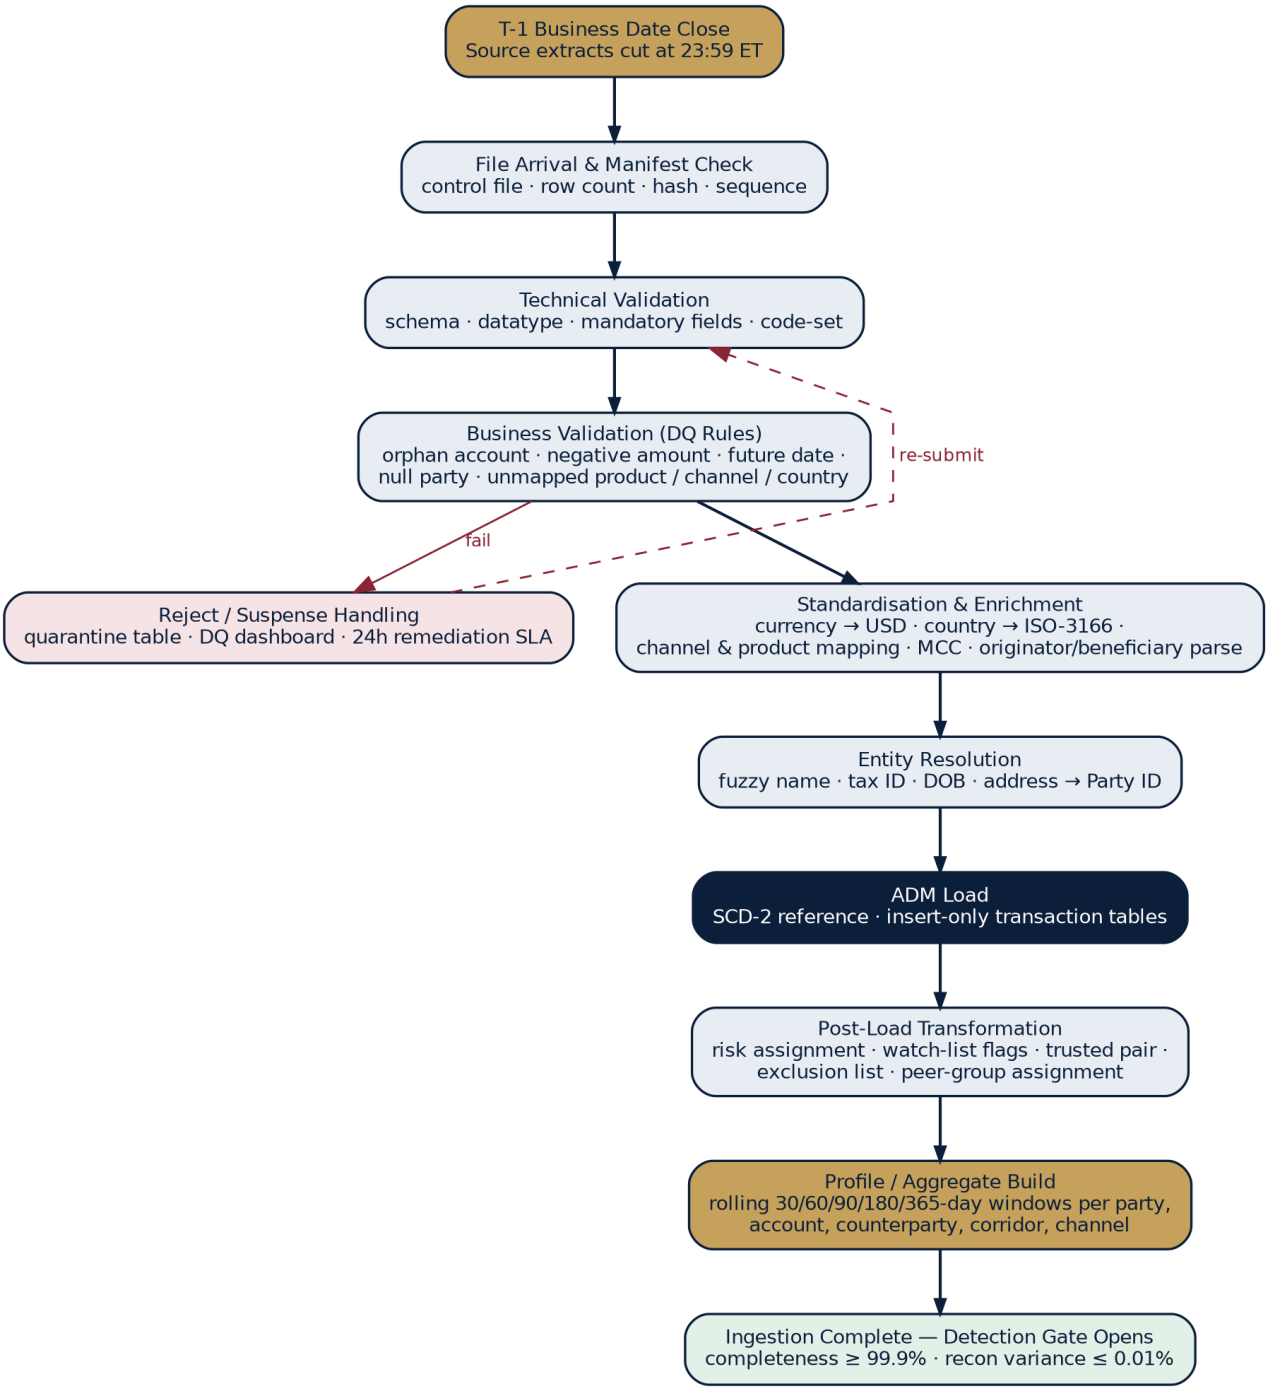

Figure 2 — Data ingestion pipeline: from source cut-off through validation, enrichment, entity resolution, ADM load and profile build to the detection gate.

4.2 The Actimize Data Model — Core Schema

Detection reads a normalised model, not source-system tables. The ERD below shows the core structures; the detail that matters for scenario authoring is which fields are *mandatory*, because a scenario that keys on an optional field silently under-detects wherever that field is null.

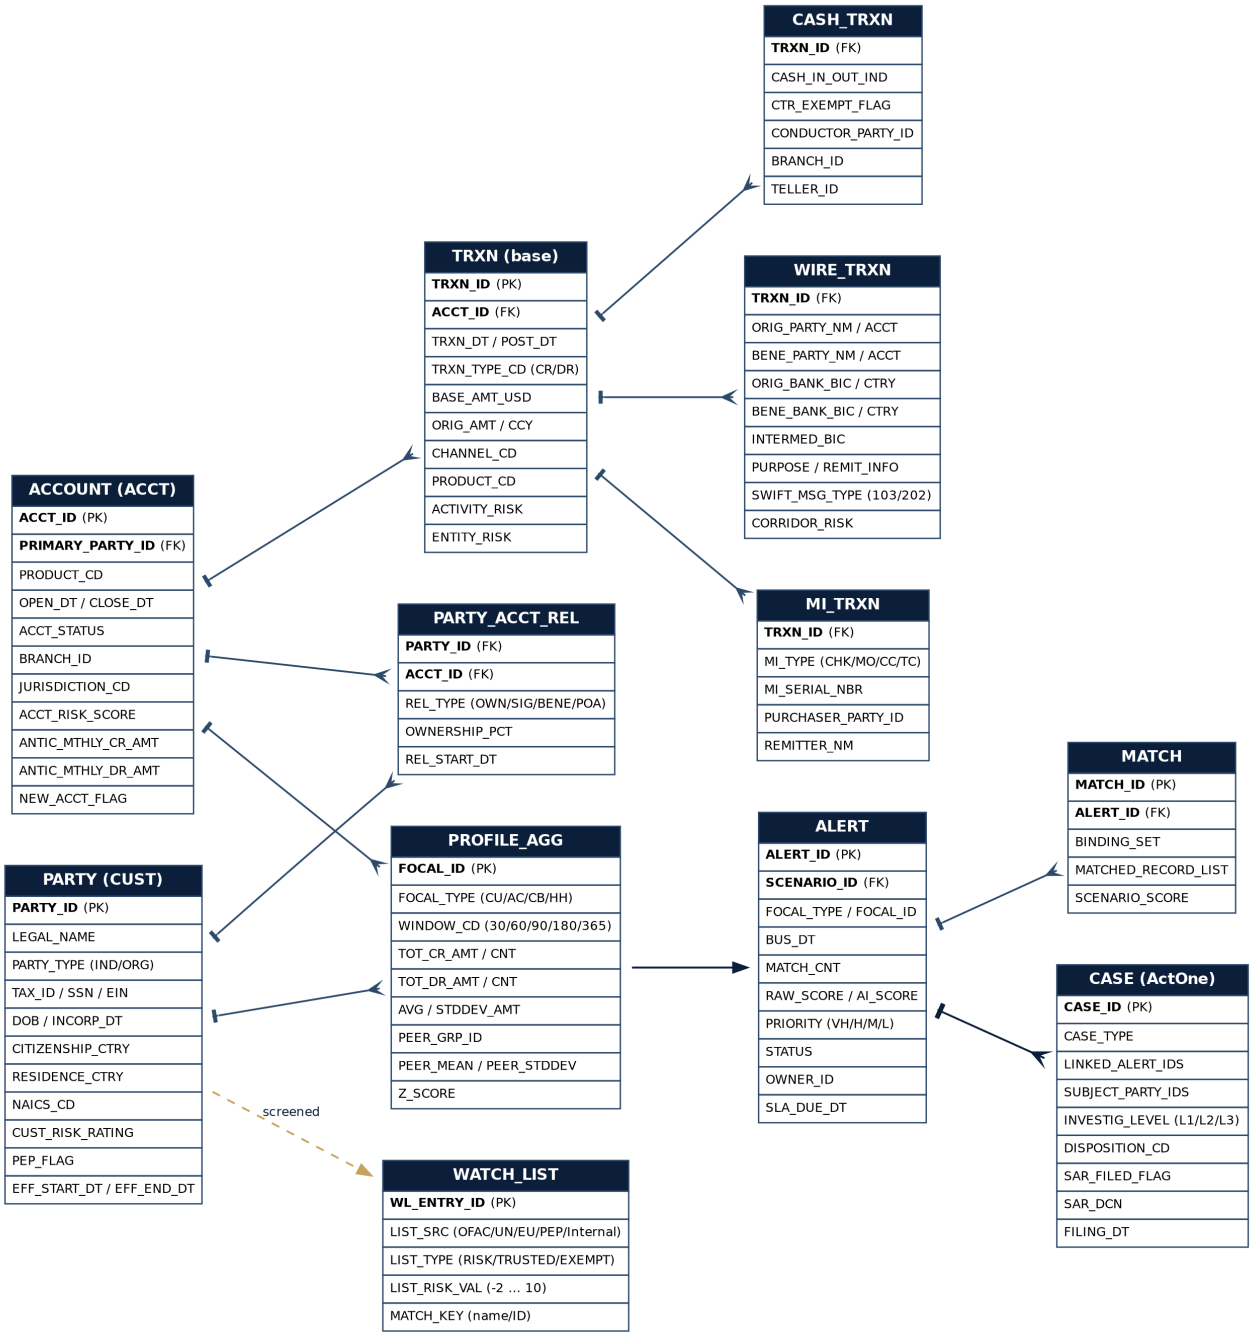

Figure 3 — Actimize Data Model (ADM): core party, account, transaction, profile and alert/case entities with cardinality.

4.2.1 Selected field-level specification

Table.Field	Type	Mand.	Description / detection significance
PARTY.PARTY_ID	VARCHAR(20)	Y	Golden party key produced by entity resolution. The aggregation key for every CU-focus scenario.
PARTY.TAX_ID	VARCHAR(20)	C	SSN/EIN. Conditionally mandatory (US persons). Drives deterministic ER and 'unrelated party' tests.
PARTY.NAICS_CD	CHAR(6)	Y	Industry code. Drives peer grouping and the cash-intensive-business flag. A null NAICS on a business customer is a DQ failure, not a default.
PARTY.CUST_RISK_RATING	CHAR(2)	Y	LO/ME/HI/VH from the CDD engine. **Binds the threshold set at run time.**
PARTY.PEP_FLAG	CHAR(1)	Y	Y/N, including close associates and family members.

Table.Field	Type	Mand.	Description / detection significance
ACCOUNT.OPEN_DT	DATE	Y	Drives the NEW_ACCT_FLAG (MNB: open ≤ 180 days) used to tighten thresholds on unseasoned relationships.
ACCOUNT.ANTIC_MTHLY_CR_AMT	NUMBER(18, 2)	Y	Anticipated monthly credits captured at onboarding. The denominator of the CDD variance scenario. If this is a copy-paste default across a portfolio, the scenario is worthless — QA samples it.
TRXN.BASE_AMT_USD	NUMBER(18, 2)	Y	USD-normalised amount at the transaction-date FX rate. All thresholds are expressed in USD.
TRXN.CHANNEL_CD	VARCHAR(10)	Y	Branch/ATM/ONLINE/MOBILE/WIRE/ACH/RTP/CARD/TPPP. Drives channel risk and several patterns.
TRXN.TRXN_DT / POST_DT	DATE	Y	Both retained. Detection uses TRXN_DT; a systematic gap between them is a red flag for the DQ engine.
CASH_TRXN.CONDUCTOR_PARTY_ID	VARCHAR(20)	C	The person physically presenting the cash. **Critical for structuring detection** — smurfing is conducted by third parties into a beneficiary's account, and if this field is not captured the pattern is invisible.
WIRE_TRXN.ORIG_PARTY_NM	VARCHAR(140)	Y	Full originator name. Truncation to 35 characters by a legacy gateway is a classic, and expensive, mapping defect.
WIRE_TRXN.BENE_BANK_CTRY	CHAR(2)	Y	ISO-3166. Joins to country risk; drives all geography scenarios.
WIRE_TRXN.INTERMED_BIC	VARCHAR(11)	N	Intermediary institution — the field where stripping and nested correspondent risk becomes visible.
WIRE_TRXN.PURPOSE / REMIT_INFO	VARCHAR(255)	N	Free-text remittance information. Feeds keyword scanning (Ch. 6.7) despite being optional — coverage is measured and reported.
PROFILE_AGG.Z_SCORE	NUMBER(8,4)	Y	(observed – peer mean) / peer stddev. The core of every behavioural-deviation scenario.

4.3 Data Quality Control Framework

DQ rules run inside the ingestion pipeline and gate the detection run. Rules are tiered: a **Critical** breach halts the batch and pages the on-call; a **High** breach permits the run but raises a compliance incident and is reported to the BSA Officer the same day; a **Moderate** breach is tracked to remediation on a 30-day clock.

ID	Control	Test	Tier	Tolerance
DQ-01	File completeness	All expected feeds present and manifest-matched (row count + hash)	Critical	100%
DQ-02	Volumetric variance	Daily record count within ±3 standard deviations of the trailing 30-day mean for that weekday	Critical	±3σ
DQ-03	Value reconciliation	Sum of TRXN.BASE_AMT_USD reconciles to the general ledger control total	Critical	≤ 0.01%
DQ-04	Orphan transactions	Every TRXN.ACCT_ID resolves to an ACCOUNT row	Critical	0 records
DQ-05	Unlinked accounts	Every ACCOUNT has ≥ 1 PARTY_ACCT_REL row	Critical	0 records
DQ-06	Mandatory field population	All 'Y' fields in the DIS non-null	High	≥ 99.9%
DQ-07	Wire chain completeness	Originator and beneficiary name + account + country populated on cross-border wires	High	≥ 99.5%
DQ-08	Conductor capture	CONDUCTOR_PARTY_ID populated on branch cash transactions ≥ $3,000	High	≥ 98%

ID	Control	Test	Tier	Tolerance
DQ-09	Code-set validity	Country, currency, channel, product, NAICS values exist in the reference domain	High	≥ 99.9%
DQ-10	Anticipated-activity plausibility	ANTIC_MTHLY_CR_AMT non-null, non-zero, and not identical across > 5% of a portfolio	High	< 5% clustering
DQ-11	FX normalisation	Every non-USD transaction has a rate applied for its TRXN_DT	High	100%
DQ-12	Duplicate detection	No duplicate TRXN_ID within a business date	Critical	0 records
DQ-13	Late-arriving data	Records with TRXN_DT more than 2 days prior to the batch date	Moderate	< 0.5%; triggers back-fill run
DQ-14	Watch-list currency	OFAC list version loaded ≤ 24 h after publication	Critical	≤ 24 h

4.3.1 Back-fill and missed-batch handling

If a feed fails and is remediated after the detection gate has closed, the FIU does not simply resume the next night — that would leave a permanent blind window. The controlled response is a **back-fill detection run**: the missing business dates are loaded, the profiles are rebuilt for the affected focal entities, and the scenarios are re-executed for those dates only, with the resulting alerts flagged BACKFILL=Y and routed to a dedicated queue so that SAR timeliness is measured from the corrected detection date. Every back-fill is logged, reported to the BSA Officer, and disclosed in the quarterly board MIS. A missed window that is never backfilled is a reportable programme deficiency.

4.4 Retention and Lineage

- **Five-year retention** (31 CFR 1010.430) for all alerts, cases, supporting transactions, analyst notes, SAR filings and E-Filing acknowledgements — including alerts that were auto-closed or suppressed.
- **Immutable audit trail** on every alert and case action: actor, timestamp, prior value, new value, justification. ActOne audit records are write-once.
- **Full lineage** from a value shown in an alert back to the source-system record and the feed file that carried it. When an investigator says 'the beneficiary bank was in Country X', the bank must be able to produce the MT103 field it came from.
- **Threshold history** retained indefinitely: the bank must be able to state what threshold was in force on any past date, and reproduce why a historical transaction did or did not alert.

5. Entity Resolution, Risk Rating and Segmentation

Three upstream decisions determine whether a scenario can work at all: *who is this?* (entity resolution), *how risky are they?* (customer risk rating), and *who should they be compared to?* (segmentation). Get these wrong and no amount of scenario tuning recovers the programme.

5.1 Entity Resolution

MNB's ER engine collapses source-system party records into a single golden Party ID. Without it, a structurer who opened three accounts under 'Robert J. Chen', 'Bob Chen' and 'R. Chen' presents as three unremarkable customers instead of one $28,000-a-week cash pattern.

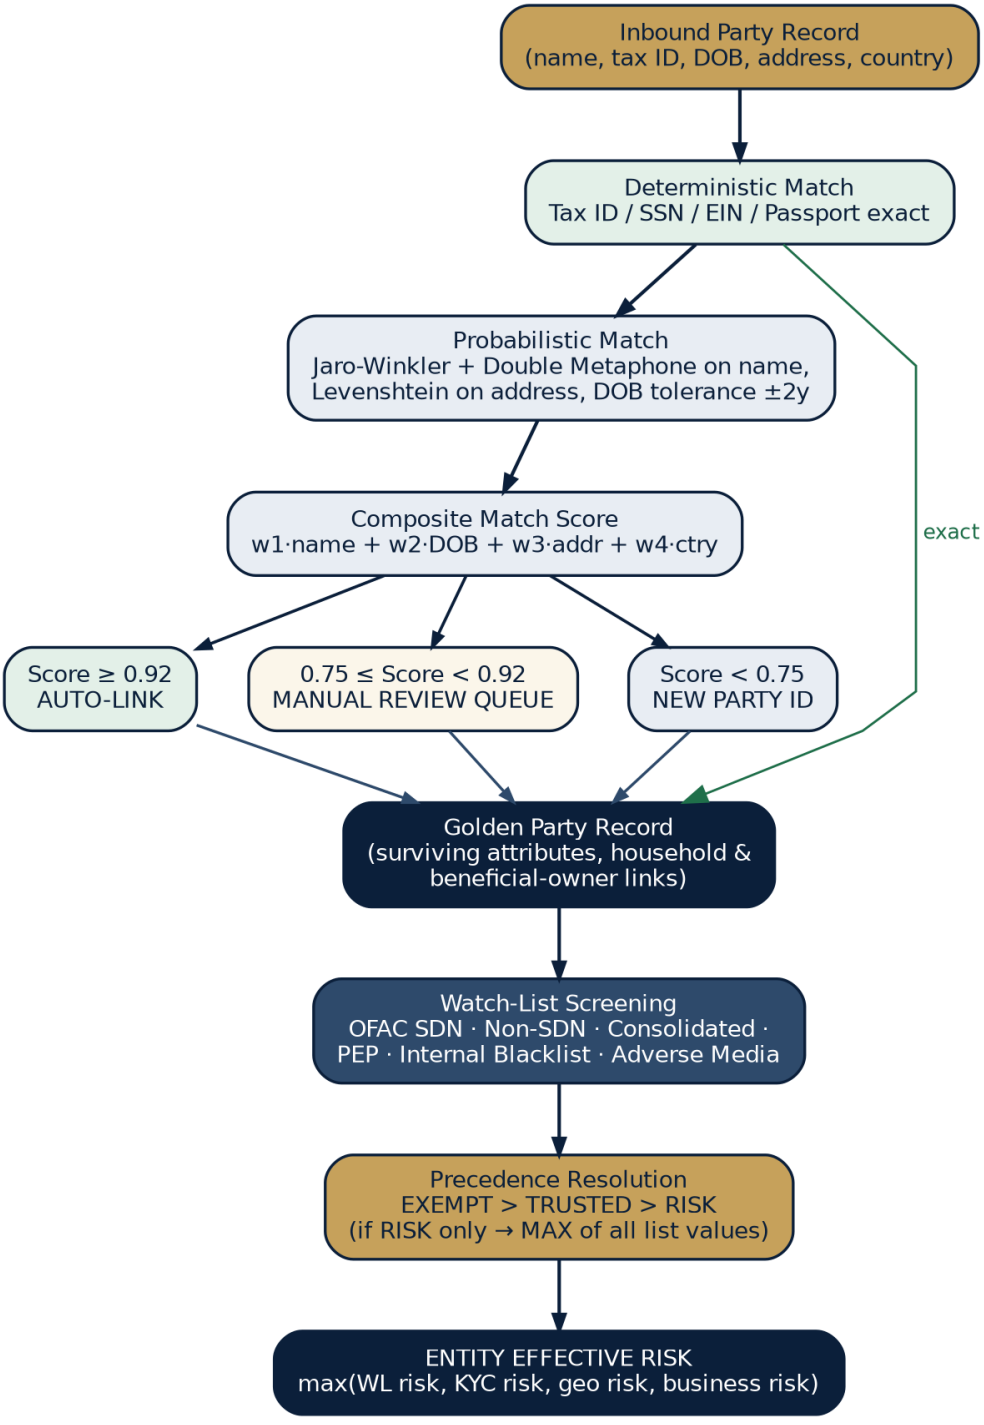

Figure 4 — Entity resolution and effective-risk determination, including watch-list precedence (exempt > trusted > risk).

5.1.1 Matching algorithm

Stage	Method	Weight	Notes
Deterministic	Exact match on SSN / EIN / passport + DOB	—	Auto-link on hit; bypasses scoring
Name similarity	Jaro-Winkler on standardised name + Double Metaphone phonetic key	0.40	Handles transliteration, nicknames, transposition
Date of birth	Exact, or within ±2 years, or day/month transposition	0.20	Transposition is a deliberate evasion signal, not just an error
Address	Levenshtein on normalised address + exact ZIP	0.20	Postal standardisation (USPS) applied first
Identifiers (partial)	Last-4 SSN, phone, email, device	0.15	Weak alone; strong in combination
Country / nationality	Exact ISO match	0.05	Tie-breaker only

Entity resolution scoring — pseudocode

```
COMPOSITE_SCORE = 0.40*jaro_winkler(name_a, name_b)
                + 0.20*dob_similarity(dob_a, dob_b)
                + 0.20*addr_similarity(addr_a, addr_b)
                + 0.15*partial_id_similarity(ids_a, ids_b)
                + 0.05*country_match(ctry_a, ctry_b)

IF   COMPOSITE_SCORE >= 0.92  -> AUTO-LINK to existing PARTY_ID
ELIF COMPOSITE_SCORE >= 0.75  -> ER MANUAL REVIEW QUEUE (data steward, 2-day SLA)
ELSE                          -> CREATE NEW PARTY_ID

# Governance: the 0.92 / 0.75 cut-points are MODEL PARAMETERS under SR 11-7.
# They are validated annually against a labelled sample of 1,000 record pairs,
# reporting precision, recall and F1. Target: recall >= 0.97 at precision >= 0.95.
```

5.2 Watch-List Screening and Effective Risk

Every resolved party is screened against risk, trusted and exempt lists. Where a party appears on several lists, precedence resolves the conflict: **Exempt > Trusted > Risk**. Where a party appears only on risk lists, effective risk is the **maximum** of those list values — never the average, because averaging a single sanctions hit against nine benign entries would extinguish it. List-membership records based on a matched *identifier* take precedence over records based on a matched *name*.

Risk value	Meaning	Effect on detection
10	Sanctioned / OFAC SDN nexus	Immediate WLF interdiction; SAM thresholds at their most stringent; case auto-created
8–9	PEP, adverse media (financial crime), law-enforcement subject	Very High segment thresholds; EDD required
5–7	High-risk jurisdiction, high-risk industry (MSB, casino, crypto, precious metals)	High segment thresholds
1–4	Elevated but managed risk	Medium segment thresholds
0	Insufficient information to assess	**Treated as Medium, never as Low.** Absence of evidence is not evidence of absence — and a systematic population of zeros indicates a data defect
–1	Low risk	Low segment thresholds
–2	Trusted / exempt	Suppressed from AML scenarios (bank's own internal accounts only, save for BSA-Officer-approved exceptions)

5.3 Customer Risk Rating Model

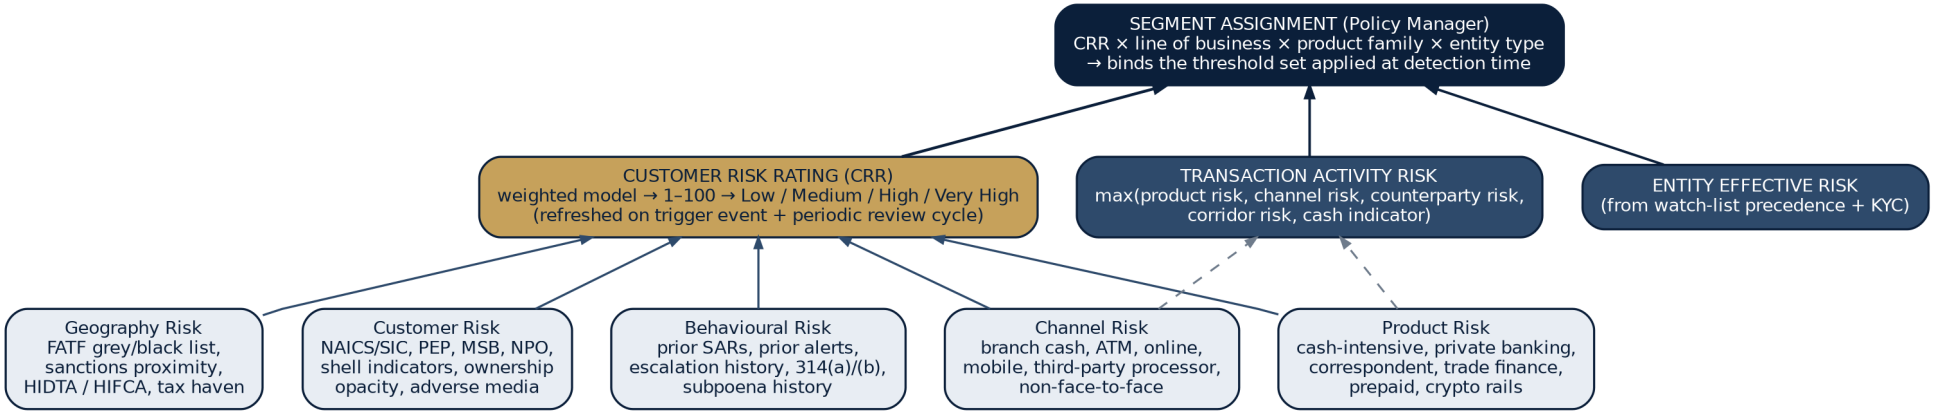

Figure 5 — Risk aggregation: customer, geography, product, channel and behavioural risk into the CRR, and the CRR into the segment that binds the threshold set.

Factor	Weight	Illustrative scoring
Customer type & industry (NAICS)	25%	Salaried individual 5 · SME retail 25 · Import/export 60 · MSB 85 · Crypto/VASP 90 · Cash-intensive (car wash, restaurant, convenience) 70 · NPO/charity 65 · Shell indicators 95
Geography (residence, citizenship, operating footprint)	20%	Domestic 5 · FATF-compliant 15 · FATF grey list 60 · FATF black list / comprehensively sanctioned 100 · HIDTA/HIFCA ZIP 50 · Offshore secrecy jurisdiction 70
Product & service	20%	Basic DDA 10 · Wealth/private banking 55 · Trade finance 70 · Correspondent (nested) 90 · Prepaid/stored value 65 · Wire-intensive 60
Channel & delivery	10%	Branch face-to-face 10 · Digital onboarding 45 · Introduced by third party 65 · Non-face-to-face foreign 75
Ownership & control transparency	15%	Natural person 5 · Simple LLC with disclosed UBO 25 · Multi-layer structure 70 · Nominee directors / bearer instruments 95
Behavioural & historical	10%	No history 0 · Prior alerts closed NFA 20 · Prior SAR 80 · 314(a) match 90 · Subpoena / LE enquiry 85

The weighted score maps to **Low (1–25) · Medium (26–50) · High (51–75) · Very High (76–100)**. Two overrides operate outside the arithmetic: any OFAC nexus forces Very High irrespective of score, and any BSA Officer manual override — up or down — is permitted but must be documented with rationale and is 100% sampled by QA. **Downward overrides are the highest-risk control point in the entire CDD programme** and are reported to the Steering Committee individually, by name.

CRR is recalculated on: onboarding; periodic review (Very High 12m / High 12m / Medium 24m / Low 36m); and on trigger events — SAR filing, watch-list hit, adverse media, beneficial-ownership change, material change in activity profile, or a 314(a) match.

5.4 Segmentation Methodology

Segmentation is the mechanism that makes a single scenario risk-sensitive. One structuring scenario, correctly segmented, behaves as thirty. Segments at MNB are constructed as **CRR × line of business × entity type × product family**, then validated statistically.

Step	Method	Evidence produced
1. Candidate segments	Business-driven partition (LoB × entity type × cash-intensity flag × CRR)	Segment definition register
2. Statistical validation	K-means / hierarchical clustering on behavioural features (monthly credit volume, txn count, cash ratio, cross-border ratio, counterparty diversity); silhouette coefficient and elbow analysis	Clustering report; target silhouette ≥ 0.45
3. Homogeneity test	Within-segment coefficient of variation on key detection variables	CV report; segments with CV > 1.5 are split
4. Population test	Minimum viable segment size	No production segment below 500 focal entities (below that, thresholds cannot be statistically supported and the segment is merged upward)
5. Stability test	Segment migration rate quarter-on-quarter	Migration > 15% triggers re-segmentation review
6. Approval	AML Model Committee, with MRM opinion	Approved segmentation model, versioned

Why segmentation is the most-cited weakness in TM validations

A bank that applies a single $10,000 cash threshold across its entire book will simultaneously drown in alerts from its supermarket customers and miss a $9,200 pattern on a salaried retail customer whose entire recorded income is $4,000 a month. Both failures are invisible in aggregate alert-volume MIS. Only segment-level productivity reporting exposes them — which is why Chapter 12 requires productivity to be reported *per scenario per segment*, never in total.

6. Detection Scenario Catalogue

MNB runs 24 production scenarios grouped into eight scenario classes. Each is documented below with its focus, frequency, lookback, detection logic and the typology it addresses. Thresholds shown are illustrative Medium-risk-segment values; the complete segment matrix is in Annex A. Detection logic is presented as pseudocode because that is the form in which it can be reviewed by a non-technologist and validated by MRM.

Coverage assessment
The scenario inventory is not a shopping list — it is a *response* to the Enterprise-Wide AML Risk Assessment. Annually, every inherent-risk typology identified in the EWRA is mapped to the scenario(s) that cover it. Any typology with no covering scenario is a documented coverage gap with a remediation owner and date. This mapping is the first artefact an examiner requests.

6.1 Scenario Class — Cash and Structuring

The most productive class in a US retail-heavy book, and the one with the tightest regulatory nexus (31 U.S.C. 5324 makes structuring a substantive federal offence in itself — the underlying funds need not be criminal proceeds for a SAR to be required).

SAM-CS-01 — Structuring — Avoidance of the CTR Reporting Threshold

Focus	Frequency	Lookback	Typology addressed
CU	Daily	7 days	Smurfing; deliberate avoidance of 31 CFR 1010.311 reporting

Parameter	Medium-segment value	Rationale
CTR_THRESHOLD	$10,000	Statutory
STRUCT_BAND_LOW	$7,000	Lower bound of the 'just below' band; set from the observed distribution of cash amounts, which shows a pronounced spike between $8,500 and $9,900
STRUCT_BAND_HIGH	$9,999	Upper bound
MIN_TXN_COUNT	3	Two could be coincidence; three within a week is a pattern
AGG_AMOUNT_MIN	$18,000	Aggregate floor — suppresses trivial multi-deposit noise
LOOKBACK_DAYS	7	Aligned to the typical smurfing cadence observed in prior SARs
INCLUDE_MI	TRUE	Cashier's cheques and money orders purchased with cash are in scope — an omission that has featured in multiple US enforcement actions

```
FOR EACH party p IN active_parties:

  txns = SELECT * FROM TRXN t JOIN CASH_TRXN c USING (TRXN_ID)
         WHERE t.PARTY_ID = p                       -- CU focus: ALL accounts of the party
           AND t.TRXN_DT BETWEEN (run_dt - LOOKBACK_DAYS) AND run_dt
           AND t.BASE_AMT_USD BETWEEN STRUCT_BAND_LOW AND STRUCT_BAND_HIGH
           AND c.CASH_IN_OUT_IND IN ('IN','OUT')
           AND (INCLUDE_MI OR t.TRXN_TYPE <> 'MI_PURCHASE')

  IF COUNT(txns) >= MIN_TXN_COUNT
     AND SUM(txns.BASE_AMT_USD) >= AGG_AMOUNT_MIN:

     -- suppress the trivially benign: same-day single-branch payroll deposits
     IF NOT (all_same_day(txns) AND all_same_branch(txns) AND p.NAICS IN CASH_INTENSIVE_SET):

        RAISE MATCH(
          focal      = p,
          binding    = (p, run_dt - LOOKBACK_DAYS .. run_dt),
          records    = txns,
          highlights = { txn_count, total_amt, distinct_branches, distinct_days,
                         distinct_conductors, max_single_amt, proximity_to_10k })

# NOTE: distinct_conductors is the strongest signal in this scenario. Cash deposited
# into one party's accounts by several different physical conductors on consecutive
# days is the classic smurfing signature. It is only detectable if the teller system
# populates CASH_TRXN.CONDUCTOR_PARTY_ID (see DQ-08).
```

Escalation red flags for the investigator

- Deposits made at three or more distinct branches, particularly geographically dispersed ones.
- Multiple distinct conductors depositing into the same beneficiary account.
- Amounts clustering tightly just below $10,000 (e.g. $9,800, $9,500, $9,700).
- Customer or conductor asks a teller about the reporting threshold, or reduces a deposit on being told a CTR will be filed — this must be reported by the branch and is independently SAR-reportable.
- Cash deposits followed within 24–72 hours by outbound wires of a similar aggregate amount.

SAM-CS-02 — Large Cash Activity — Single or Aggregated

Focus	Frequency	Lookback	Typology addressed
CU	Daily	1 day	Bulk cash placement; proceeds of drug trafficking and other cash-generating crime

Parameter	Medium-segment value	Rationale
SINGLE_TXN_AMT	$25,000	Well above CTR; captures blatant placement
AGG_DAILY_AMT	$40,000	Aggregate across all accounts of the party in one business day
PEER_MULTIPLE	4×	Alert if daily cash exceeds 4× the peer-group mean, even below the absolute thresholds

```
cash_today = SUM(BASE_AMT_USD) WHERE party = p AND cash = TRUE AND TRXN_DT = run_dt

IF cash_today >= AGG_DAILY_AMT
   OR MAX(single cash txn) >= SINGLE_TXN_AMT
   OR cash_today >= PEER_MULTIPLE * peer_profile(p.segment).mean_daily_cash:
      RAISE MATCH(...)

# The PEER_MULTIPLE limb is what makes this scenario risk-sensitive rather than
# merely a second CTR. A $30,000 cash day is unremarkable for a supermarket and
# extraordinary for a software consultancy. Absolute thresholds alone cannot
# express that; a peer-relative limb can.
```

Escalation red flags for the investigator

- Cash volume with no plausible business explanation in the NAICS profile.
- Cash deposits that materially exceed the anticipated activity captured at onboarding.
- Rapid dispersal of the cash within days (see SAM-RM-01).

SAM-CS-03 — Cash Activity Inconsistent with the Business Profile

Focus	Frequency	Lookback	Typology addressed
CU	Weekly	90 days	Trade-based laundering front companies; unlicensed MSBs; cash-intensive business misuse

Parameter	Medium-segment value	Rationale
CASH_RATIO_THRESHOLD	0.60	Cash as a proportion of total credits
PEER_Z_SCORE	2.5	Standard deviations above the NAICS peer mean
MIN_TOTAL_CREDITS	$50,000	Materiality floor over the lookback

```
cash_ratio = profile(p, 90d).cash_credits / profile(p, 90d).total_credits
z          = (cash_ratio - peer(p.naics, p.segment).mean_cash_ratio)
             / peer(p.naics, p.segment).stddev_cash_ratio

IF profile(p, 90d).total_credits >= MIN_TOTAL_CREDITS
   AND cash_ratio > CASH_RATIO_THRESHOLD
   AND z > PEER_Z_SCORE:
      RAISE MATCH(highlights = { cash_ratio, peer_mean, z, naics_desc })

# A law firm, a freight broker or an IT consultancy running a 70% cash ratio is not
# 'a busy business'. It is a business that is either unlicensed money transmission
# or a placement vehicle. The NAICS peer mean is what turns that intuition into a control.
```

Escalation red flags for the investigator

- A non-cash-intensive NAICS code with a high cash ratio.
- Cash deposits without corresponding operating outflows (rent, payroll, suppliers).
- Business address that is a mailbox, a residence, or shared with dozens of other entities.

6.2 Scenario Class — Reportable Transactions

SAM-RT-01 — Possible CTR — Currency Transactions Over the Reporting Threshold

Focus	Frequency	Lookback	Typology addressed
CU	Daily	2 days	CTR completeness assurance (a control on the reporting process, not a suspicion scenario)

Parameter	Medium-segment value	Rationale
CTR_THRESHOLD	$10,000	Statutory (31 CFR 1010.311)
AGGREGATION	Same person, same business day, across all accounts and branches	Statutory aggregation rule

```
-- This scenario exists to prove the CTR filing process is complete.
-- It reconciles what SHOULD have been reported against what WAS reported.

reportable = SUM(cash) BY (party, business_date) WHERE SUM(cash) > CTR_THRESHOLD
filed      = SELECT * FROM CTR_FILINGS WHERE business_date = run_dt

exceptions = reportable  MINUS  filed          -- should be EMPTY

IF exceptions IS NOT EMPTY:
   RAISE OPERATIONAL_EXCEPTION -> BSA Reporting Team, same-day remediation
   -- a missed CTR is a civil-money-penalty exposure; this is a Critical control

-- Separately: parties with a valid 31 CFR 1020.315 CTR EXEMPTION are excluded from
-- the filing requirement but REMAIN FULLY IN SCOPE for every SAM suspicion scenario.
-- Conflating CTR exemption with monitoring exemption is a recurring, serious error.
```

Escalation red flags for the investigator

- Any exception at all — this scenario should normally produce zero.
- A customer newly granted a CTR exemption whose cash volume then materially increases.

6.3 Scenario Class — Rapid Movement of Funds and Layering

SAM-RM-01 — Rapid Movement of Funds — All Activity (Pass-Through)

Focus	Frequency	Lookback	Typology addressed
AC	Daily	10 days	Layering; funnel accounts; money-mule networks; flow-through / pass-through accounts

Parameter	Medium-segment value	Rationale
MIN_CREDIT_AMT	$25,000	Materiality floor
FLOW_THROUGH_PCT	80%	Debits as a percentage of credits within the window
MAX_HOLD_DAYS	5	Funds out within N days of arrival
MIN_TXN_COUNT	4	Excludes a single legitimate in-and-out (e.g. a property purchase)
AVG_BALANCE_RATIO	0.15	Average balance as a fraction of throughput — a true pass-through account holds almost nothing

```
FOR EACH account a:
  credits = SUM(CR) OVER (a, run_dt-10 .. run_dt)
  debits  = SUM(DR) OVER (a, run_dt-10 .. run_dt)
  avg_bal = AVG(daily_balance) OVER (a, run_dt-10 .. run_dt)

  IF credits >= MIN_CREDIT_AMT
     AND debits >= FLOW_THROUGH_PCT * credits
     AND (credits + debits) txn_count >= MIN_TXN_COUNT
     AND avg_bal <= AVG_BALANCE_RATIO * credits
     AND median_hold_period(a) <= MAX_HOLD_DAYS:

        RAISE MATCH(highlights = { credits, debits, flow_pct, avg_bal,
                                   median_hold_days, distinct_originators,
                                   distinct_beneficiaries, channel_mix })

# The four limbs together are what distinguishes a mule account from a busy
# operating account. A trading company also moves money in and out quickly - but
# it holds working capital (avg_bal fails) and its counterparties repeat (distinct
# counterparty count fails). Any single limb on its own is a false-positive engine.
```

Escalation red flags for the investigator

- Many-to-one inbound (dozens of small credits) followed by one-to-few outbound.
- Counterparties with no discernible commercial relationship to the account holder.
- Account opened recently, high throughput, near-zero retained balance.
- Inbound P2P (Zelle/RTP) aggregating and departing by wire — the classic mule signature.

SAM-RM-02 — Deposits Followed by Rapid Wire Transfer Out

Focus	Frequency	Lookback	Typology addressed
CU	Daily	10 days	Placement-to-layering conversion; funnel accounts serving TCO/TBML networks

Parameter	Medium-segment value	Rationale
MIN_CASH_IN	$15,000	Aggregate cash-in over the window
WIRE_OUT_PCT	70%	Outbound wire value as a percentage of cash-in
MAX_LAG_DAYS	5	Between the cash and the wire

```
cash_in   = SUM(cash credits)  OVER (p, run_dt-10 .. run_dt)
wire_out  = SUM(wire debits)   OVER (p, run_dt-10 .. run_dt)

IF cash_in >= MIN_CASH_IN
   AND wire_out >= WIRE_OUT_PCT * cash_in
   AND temporal_sequence(cash_in, wire_out, MAX_LAG_DAYS):   -- cash precedes wire

      score_uplift = 0
      IF beneficiary_country IN HIGH_RISK_JURISDICTIONS: score_uplift += 20
      IF distinct_deposit_branches >= 3:                    score_uplift += 15
      IF beneficiary is a first-time counterparty:          score_uplift += 10

      RAISE MATCH(score_uplift = score_uplift, ...)
```

Escalation red flags for the investigator

- Cash deposited in one geography and wired to an unrelated one.
- Beneficiary in a jurisdiction with no commercial nexus to the customer's stated business.
- Round-dollar wires immediately after non-round cash aggregation.

SAM-RM-03 — Wire Transfers Between Unrelated Parties / Journals Between Unrelated Accounts

Focus	Frequency	Lookback	Typology addressed
AC	Daily	30 days	Layering through internal book transfers; nominee and mule networks

Parameter	Medium-segment value	Rationale
MIN_AGG_AMT	$20,000	Aggregate over the window
MIN_COUNTERPARTY_CNT	3	Distinct unrelated counterparties
UNRELATED_TEST	No shared tax ID · no shared party · no household link · no formal relationship on file	Definition of 'unrelated'

```
FOR EACH internal transfer (from_acct, to_acct, amt) IN window:

  related = shares_tax_id(from, to)
         OR shares_party(from, to)
         OR same_household(from, to)
         OR formal_relationship_on_file(from, to)   -- e.g. parent/subsidiary, POA

  IF NOT related: add to unrelated_set[from_acct]

IF SUM(unrelated_set[a].amt) >= MIN_AGG_AMT
   AND COUNT(DISTINCT unrelated_set[a].counterparty) >= MIN_COUNTERPARTY_CNT:
      RAISE MATCH(...)

# Internal book transfers leave no external footprint and are therefore attractive
# for layering. They are also the transfers most often EXCLUDED from monitoring on
# the reasoning that 'the money never left the bank'. That reasoning is wrong and
# has been cited in enforcement actions.
```

Escalation red flags for the investigator

- Transfers between customers with no commercial or familial link.
- Circular flows (A→B→C→A) — visible in the ActOne link analysis view.
- Transfers that consistently sit just below an internal review threshold.

6.4 Scenario Class — Behavioural Deviation

SAM-BD-01 — Significant Change from Previous Average Activity

Focus	Frequency	Lookback	Typology addressed
AC	Month-start	12 months	Account takeover; sale/rental of an account; a legitimate business converted to a laundering vehicle

Parameter	Medium-segment value	Rationale
PCT_INCREASE	300%	Current-month activity vs the trailing 12-month mean
MIN_CURRENT_AMT	$50,000	Materiality floor — prevents alerting on tiny accounts where percentages explode
Z_SCORE_THRESHOLD	3.0	Standard deviations above the entity's own trailing mean
MIN_HISTORY_MONTHS	6	Without adequate history there is no baseline; such accounts are covered by SAM-BD-03 instead

```
hist = profile(a, trailing 12 months EXCLUDING current month)
cur  = profile(a, current month)

IF hist.months_observed >= MIN_HISTORY_MONTHS
   AND cur.total_amt >= MIN_CURRENT_AMT
   AND cur.total_amt >= (1 + PCT_INCREASE/100) * hist.mean_monthly_amt
   AND (cur.total_amt - hist.mean_monthly_amt) / NULLIF(hist.stddev_monthly_amt,0)
       > Z_SCORE_THRESHOLD:

      RAISE MATCH(highlights = { cur_amt, hist_mean, hist_stddev, z_score,
                                 new_counterparties, new_channels, new_countries })

# The dual test (percentage AND z-score) is deliberate. Percentage alone over-alerts
# on volatile accounts whose normal state IS variability; z-score alone under-alerts
# on accounts with a large historical standard deviation. Requiring both means the
# activity must be large in absolute terms AND abnormal relative to the account's
# own demonstrated volatility.
```

Escalation red flags for the investigator

- Change coincides with a new signatory, a new address, or a new device.
- New counterparties concentrated in high-risk jurisdictions.
- A dormant or low-activity account suddenly transacting at scale — cross-reference SAM-BD-03.

SAM-BD-02 — Activity Materially Exceeds Anticipated / Declared Activity (CDD Variance)

Focus	Frequency	Lookback	Typology addressed
CU	Month-end	1 month	Misrepresentation at onboarding; undisclosed business lines; income inconsistent with declared profile

Parameter	Medium-segment value	Rationale
VARIANCE_PCT	200%	Actual credits vs anticipated monthly credits captured at onboarding
MIN_ABSOLUTE_VAR	$20,000	Absolute floor
CONSECUTIVE_MONTHS	2	Two consecutive breaches — one month may be a genuine one-off

```
antic  = ACCOUNT.ANTIC_MTHLY_CR_AMT (aggregated across the party's accounts)
actual = profile(p, current month).total_credits

IF antic IS NULL OR antic = 0:
   RAISE DQ_EXCEPTION('anticipated activity not captured')   -- see DQ-10
   -- NOT a silent skip. A null denominator must never quietly disable a scenario.

ELIF actual >= (1 + VARIANCE_PCT/100) * antic
     AND (actual - antic) >= MIN_ABSOLUTE_VAR
     AND breach_streak(p) >= CONSECUTIVE_MONTHS:
        RAISE MATCH(highlights = { antic, actual, variance_pct, streak_months })

# This scenario is the operational teeth of the CDD Rule's 'ongoing monitoring to
# maintain and update customer information' obligation. It is also the scenario that
# most ruthlessly exposes a weak onboarding process: if relationship managers enter
# a default figure for anticipated activity to clear a mandatory field, this scenario
# either floods or goes silent - and DQ-10 exists precisely to catch that.
```

Escalation red flags for the investigator

- A material and sustained gap between declared and actual activity.
- The customer cannot explain the source of the excess when asked via RFI.
- Activity types (e.g. international wires) that were never disclosed at onboarding.

SAM-BD-03 — Escalation in a Dormant or Inactive Account

Focus	Frequency	Lookback	Typology addressed
AC	Daily	1 day (against a 180-day dormancy profile)	Account takeover; sold or rented accounts; sleeper mule accounts activated on demand

Parameter	Medium-segment value	Rationale
DORMANCY_DAYS	180	No customer-initiated activity
REACTIVATION_AMT	$10,000	Single or aggregated post-dormancy activity
WITHDRAWAL_MULTIPLIER	0.5×	Withdrawal thresholds are set at half the deposit thresholds — the bank's exposure is greater on outbound funds

Parameter	Medium-segment value	Rationale
OBSERVATION_WINDOW	30 days post-reactivation	Heightened monitoring window

```
IF account_dormant_for(a) >= DORMANCY_DAYS
   AND activity_today(a) > 0:

      IF outbound_amt >= REACTIVATION_AMT * WITHDRAWAL_MULTIPLIER
         OR inbound_amt >= REACTIVATION_AMT:

            RAISE MATCH(priority_uplift = HIGH)
            SET a.enhanced_monitoring_until = run_dt + OBSERVATION_WINDOW

# Asymmetric thresholds are a deliberate design choice, not an oversight. A dormant
# account receiving $10,000 may be a forgotten inheritance. A dormant account SENDING
# $5,000 to a new beneficiary is materially more likely to be an account that has
# changed hands. The control is calibrated to the direction of the bank's exposure.
```

Escalation red flags for the investigator

- Reactivation immediately preceded by a credential reset, a new device, or a change of address.
- Beneficiaries entirely new to the account's history.
- Reactivation across a cluster of dormant accounts in the same period — check the link analysis.

6.5 Scenario Class — High-Risk Entity, Geography and Correspondent Banking

SAM-HR-01 — High-Risk Transactions — Focal High-Risk Entity

Focus	Frequency	Lookback	Typology addressed
AC	Daily	1 day	Elevated surveillance of parties already assessed as high risk

Parameter	Medium-segment value	Rationale
ENTITY_RISK_MIN	7	Effective risk value of the focal entity
AMOUNT_THRESHOLD	$5,000	Deliberately low — the customer's risk is the trigger, the amount is only a materiality gate
ANY_CROSS_BORDER	TRUE	Any cross-border movement by a risk-7+ entity is in scope regardless of amount

```
IF entity_effective_risk(a.party) >= ENTITY_RISK_MIN:
   IF txn.BASE_AMT_USD >= AMOUNT_THRESHOLD
      OR (ANY_CROSS_BORDER AND txn.is_cross_border):
         RAISE MATCH(...)
```

Escalation red flags for the investigator

- Any activity at all is reviewable; the question is not 'is this large' but 'is this consistent'.
- Newly-designated risk (a recent PEP identification or adverse-media hit) with pre-existing activity that was never reviewed under the elevated risk lens.

SAM-HR-02 — High-Risk Transactions — High-Risk Counterparty or Corridor

Focus	Frequency	Lookback	Typology addressed
AC	Daily	14 days	Exposure to sanctioned, FATF-listed, or otherwise high-risk jurisdictions and counterparties

Parameter	Medium-segment value	Rationale
CORRIDOR_RISK_MIN	6	Country risk value of the originating or beneficiary jurisdiction
AGG_AMOUNT	$15,000	Aggregate to/from high-risk corridors over the lookback
TXN_COUNT_MIN	2	Distinguishes a pattern from a one-off

```
hr_txns = txns WHERE country_risk(counterparty_country) >= CORRIDOR_RISK_MIN
               AND TRXN_DT >= run_dt - 14

IF SUM(hr_txns.amt) >= AGG_AMOUNT AND COUNT(hr_txns) >= TXN_COUNT_MIN:
   RAISE MATCH(highlights = { corridor_list, total_amt, txn_count,
                              beneficiary_names, intermediary_banks })

# The intermediary bank chain matters as much as the endpoints. Funds routed to a
# low-risk country THROUGH an intermediary in a high-risk one is a nesting pattern.
# WIRE_TRXN.INTERMED_BIC is therefore in scope for the corridor test, not just
# ORIG_BANK_CTRY and BENE_BANK_CTRY.
```

Escalation red flags for the investigator

- A jurisdiction with no commercial nexus to the customer's stated business.
- Payments routed through an unnecessary intermediary in a secrecy jurisdiction.
- Beneficiary names that are generic ('Trading Ltd', 'Global Import Export').

SAM-CB-01 — Correspondent Banking — Nesting and Unexpected Volume

Focus	Frequency	Lookback	Typology addressed
CB	Weekly	30 days	PATRIOT Act §312/§313: nested correspondent relationships, downstream clearing for undisclosed banks

Parameter	Medium-segment value	Rationale
VOLUME_DEVIATION	250%	Versus the volume declared in the correspondent's due diligence questionnaire
UNIQUE_ORIGINATOR_CNT	50	Distinct originators per week routed through one respondent
FOREIGN_FI_ORIG_FLAG	TRUE	Originator itself appears to be a financial institution not disclosed in the relationship

```
FOR EACH respondent bank r:

  declared = r.due_diligence.expected_monthly_volume
  actual   = profile(r, 30d).total_cleared_volume

  nested_indicator = COUNT(DISTINCT originators WHERE originator_is_FI(originator)
                                                  AND originator NOT IN r.disclosed_affiliates)

  IF actual >= (1 + VOLUME_DEVIATION/100) * declared
     OR COUNT(DISTINCT originators) >= UNIQUE_ORIGINATOR_CNT
     OR nested_indicator > 0:
        RAISE MATCH(priority = VERY_HIGH)

# Nesting is the single largest USD-clearing risk a Tier 1 bank carries. The bank
# knows its respondent; it does not know its respondent's respondent. Detecting a
# financial institution appearing as an ORIGINATOR in a respondent's traffic is the
# principal automated control available, and it depends entirely on the completeness
# of the originator fields in the MT202/pacs.009 message.
```

Escalation red flags for the investigator

- Originators that are themselves banks or MSBs not disclosed in the correspondent questionnaire.
- Volume materially in excess of the declared expectation.
- Payment volume from jurisdictions the respondent stated it does not serve.
- Any indication of a shell bank in the chain — prohibited outright under §313.

6.6 Scenario Class — Internal Accounts, Employees and Structural Patterns

SAM-IN-01 — Patterns of Funds Transfer Between Internal Accounts and Customers

Focus	Frequency	Lookback	Typology addressed
AC	Weekly	7 days	Insider facilitation; exclusive or concealed relationships between staff-controlled and customer accounts

Parameter	Medium-segment value	Rationale
MIN_TXN_COUNT	5	Over the lookback
EXCLUSIVITY_PCT	75%	Proportion of an account's activity concentrated with a single counterparty
EMPLOYEE_NEXUS	TRUE	Either party linked to an employee record

```
FOR EACH pair (internal_or_employee_acct i, customer_acct c):
  pair_vol  = SUM(amt) BETWEEN i AND c OVER 7d
  total_vol = SUM(amt) FOR c OVER 7d

  IF COUNT(pair_txns) >= MIN_TXN_COUNT
     AND pair_vol >= EXCLUSIVITY_PCT * total_vol:
        RAISE MATCH(route_to = SPECIAL_INVESTIGATIONS_UNIT)

# Employee-nexus alerts NEVER enter the standard queue. They route directly to the
# SIU under restricted visibility, because the standard queue may contain the
# subject's own colleagues - or the subject.
```

Escalation red flags for the investigator

- Concentration of a customer's activity with a single internal or employee-linked account.
- Employee accessing or servicing the account of a customer with whom they transact.
- Suppression or trusted-pair requests raised by, or on behalf of, a party with an employee nexus — an immediate escalation trigger.

SAM-IN-02 — Round-Dollar and Repeating-Amount Transactions

Focus	Frequency	Lookback	Typology addressed
CU	Weekly	30 days	Invoice fabrication; layering with artificial amounts; contrived commercial narratives

Parameter	Medium-segment value	Rationale
ROUND_AMOUNT_RULE	Divisible by $1,000 with no cents	Definition
MIN_COUNT	6	Round-dollar transactions in the window
ROUND_RATIO	70%	Round-dollar as a proportion of all transactions
MIN_AGG	$50,000	Materiality floor

```
round_txns = txns WHERE MOD(BASE_AMT_USD, 1000) = 0 AND cents = 0

IF COUNT(round_txns) >= MIN_COUNT
   AND COUNT(round_txns) >= ROUND_RATIO * COUNT(all_txns)
   AND SUM(round_txns.amt) >= MIN_AGG:
      RAISE MATCH(...)

# Genuine commercial activity is untidy: $47,283.19, $12,904.55. Sustained perfectly
# round amounts across many counterparties indicate amounts that were decided rather
# than earned. Weak alone; powerful in combination with SAM-RM-01 and SAM-TB-01,
# which is precisely what alert correlation (Ch. 9.3) exists to surface.
```

Escalation red flags for the investigator

- Round-dollar wires against invoices that are themselves round.
- Identical amounts repeating across different counterparties.
- Round amounts with vague or absent remittance information.

6.7 Scenario Class — FinCEN Advisory-Driven Typologies

These scenarios exist because FinCEN has published specific red-flag guidance and expects it to be reflected in monitoring. When one of them escalates to a SAR, the corresponding FinCEN **key term** must be entered in the narrative and, where applicable, in the SAR field 'Filing Institution Note to FinCEN' — this is how FinCEN aggregates filings by typology and it is checked at examination.

SAM-HT-01 — Human Trafficking and Human Smuggling Indicators

Focus	Frequency	Lookback	Typology addressed
CU	Weekly	90 days	FinCEN Advisory FIN-2020-A008 (human trafficking); key term: HUMAN TRAFFICKING FIN-2020-A008

Parameter	Medium-segment value	Rationale
COMPOSITE_TRIGGER	≥ 3 indicators	No single indicator alerts; the typology is a constellation
INDICATORS	Frequent late-night ATM cash withdrawals near hotels/motels · third-party deposits from many individuals · payments to online classified-advertising platforms · transactions in multiple states along known corridors · shared address across many young account holders · minimal ordinary living expenses (no grocery, utility or rent)	Derived from the advisory

```
indicator_count = 0
IF late_night_atm_near_lodging(p, 90d) >= 6:             indicator_count += 1
IF distinct_third_party_cash_depositors(p, 90d) >= 5:    indicator_count += 1
IF payments_to_classified_ad_platforms(p, 90d) > 0:      indicator_count += 1
IF distinct_states_of_activity(p, 90d) >= 4:             indicator_count += 1
IF shared_address_party_count(p) >= 4:                   indicator_count += 1
IF normal_living_expense_ratio(p, 90d) < 0.10:           indicator_count += 1

IF indicator_count >= 3:
   RAISE MATCH(priority = VERY_HIGH, route_to = SIU,
               narrative_key_term = 'HUMAN TRAFFICKING FIN-2020-A008')

# Composite scenarios are the correct architecture for advisory typologies. Each
# individual behaviour is common and innocuous; the co-occurrence is not. Building
# these as six separate scenarios would generate six times the alerts and none of
# the insight.
```

Escalation red flags for the investigator

- Any indication that a third party controls the customer's account or accompanies them to the branch.
- The customer is unable to speak for themselves or appears coached.
- Escalate to the SIU immediately — never send an RFI to the branch that could tip off a controller.

SAM-EE-01 — Elder Financial Exploitation

Focus	Frequency	Lookback	Typology addressed
CU	Weekly	60 days	FinCEN Advisory FIN-2022-A002; key term: ELDER FINANCIAL EXPLOITATION FIN-2022-A002

Parameter	Medium-segment value	Rationale
AGE_THRESHOLD	65	Customer age
BEHAVIOUR_CHANGE	z ≥ 2.5	Versus the customer's own trailing baseline
NEW_PARTY_ON_ACCOUNT	Added within 90 days	New signatory, POA or joint holder
OUTBOUND_INDICATORS	Wires abroad · crypto exchange purchases · gift cards · new unrelated beneficiary	Common scam rails

```
IF p.age >= AGE_THRESHOLD:
   uplift = 0
   IF behaviour_z_score(p, 60d) >= 2.5:                      uplift += 1
   IF new_signatory_or_poa_within(p, 90d):                   uplift += 1
   IF outbound_to(crypto_exchange | gift_card | new_foreign_bene): uplift += 1
   IF account_takeover_signals(p):                           uplift += 1
   IF uplift >= 2: RAISE MATCH(priority = HIGH,
                    narrative_key_term = 'ELDER FINANCIAL EXPLOITATION FIN-2022-A002')

# NOTE: This scenario protects a customer from a third party, and often that third
# party is a family member with signing authority. RFI handling therefore differs:
# contact is made with the CUSTOMER directly, never through the new signatory.
```

Escalation red flags for the investigator

- A new signatory or POA added shortly before activity changes.
- Uncharacteristic wires abroad, crypto purchases, or gift-card spending.
- The customer appears confused about transactions on their own account.

- Consider a referral to Adult Protective Services in parallel with the SAR (permitted, and encouraged by the advisory — the SAR itself remains confidential).

SAM-VA-01 — Virtual Asset Exposure and Convertible Virtual Currency

Focus	Frequency	Lookback	Typology addressed
CU	Weekly	30 days	FinCEN CVC guidance; ransomware (FIN-2021-A004); unregistered money transmission

Parameter	Medium-segment value	Rationale
VASP_AGG_AMT	$10,000	Aggregate to/from identified virtual-asset service providers
UNREGISTERED_VASP	Any amount	Counterparty is a VASP not registered as an MSB with FinCEN
P2P_PATTERN	≥ 5 counterparties	Many-to-one inbound followed by a single VASP outbound — the P2P-exchange / unlicensed-transmitter signature
MIXER_NEXUS	Any	Counterparty attributed by the blockchain-analytics vendor to a mixer, tumbler or sanctioned service

```
vasp_txns = txns WHERE counterparty IN vasp_registry

IF SUM(vasp_txns.amt, 30d) >= VASP_AGG_AMT
   OR ANY(vasp_txns.counterparty NOT IN fincen_msb_registry)
   OR ANY(vasp_txns.counterparty IN mixer_or_sanctioned_service_list)
   OR (COUNT(DISTINCT inbound_p2p_counterparties) >= 5
       AND single_large_vasp_outbound_within(3 days)):

      RAISE MATCH(enrich_with = blockchain_analytics_attribution)

# The bank does not monitor the chain; it monitors the FIAT ON-RAMP AND OFF-RAMP,
# which is the point at which it has visibility and legal obligation. Attribution
# data from the analytics vendor is enrichment, not detection - the alert is raised
# on the bank-visible fiat leg.
```

Escalation red flags for the investigator

- Counterparty VASP not registered with FinCEN as a money transmitter.
- Exposure attributed to a mixer, a darknet market, or a sanctioned exchange.
- A pattern of inbound P2P payments consolidated and sent to an exchange — likely unlicensed transmission.
- Ransomware profile: sudden large fiat-to-crypto conversion by a corporate customer under duress.

SAM-TB-01 — Trade-Based Money Laundering Indicators

Focus	Frequency	Lookback	Typology addressed
CU	Monthly	180 days	TBML: over/under-invoicing, phantom shipments, carousel trade; FATF and FinCEN trade advisories

Parameter	Medium-segment value	Rationale
PRICE_DEVIATION	±40%	Unit price versus the world-trade reference price for the HS code
ROUTE_ANOMALY	TRUE	Goods routed via a jurisdiction with no logistical rationale
DUAL_USE_HS	Any	HS code appears on a dual-use or strategic-goods control list
PAYMENT_MISMATCH	±25%	Payment value versus documented invoice value
CIRCULAR_TRADE	TRUE	The same goods invoiced between related parties more than once

```
FOR EACH trade transaction t (LC, guarantee, documentary collection, open account):

  ref_price = world_reference_price(t.hs_code, t.period)
  deviation = ABS(t.unit_price - ref_price) / ref_price

  flags = []
  IF deviation > PRICE_DEVIATION:                          flags += 'PRICE_ANOMALY'
  IF NOT route_is_logical(t.origin, t.transit, t.dest):    flags += 'ROUTE_ANOMALY'
  IF t.hs_code IN DUAL_USE_LIST:                           flags += 'DUAL_USE'
  IF ABS(t.payment_amt - t.invoice_amt)/t.invoice_amt
        > PAYMENT_MISMATCH:                                flags += 'PAYMENT_MISMATCH'
  IF circular_trade_detected(t.buyer, t.seller, t.goods):  flags += 'CIRCULAR'
  IF vessel_or_port IN SANCTIONED_PORTS:                   flags += 'SANCTIONS_NEXUS'

  IF COUNT(flags) >= 2: RAISE MATCH(highlights = flags, route_to = TRADE_AML_TEAM)

# TBML detection is only as good as the structured capture of goods description,
# HS code, quantity, unit price, ports and vessel. Where the trade system stores
# these as free text, the scenario degrades to keyword scanning - which is why the
# structured capture of these fields is itself a tracked remediation item.
```

Escalation red flags for the investigator

- Unit prices materially off market for the commodity.
- Goods description inconsistent with the customer's stated business.
- Shipping route with no commercial logic; transshipment through a sanctions-adjacent port.
- The same goods appearing to move repeatedly between related parties (carousel).
- Payment from, or to, a third party unrelated to the documented trade.

6.8 Scenario Class — Sanctions and Watch-List Nexus

SAM-WL-01 — Watch-List Nexus Cross-Referral from WLF

Focus	Frequency	Lookback	Typology addressed
CU	Event-driven (real time)	n/a	OFAC nexus; 314(a) subject match; PEP identification

Parameter	Medium-segment value	Rationale
TRIGGER	Confirmed WLF true match, or 314(a) subject match	Event
ACTION	Auto-create an ActOne case, priority Very High	No triage tier — straight to L2/SIU

```
ON EVENT wlf_true_match(party, list, entry) OR fincen_314a_match(party):

   1. Block / reject the payment per OFAC rules (WLF, real time)
   2. Force party.CRR = VERY_HIGH (override the CDD model)
   3. Auto-create ActOne CASE(type = SANCTIONS_NEXUS, priority = VERY_HIGH)
   4. Retro-scan: pull ALL activity for the party over the trailing 12 months
      into the case as evidence (a 314(a) match is a signal about the PAST,
      not merely the present)
   5. Notify BSA Officer within 1 hour; OFAC report within 10 business days

# A 314(a) match does not itself require a SAR. It requires the bank to search its
# records and report matches to FinCEN. But it is highly probative information that
# must be fed back into the risk model and the monitoring lens - and MNB policy is
# that a 314(a) match always triggers a documented review of historical activity.
```

Escalation red flags for the investigator

- Any true positive is Very High by definition.
- Consider whether prior alerts on the same party, previously closed NFA, should now be reopened in light of the new information — this is an explicit checklist item in the case template.

7. Anatomy of a Behavioural Detection Run

This chapter is the technical core of the framework. It describes what physically happens between midnight and 07:30 every night, in what order, with what dependencies, and what happens when a step fails. If an examiner asks the FIU to prove that a specific transaction on a specific date was evaluated against a specific threshold, the answer is reconstructed from the artefacts described here.

7.1 The Nightly Detection Cycle

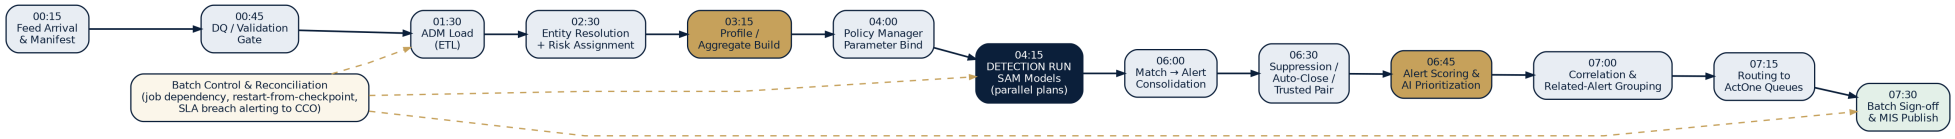

Figure 6 — The nightly detection cycle with target completion times. Each stage is a gated dependency; the detection run does not start until the data gate passes.

Stage	Job	Depends on	Target	Failure action
1	Feed arrival & manifest validation	Source cut-off 23:59	00:15	Page on-call; escalate to feed owner at +30 min
2	Technical & business DQ validation	Stage 1	00:45	Critical breach → halt batch, notify BSA Officer
3	Standardisation, FX normalisation, enrichment	Stage 2	01:15	Retry ×2, then halt
4	ADM load (SCD-2 reference, insert-only transactions)	Stage 3	01:30	Rollback to checkpoint, re-run
5	Entity resolution & party linking	Stage 4	02:30	Unresolved records queue to the data steward; batch proceeds with prior-day Party IDs
6	Risk assignment (entity effective risk, activity risk)	Stage 5 + F-WL-01	03:00	Halt — stale watch-list data is a Critical breach (DQ-14)
7	Profile / aggregate build	Stage 6	03:15	Halt; profiles are the substrate of every scenario
8	Policy Manager parameter bind	Stage 7	04:00	Halt; a run with unbound parameters is prohibited
9	**Detection run — SAM scenario plans (parallel)**	Stage 8	04:15–06:00	Per-plan restart from checkpoint; partial completion is logged and the incomplete plans are re-run
10	Match → alert consolidation	Stage 9	06:00	Retry
11	Suppression, trusted pair, auto-close	Stage 10	06:30	Retry
12	Match & alert scoring; AI prioritisation	Stage 11	06:45	Fall back to rules-only scoring; alerts are never withheld because the AI layer failed
13	Correlation & related-alert grouping	Stage 12	07:00	Non-blocking; alerts release ungrouped
14	Routing to ActOne queues	Stage 13	07:15	Retry; manual release available to the Actimize admin under dual control
15	Batch sign-off & MIS publication	Stage 14	07:30	Sign-off withheld; incident raised

The detection gate
Stages 1–8 exist to protect stage 9. A detection run executed on incomplete data produces a *false assurance of coverage* — the worst possible outcome, because the bank now has a clean batch record and an undetected pattern. MNB's standing rule: **it is always better to run late than to run wrong.** A delayed batch is an operational issue; a silently incomplete batch is a regulatory one.

7.2 Execution Semantics — From Transaction to Alert

Actimize executes detection through a hierarchy: a **plan** schedules one or more **jobs**; each job executes a **model** (scenario); each model evaluates one or more **patterns** over a set of **binding sets**; each binding set that satisfies the pattern produces a **match**; matches consolidate into **alerts**. Understanding this hierarchy is what allows an investigator to answer 'why did this alert fire?' precisely rather than approximately.

Plan → job → model → pattern → binding → match → alert

```
PLAN  AML_NIGHTLY_RETAIL                    (schedule: daily 04:15, parallelism 8)
 |
 +-- JOB  J_STRUCTURING          -> MODEL SAM-CS-01
 |     |
 |     +-- PATTERN P1  'multiple cash in the sub-CTR band'
 |     +-- PATTERN P2  'cash + monetary instrument combination'
 |           |
 |           +-- BINDING SET  (PARTY_ID, run_dt-7 .. run_dt)
 |                 |
 |                 +-- evaluate  -> MATCH  (matched_records = [txn1..txnN])
 |                                         (match_score     = 62)
 |
 +-- JOB  J_RAPID_MOVEMENT       -> MODEL SAM-RM-01
 +-- JOB  J_DORMANCY             -> MODEL SAM-BD-03
 +-- JOB  J_HIGH_RISK_ENTITY     -> MODEL SAM-HR-01
 ...

POST-PROCESSING (after all jobs in the plan complete):
  consolidate:   matches GROUP BY (scenario_id, focal_type, focal_id, business_dt)
                 -> one MULTI-MATCH ALERT per group
  score:         alert_score = MAX(match_score)      [Highest Match Score strategy]
  suppress:      apply suppression rules and trusted pairs -> auto-close where matched
  auto-close:    apply auto-close rules              -> close with reason code
  correlate:     group alerts across scenarios by focal entity and by network linkage
  route:         assign to queue by (priority, LoB, segment, skill)
```

7.2.1 Frequency, lookback and de-duplication

These three parameters interact in ways that are the source of most scenario-design defects.

Configuration	Behaviour	Assessment
Frequency = Lookback (e.g. 1 day / 1 day)	Each run examines exactly the data since the last run. No overlap, no gaps.	Correct for point-in-time scenarios (large cash, high-risk entity).
Frequency < Lookback (e.g. daily / 30 days)	Runs overlap; the same transaction is examined on 30 consecutive nights.	**The normal configuration for pattern scenarios.** Requires de-duplication, otherwise the same behaviour alerts 30 times.
Frequency > Lookback (e.g. weekly / 1 day)	Six days out of every seven are never examined.	**Prohibited.** This is a blind window. It is also easy to configure by accident, and it is silent — no error is raised, alerts simply never fire. Policy Manager enforces a hard validation rule preventing this configuration from being saved.

```
-- De-duplication logic applied at match consolidation.
-- Objective: the same behaviour by the same entity must not re-alert every night
-- for the duration of the lookback window.

FOR EACH candidate_match m:

  prior = SELECT * FROM ALERT
          WHERE scenario_id = m.scenario_id
            AND focal_id    = m.focal_id
            AND business_dt >= (m.business_dt - scenario.frequency_period)
            AND status NOT IN ('REOPENED')

  IF prior EXISTS:
     IF m.matched_records IS A SUBSET OF prior.matched_records:
        SUPPRESS m  (reason = 'DUPLICATE_OF_' || prior.alert_id)   -- no new behaviour
     ELIF m.new_record_amt >= scenario.material_increment_threshold:
        LINK m TO prior AS 'ESCALATING_ACTIVITY'  -- new, material activity: re-alert
        RAISE ALERT(priority_uplift = TRUE)
     ELSE:
        APPEND m.matched_records TO prior          -- enrich the open alert in place
  ELSE:
     RAISE ALERT(m)

-- The middle branch matters. A structurer who deposits $9,500 on Monday (alert
-- raised) and another $9,500 on Wednesday has ESCALATED. Naive de-duplication
-- would silently swallow Wednesday because an alert already exists for the
-- pattern. MNB's rule: new material activity re-alerts and uplifts priority.
```

7.2.2 The profile build — where the real cost lives

Scenarios do not scan raw transaction history at run time; at MNB's volumes (~48 million postings per day) that would be impossible inside the batch window. Instead, the profile build pre-aggregates rolling windows once, and every scenario reads from those aggregates. The profile build is therefore the longest and most performance-critical job in the batch.

Profile / aggregate build — illustrative SQL

```
-- Illustrative profile build (Oracle syntax). Executed once per night; every
-- CU-focus scenario reads from the resulting PROFILE_AGG table.

MERGE INTO PROFILE_AGG tgt
USING (
  SELECT  p.PARTY_ID                                          AS FOCAL_ID,
          'CU'                                                AS FOCAL_TYPE,
          w.WINDOW_CD,
          SUM(CASE WHEN t.DRCR = 'CR' THEN t.BASE_AMT_USD ELSE 0 END)  AS TOT_CR_AMT,
          SUM(CASE WHEN t.DRCR = 'DR' THEN t.BASE_AMT_USD ELSE 0 END)  AS TOT_DR_AMT,
          SUM(CASE WHEN c.TRXN_ID IS NOT NULL AND t.DRCR='CR'
                   THEN t.BASE_AMT_USD ELSE 0 END)                     AS CASH_CR_AMT,
          COUNT(*)                                            AS TXN_CNT,
          COUNT(DISTINCT t.COUNTERPARTY_ID)                   AS CPTY_CNT,
          COUNT(DISTINCT t.CHANNEL_CD)                        AS CHANNEL_CNT,
          AVG(t.BASE_AMT_USD)                                 AS AVG_AMT,
          STDDEV(t.BASE_AMT_USD)                              AS STDDEV_AMT,
          MAX(t.BASE_AMT_USD)                                 AS MAX_AMT
  FROM        TRXN            t
  JOIN        PARTY_ACCT_REL r  ON r.ACCT_ID  = t.ACCT_ID
  JOIN        PARTY          p  ON p.PARTY_ID = r.PARTY_ID
  LEFT JOIN   CASH_TRXN      c  ON c.TRXN_ID  = t.TRXN_ID
  CROSS JOIN  WINDOW_DIM     w                       -- 30 / 60 / 90 / 180 / 365
  WHERE  t.TRXN_DT BETWEEN :RUN_DT - w.DAYS AND :RUN_DT
    AND  r.REL_TYPE IN ('OWNER','SIGNATORY')          -- exclude view-only relationships
    AND  t.EXCLUDE_FLAG = 'N'                         -- internal GL contra entries
  GROUP BY p.PARTY_ID, w.WINDOW_CD
) src
ON (tgt.FOCAL_ID = src.FOCAL_ID AND tgt.WINDOW_CD = src.WINDOW_CD)
WHEN MATCHED THEN UPDATE SET ...
WHEN NOT MATCHED THEN INSERT ...;

-- Peer statistics are then computed over the resulting profiles:
UPDATE PROFILE_AGG pa
SET  PEER_MEAN   = (SELECT AVG(TOT_CR_AMT)    FROM PROFILE_AGG
                    WHERE PEER_GRP_ID = pa.PEER_GRP_ID AND WINDOW_CD = pa.WINDOW_CD),
     PEER_STDDEV = (SELECT STDDEV(TOT_CR_AMT) FROM PROFILE_AGG
                    WHERE PEER_GRP_ID = pa.PEER_GRP_ID AND WINDOW_CD = pa.WINDOW_CD);

UPDATE PROFILE_AGG
SET Z_SCORE = (TOT_CR_AMT - PEER_MEAN) / NULLIF(PEER_STDDEV, 0);

-- PERFORMANCE NOTES (why this completes in ~15 minutes rather than ~6 hours):
--   * TRXN is RANGE-partitioned by TRXN_DT, sub-partitioned by HASH(ACCT_ID)
--   * the MERGE runs with PARALLEL(16) and direct-path writes
--   * PROFILE_AGG is truncate-and-rebuild per window, not incrementally updated,
--     because rolling-window decay is more expensive than a partition rebuild
--   * peer statistics are computed AFTER the profile MERGE, never inside it
```

7.3 Performance and Capacity

Metric	MNB volume	Design implication
Daily postings ingested	~48 million	Partitioned load; direct-path insert; no row-by-row processing anywhere in the pipeline
Parties monitored	~28 million	CU-focus scenarios must read profiles, never raw transactions
Accounts monitored	~41 million	Account-focus profiles are the larger of the two aggregate sets
Scenarios in production	24	Executed across 5 plans; parallelism 8; longest-running scenario governs the critical path
Batch window	04:15 → 06:00	1h 45m for detection. Any scenario exceeding 25 minutes is flagged for optimisation review
Alerts generated (pre-suppression)	~9,500 / night	Volume before auto-close and suppression

Metric	MNB volume	Design implication
Alerts released to queues	~1,150 / night	≈88% reduced by suppression, trusted pairs and auto-close — every one of which is an approved, evidenced, sampled rule
FIU L1 capacity	~1,200 alerts / day	Capacity is a hard constraint on tuning: any proposed threshold change is simulated against it before approval

The capacity trap
Because L1 capacity is finite, there is a permanent, structural incentive to tune thresholds *upward* until alert volume fits the headcount. This is the single most dangerous dynamic in transaction monitoring, and it has been at the centre of several US enforcement actions. MNB's control is procedural and absolute: **capacity may never be cited as a justification in a threshold-change rationale.** If the risk-based threshold produces more alerts than the FIU can handle, the correct response is to increase capacity, improve alert quality through better segmentation, or escalate the resourcing gap to the Board — never to raise the threshold. Every threshold-change submission carries an explicit attestation to this effect.

7.4 Alert Payload

The alert record handed to ActOne carries everything the investigator needs to begin — and everything the bank needs to reconstruct the decision five years later.

Alert payload delivered to ActOne (illustrative, redacted)

```
{
  "alert_id":        "ALT-2026-0712-0043981",
  "scenario_id":     "SAM-CS-01",
  "scenario_name":   "Structuring - Avoidance of Reporting Threshold",
  "business_dt":     "2026-07-11",
  "generated_ts":    "2026-07-12T06:12:44Z",
  "focal": {
     "focal_type":   "CU",
     "party_id":     "P-8842391",
     "legal_name":   "[REDACTED - Individual]",
     "crr":          "ME",
     "segment_id":   "RET-IND-ME-DDA",
     "pep_flag":     "N",
     "effective_risk": 3,
     "acct_ids":     ["0019-448220", "0019-551903"]
  },
  "threshold_set": {
     "version":      "TS-2026-Q2-v4",
     "approved_by":  "AML Model Committee 2026-04-18",
     "params": { "STRUCT_BAND_LOW": 7000, "STRUCT_BAND_HIGH": 9999,
                 "MIN_TXN_COUNT": 3, "AGG_AMOUNT_MIN": 18000,
                 "LOOKBACK_DAYS": 7, "INCLUDE_MI": true }
  },
  "matches": [
    { "match_id": "MTC-...-01", "pattern": "P1_MULTI_CASH_SUB_CTR",
      "match_score": 62,
      "matched_records": [
        {"trxn_id":"T-9931204","dt":"2026-07-06","amt":9500.00,"chan":"BRANCH",
         "branch":"BR-0182","conductor":"P-8842391"},
        {"trxn_id":"T-9944871","dt":"2026-07-08","amt":9200.00,"chan":"BRANCH",
         "branch":"BR-0407","conductor":"P-9910233"},
        {"trxn_id":"T-9958110","dt":"2026-07-09","amt":9750.00,"chan":"BRANCH",
         "branch":"BR-0311","conductor":"P-9910233"}
      ]}
  ],
  "highlights": {
     "txn_count": 3, "total_amt": 28450.00,
     "distinct_branches": 3, "distinct_conductors": 2, "distinct_days": 3,
     "max_single_amt": 9750.00, "closest_to_10k": 9750.00,
     "pct_of_anticipated_monthly": 712.0
  },
  "priors": { "all_prior": 1, "same_scenario_prior": 0,
              "prior_sar": false, "prior_314a": false },
  "scoring": { "rule_score": 62, "ai_score": 0.78, "final_priority": "HIGH" },
  "correlation": { "related_alert_ids": ["ALT-2026-0712-0044102"],
                   "relation": "SAME_FOCAL_ENTITY_DIFFERENT_SCENARIO" },
  "routing": { "queue": "Q-RET-STRUCTURING-HIGH", "sla_due_dt": "2026-07-17",
               "assigned_to": null },
  "lineage": { "batch_id": "BATCH-20260712-01", "feed_ids": ["F-CASH-01","F-TXN-01"],
               "adm_load_ts": "2026-07-12T01:28:11Z" }
}
```


Two elements of this payload deserve emphasis. First, **the threshold set is embedded in the alert itself**, with its version and approval reference — so the bank can prove, years later, exactly what parameters were in force when this alert fired, without reconstructing them from a change log. Second, **lineage** ties the alert back to the specific batch and feed files, which is what makes an end-to-end examiner trace possible.

8. Alert Scoring and Prioritisation

Scoring is a workload-management device. It answers 'what should be looked at first?', not 'is this person a criminal?'. That distinction is not pedantry — an alert score that appears in a SAR narrative, or that is used to justify closing an alert without review, converts a triage tool into an unvalidated decisioning model and creates immediate regulatory exposure.

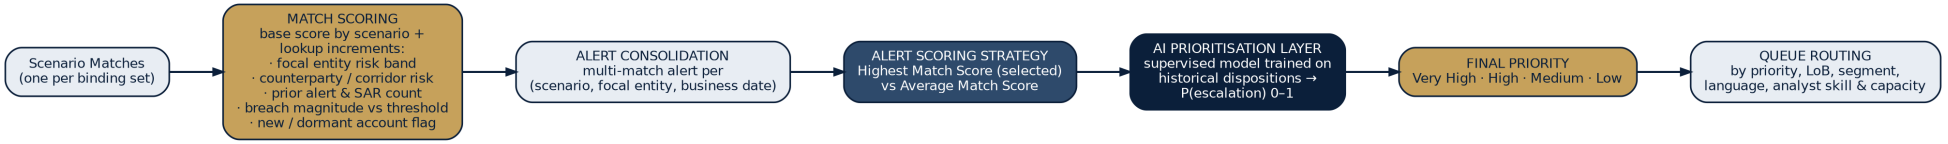

Figure 7 — Scoring pipeline: match scoring → alert consolidation → alert scoring strategy → AI prioritisation → final priority → queue routing.

8.1 Match Scoring

Each match receives a base score from its scenario, then accrues increments from a lookup-based scoring rule. The increments are additive, capped at 100, and every one is a governed parameter.

Scoring factor	Condition	Increment
Scenario base score	Set per scenario by inherent typology severity (e.g. Structuring 40, Human Trafficking 60, Round-Dollar 20)	20 – 60
Focal entity risk	CRR = Very High / High / Medium / Low	+25 / +15 / +5 / 0
Watch-list nexus	Party or counterparty on any risk list	+30
PEP status	Focal entity or beneficial owner is a PEP	+15
Breach magnitude	Aggregate amount ≥ 2× / ≥ 5× the scenario threshold	+10 / +20
Counterparty / corridor risk	Any counterparty in a jurisdiction with country risk ≥ 6	+15
Prior alert history	Same-scenario prior in the last 12 months	+10
Prior SAR	A SAR has previously been filed on the focal entity	+25
314(a) history	Focal entity previously matched a 314(a) subject list	+25
New account	Account open ≤ 180 days	+10
Dormancy reactivation	Behaviour follows a dormancy period	+10
Cash intensity	Cash ratio > peer mean + 2σ	+10
Employee nexus	Any party in the match linked to an employee record	+30 *and* forced route to SIU
Structuring proximity	Amounts within 5% of a reporting threshold	+15

8.2 Alert Scoring Strategy

Where an alert consolidates several matches, two strategies are available. MNB has selected **Highest Match Score**.

Strategy	Definition	Assessment
Highest Match Score (selected)	The alert takes the score of its most severe constituent match.	Conservative and correct for AML. Averaging would let a cluster of low-severity matches dilute a single serious one — the arithmetic would *hide* the very thing the bank is looking for.
Average Match Score	The alert takes the mean of its constituent match scores.	Appropriate in fraud contexts where volume itself is the signal. Rejected for AML at MNB, with the rationale documented in the model development file.

8.3 Priority Bands and Routing

Priority	Score band	Queue & tier	SLA to first action	SLA to disposition
Very High	80 – 100	SIU / L2 direct (no L1 triage)	4 hours	10 business days
High	55 – 79	L1 triage, senior analysts	1 business day	15 business days
Medium	30 – 54	L1 triage, standard	2 business days	20 business days
Low	< 30	L1 triage, standard / auto-close candidate	5 business days	30 business days

Routing is not by priority alone. The queue assignment considers line of business, segment, language, analyst skill certification (e.g. only trade-certified analysts receive SAM-TB-01 alerts) and current workload. Employee-nexus and sanctions-nexus alerts bypass the standard queue structure entirely.

8.4 The AI Prioritisation Layer

MNB deploys a supervised model, trained on ~3 years of historical alert dispositions, that outputs P(escalation) for each alert — the estimated probability that a human investigator would escalate it. The model re-orders the queue within priority bands. It is **advisory only**.

Control	Requirement
Scope of authority	The model **re-orders** alerts. It **cannot close** an alert, cannot suppress one, and cannot prevent one from being generated. Every alert generated by a scenario is still worked by a human or closed under an approved, sampled auto-close rule.
Explainability	Every score is accompanied by its top contributing features in plain language ('score driven by: 3 distinct conductors; prior SAR on focal entity; amounts within 3% of threshold'). An unexplainable score is not usable in an investigation.
Model risk classification	Tier 1 model under SR 11-7. Full development documentation, independent MRM validation before deployment, and annual revalidation.
Bias and fairness testing	Tested for disparate impact across protected characteristics and proxies (ZIP code, name-origin inference, branch geography). Documented and reviewed by Compliance and Legal. A monitoring model that systematically under-prioritises alerts from a protected group is both a fair-lending and a BSA problem.
Champion / challenger	The challenger runs in shadow mode for a full quarter with no queue effect; promotion requires a demonstrated lift in escalation rate at equal or better recall of true positives, and committee approval.
Performance monitoring	Monthly: AUC, precision at k, drift in feature distributions, and — critically — **recall of alerts that ultimately became SARs**. If the model systematically deprioritises alerts that later become SARs, it is withdrawn immediately.
Feedback loop	Investigator dispositions are the training labels. This creates a self-reinforcement risk: if analysts under-escalate a typology, the model learns to deprioritise it. Mitigation: QA-sampled dispositions are weighted more heavily in training, and a stratified random sample of low-scored alerts is worked regardless of score to keep the label distribution honest.

The reason low-scored alerts are still worked
If the bank only ever investigates what the model ranks highly, the model's own blind spots become permanent and invisible — it will never receive a training label that contradicts it. MNB therefore works a mandatory stratified random sample of **5% of Low-priority and auto-closed alerts** every month, worked to full L1 standard. Escalations from that sample are the primary early-warning signal for both threshold and model degradation, and the sample results are reported to the AML Model Committee quarterly.

9. Alert and Case Management in ActOne

9.1 Alert Lifecycle

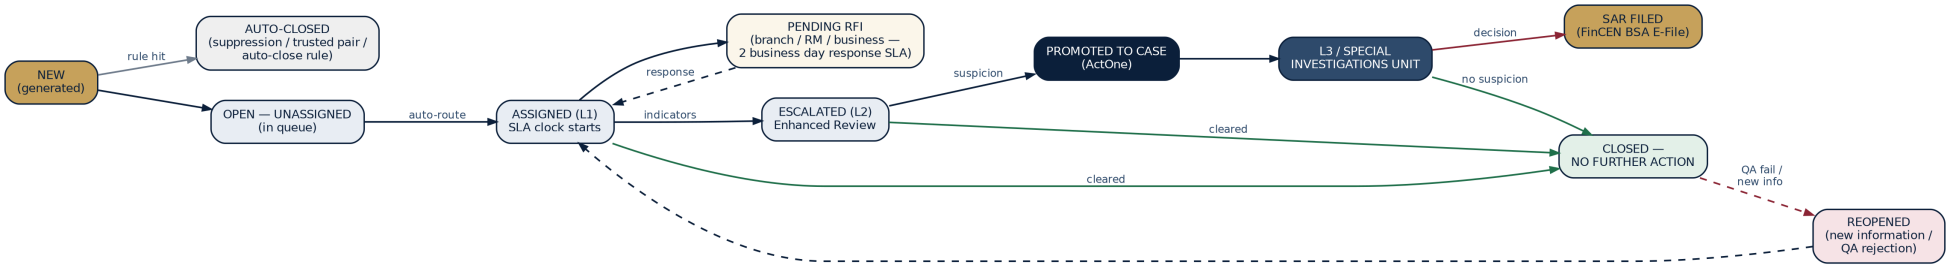

Figure 8 — Alert state machine from generation through triage, escalation, case promotion and disposition, including reopening.

9.2 Queue Model and Assignment

Queue	Population	Tier	Owner
Q-SANCTIONS-NEXUS	WLF true matches, 314(a) matches	L2 / SIU	Sanctions Investigations Lead
Q-SIU-RESTRICTED	Employee nexus, insider, litigation hold	SIU	Head of SIU (restricted visibility)
Q-RET-STRUCTURING	SAM-CS-01/02/03 on retail segments	L1 → L2	Retail FIU Team Lead
Q-RET-BEHAVIOURAL	SAM-BD-01/02/03 on retail segments	L1 → L2	Retail FIU Team Lead
Q-COMM-WIRE	SAM-RM-01/02/03, SAM-HR-02 on commercial	L1 → L2	Commercial FIU Team Lead
Q-CORRESPONDENT	SAM-CB-01	L2 (certified only)	Head of Correspondent AML
Q-TRADE	SAM-TB-01	L2 (trade-certified only)	Head of Trade AML
Q-VULNERABLE	SAM-EE-01, SAM-HT-01	L2 / SIU	Head of Vulnerable Customer Investigations
Q-CRYPTO	SAM-VA-01	L2 (crypto-certified only)	Head of Digital Assets AML
Q-QA-SAMPLE	5% stratified sample of Low / auto-closed alerts	L1 (independent analyst)	Compliance Testing

Assignment is automatic and skill-based. An analyst cannot self-select alerts (which would allow cherry-picking of easy closures), and an analyst cannot be assigned an alert on a customer with whom they have a declared relationship — a conflict check runs at assignment time.

9.3 Related Alerts and Correlation

Investigating an alert in isolation is the most common cause of a missed SAR. Correlation groups alerts before they reach the analyst:

- **Same focal entity, different scenario.** A structuring alert and a rapid-movement alert on the same party in the same week is a placement-and-layering sequence, not two coincidences.
- **Same focal entity, prior alerts.** All alerts on the party in the last 12 months are surfaced, including those closed NFA. A pattern of individually-explicable alerts is itself a pattern.
- **Network linkage.** Alerts on parties connected through shared addresses, shared phone numbers, shared devices, common beneficial owners, or transaction counterparties. This is how a mule network of 40 individually-unremarkable accounts becomes one visible structure.
- **Cross-domain.** A confirmed fraud in IFM-X or a sanctions near-match in WLF on the same party.

Closure of related alerts. Where the whole relationship has been reviewed, related alerts on the same focal entity may be closed together under a single documented rationale — but the closure must state that the relationship as a whole was assessed, not merely the triggering transactions. Bulk closure without a relationship-level review is an audit finding, and one that is easy for QA to detect from the near-identical closure narratives it produces.

9.4 Investigation Standards

9.4.1 L1 Triage — 5 business days

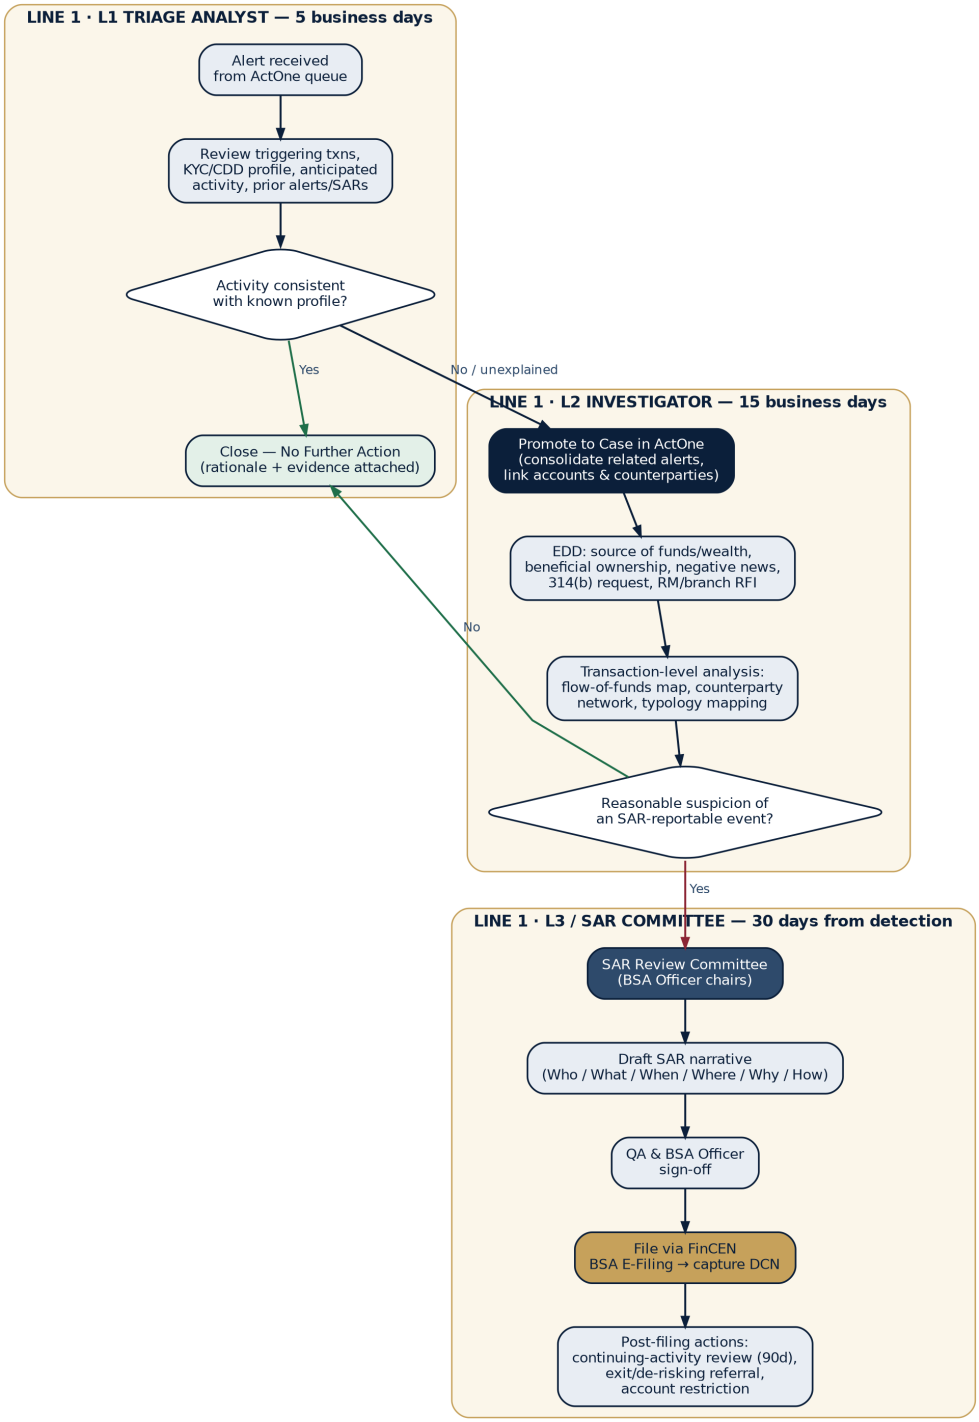

Figure 9 — Investigation workflow across L1 triage, L2 case investigation and the L3 / SAR Committee decision, with SLA at each tier.

The L1 analyst must, at minimum, evidence the following before closing an alert:

#	Mandatory L1 step	Evidence retained
1	Review the triggering transactions and the alert highlights	Screenshot / system reference
2	Review the CDD/KYC profile: occupation, business, source of funds, anticipated activity, beneficial ownership	CDD record reference and date of last review
3	Compare the activity to the anticipated profile	Explicit statement of the comparison and the variance
4	Review prior alerts and prior SARs on the party (12 months minimum; SAR history: all)	Prior-alert list with dispositions

#	Mandatory L1 step	Evidence retained
5	Review the counterparties: are they known, expected, and consistent with the customer's business?	Counterparty analysis
6	Check watch-list, PEP, adverse-media and 314(a) status as at the review date	Screening result and timestamp
7	Reach a documented conclusion	Closure narrative meeting the standard in 9.4.3, or escalation

9.4.2 The RFI to the business

Where the alert cannot be resolved from bank records, the analyst issues a Request for Information to the relationship manager or branch (template at Annex E). Rules:

- The business has **two business days** to respond. Non-response escalates to the LoB Compliance Head at day 3 and to the Country/Divisional Head at day 5.
- **The RFI must never tip off the customer.** Under 31 U.S.C. 5318(g)(2), disclosure of the existence of a SAR — or of the fact that one is being considered — is a federal offence. RFI templates are worded as routine relationship-servicing enquiries and carry a standing no-tipping-off warning. Where the typology involves a controlling third party (human trafficking, elder exploitation, insider), **no RFI is sent at all** and the matter routes directly to the SIU.
- An RFI response is *evidence*, not a conclusion. 'The RM confirms the customer is legitimate' is not a closure rationale. The RM's *explanation*, tested against the transaction record, may be.
- The RFI clock pauses the alert SLA but **does not pause the 30-day SAR clock**, which runs from initial detection irrespective of internal process.

9.4.3 Closure narrative standard

The single most common QA failure is a closure narrative that records a conclusion without the reasoning that supports it. Every closure must answer four questions in writing:

Question	Unacceptable	Acceptable
What triggered the alert?	“Structuring alert.”	“Three cash deposits totalling $28,450 across three branches over four days, each between $9,200 and $9,750, with two distinct conductors.”
What is the customer's expected activity?	“Customer is known to the branch.”	“Customer is a licensed general contractor (NAICS 236220), onboarded 2019, anticipated monthly credits $35,000, historically averaging $31,200 with a 55% cash ratio typical of the trade.”
Does the activity reconcile to that expectation?	“Activity appears normal.”	“The deposits reconcile to three signed progress-payment invoices provided by the customer (attached), for jobs at [addresses], each paid in cash by the homeowner. The deposit dates follow the invoice dates by 1–3 days. Aggregate is within the anticipated envelope.”
Why is there no suspicion?	“No red flags identified.”	“The multi-branch pattern is explained by the job sites' locations; the second conductor is the customer's foreman, identified and known to the branch. The sub-$10,000 amounts reflect the invoice values, not amounts chosen by the customer. No indication of deliberate threshold avoidance.”

The standard to write to
Write every closure as though it will be read, three years from now, by an examiner who already knows the customer was later indicted. If the narrative would still look like a reasonable decision on the information available at the time, it is adequate. If it would look like a box being ticked, it is not. This is the standard, and it is applied in QA sampling.

9.5 Case Management

Attribute	Standard
When a case is created	Automatically on any Very High alert; on any sanctions or 314(a) nexus; on L1 escalation; on the correlation of three or more alerts on one focal entity within 90 days; or at investigator discretion.
What a case contains	All linked alerts across scenarios; all subject parties and related parties; the full transaction set (not merely the triggering transactions); RFI correspondence; EDD findings; negative-news and public-records searches; 314(b) exchanges; flow-of-funds and link-analysis outputs; the disposition rationale.
Case types	AML-SUSPICION · SANCTIONS-NEXUS · FRAUD-AML-REFERRAL · INSIDER · 314(a)-MATCH · REGULATORY-REQUEST · LAW-ENFORCEMENT-SUBPOENA · CONTINUING-ACTIVITY-REVIEW
Ownership	One accountable investigator; contributors may be added. Ownership transfer requires Team Lead approval and is logged.
The four-eyes rule	No investigator may both author and approve a SAR decision — in either direction. A decision *not* to file is reviewed with the same rigour as a decision to file, and is approved at the same level.

9.6 Escalation and Delay Matrix

Trigger	Escalates to	Timing
Alert unassigned beyond SLA	FIU Team Lead	Immediately on breach
Alert open > 5 days past disposition SLA	FIU Team Lead	Daily report
Alert open > 15 days past SLA	Head of FIU	Daily report
Alert open > 30 days past SLA	BSA Officer	Individually, by alert
Alert open > 45 days past SLA	Chief Compliance Officer + noted in Board MIS	Individually, by alert
Any alert approaching day 25 of the 30-day SAR clock	BSA Officer — **hard escalation**	Automatic; cannot be dismissed by the analyst
RFI unanswered > 2 business days	LoB Compliance Head	Automatic
RFI unanswered > 5 business days	Divisional Head + logged as a business-line control breach	Automatic
Employee-nexus alert	Head of SIU	Immediate; bypasses all standard queues
Sanctions true match	BSA Officer + Head of Sanctions	Within 1 hour
Potential SAR identified	SAR Review Committee	Within 3 business days of identification
Decision not to file where an SAR was recommended by the investigator	BSA Officer, in writing, with rationale; reported to the Steering Committee	Individually

The most important line in that table
A documented decision *not* to file, where the investigating analyst recommended filing, is the highest-risk decision the FIU makes. It is the exact fact pattern that appears in consent orders. MNB requires it in writing, from the BSA Officer personally, with reasoning — and every such decision is listed by name in the quarterly report to the Board BSA & AML Committee.

9.7 SAR Decisioning and Filing

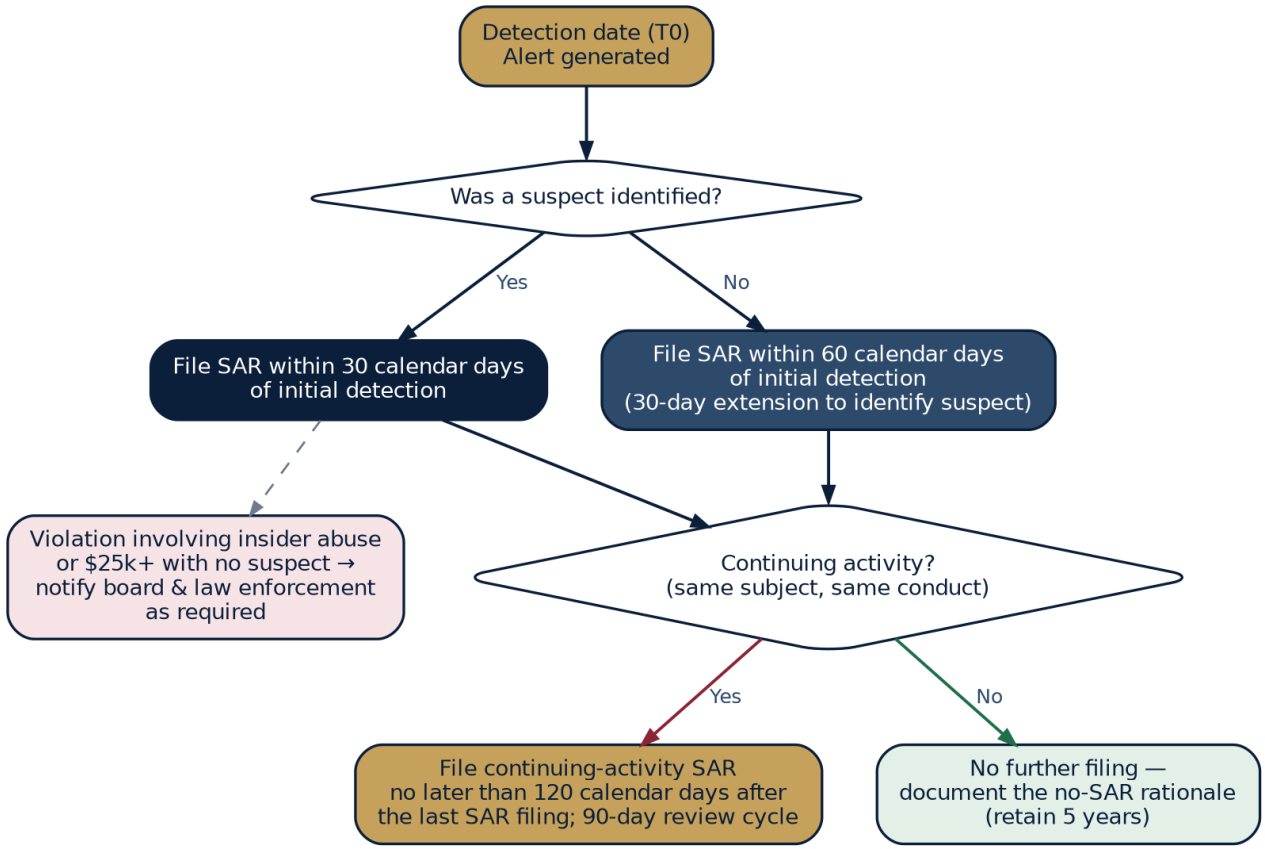

Figure 10 — SAR filing timeline: the 30/60-day rules from initial detection and the 120-day continuing-activity cycle.

Requirement	Standard at MNB
Filing trigger	Known or suspected violation aggregating $5,000+ with an identified suspect; $25,000+ with no suspect; any insider abuse regardless of amount; any transaction with no apparent lawful purpose.
The 30-day clock	Runs from **initial detection**. MNB conservatively sets initial detection = alert generation date unless the FIU documents, contemporaneously, why a later date is the correct one. Retro-dating detection to buy time is prohibited and is a disciplinary matter.
The 60-day extension	Available only where no suspect has been identified. It is not a general extension and may not be used to accommodate internal backlog.
Continuing activity	Reviewed on a 90-day cycle; a continuing-activity SAR is filed no later than 120 calendar days after the date of the previous filing.
Narrative standard	Who, what, when, where, why and how — in plain English, chronological, with amounts, dates, account numbers and counterparties. The narrative must be self-contained: a FinCEN analyst reading it, with no access to MNB's systems, must be able to understand the activity. (Template: Annex D.)
What must not appear	Alert scores, model outputs, internal scenario names, or speculation. The narrative states facts and the reasoned basis for suspicion.
Filing mechanism	FinCEN BSA E-Filing (Form 111). The Document Control Number is captured back into ActOne and reconciled — an unacknowledged filing is treated as a non-filing until the DCN is received.
Confidentiality	The existence of a SAR is disclosed to no one outside the permitted circle. 31 U.S.C. 5318(g)(2) prohibits disclosure to the subject or to any person named in it, including — importantly — the relationship manager, who is *not* told that a SAR was filed.
Post-filing	Mandatory: exit/de-risking referral where warranted; account restriction consideration; CRR forced to Very High; enhanced monitoring flag applied; and a documented review of whether prior NFA closures on the party should be reopened.
Retention	SAR and all supporting documentation retained 5 years from the filing date (31 CFR 1020.320(d)).

9.8 Information Sharing — 314(a) and 314(b)

Provision	Obligation / permission	MNB process
314(a)	**Mandatory.** Search records for named subjects on FinCEN's request; report matches within 14 days.	Automated search across PARTY, ACCOUNT and TRXN covering the statutory period. Positive matches: report to FinCEN, force CRR to Very High, auto-create a case, and retro-review 12 months of activity. A 314(a) match is **not** automatically a SAR — but it is always a documented review.
314(b)	**Voluntary**, with a safe harbour from liability, permitting information sharing between registered financial institutions on suspected ML/TF.	MNB is registered and re-registers annually. Requests are made and answered by the FIU only, under a controlled log. 314(b) is materially under-used across the industry, and MNB actively uses it — particularly on mule networks and funnel accounts, where the counterparty bank routinely holds the missing half of the picture.

10. Alert Optimisation — Suppression, Trusted Pairs and Auto-Close

These three mechanisms remove roughly 88% of raw alert volume at MNB before an analyst sees anything. That figure is both the reason the FIU is operable and the reason this chapter carries the heaviest governance in the framework. Every one of these mechanisms is, in substance, a decision not to look.

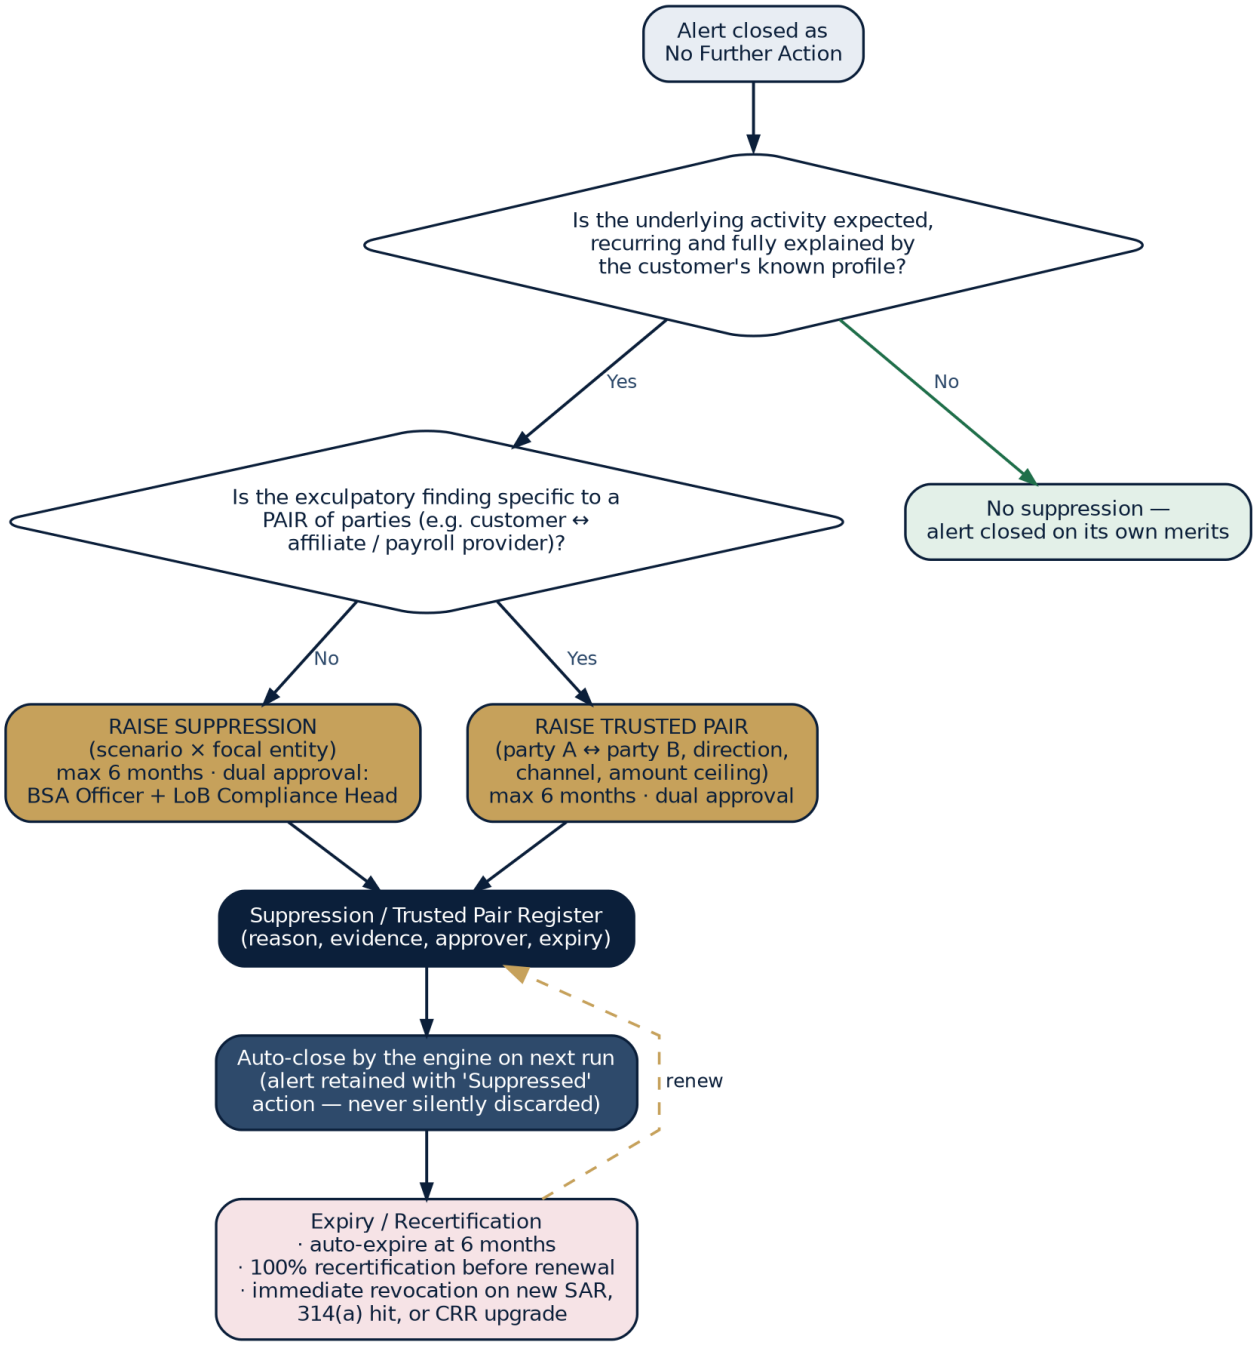

Figure 11 — Decision tree for raising a suppression rule versus a trusted pair, with approval, register and expiry controls.

10.1 The Three Mechanisms Compared

	Suppression	Trusted Pair	Auto-Close Rule
Scope	Scenario × focal entity	Party A ↔ Party B (directional)	Rule-based across the population
Basis	The behaviour is understood and expected for this specific entity under this specific scenario	The relationship between these two parties has been investigated and is legitimate	Objectively low-risk alert characteristics
Alert still generated?	Yes — auto-closed with action 'Suppressed'	Yes — auto-closed	Yes — closed with a reason code
Maximum duration	6 months	6 months	Reviewed annually; effective indefinitely until changed
Approval	BSA Officer + LoB Compliance Head (dual)	BSA Officer + LoB Compliance Head (dual)	AML Model Committee + BSA Officer
Who may request	L2 investigator and above, with evidence	L2 investigator and above, with evidence	AML Model Governance only
Auto-expiry	Hard expiry; no auto-renewal	Hard expiry; no auto-renewal	n/a
QA sampling	100% at recertification; 10% in-life monthly	100% at recertification; 10% in-life monthly	5% stratified monthly sample worked to full L1 standard
Immediate revocation	On new SAR, 314(a) match, watch-list hit, CRR upgrade, or any employee nexus	Same	Same, at rule level

10.2 Illustrative Auto-Close Rules

Auto-close rules and the prohibited conditions

```
-- Auto-close rules are conjunctive by design. Every condition must hold.
-- A rule that closes on a single condition is not a rule, it is a hole.

RULE AC-01  'Low-value low-risk single-match alert'
  CLOSE alert WHERE
        alert.priority              = 'LOW'
    AND alert.match_count           = 1
    AND focal.CRR                   = 'LO'
    AND focal.pep_flag              = 'N'
    AND focal.watchlist_nexus       = FALSE
    AND focal.prior_sar             = FALSE
    AND focal.prior_alerts_12m      = 0
    AND alert.total_amt             < 1.2 * scenario.threshold
    AND alert.counterparty_max_risk < 4
    AND focal.account_age_days      > 365
  REASON_CODE = 'AC-01_LOW_RISK_MARGINAL_BREACH'

RULE AC-02  'Recurring verified salary credit'
  CLOSE alert WHERE
        alert.scenario_id           IN ('SAM-BD-01','SAM-BD-02')
    AND credit_source               = 'ACH_PAYROLL'
    AND ach_company_id              IN VERIFIED_EMPLOYER_REGISTRY
    AND amount_within_tolerance(historical_salary, 15%)
    AND no_other_scenario_alert_in_window(focal, 30d)
  REASON_CODE = 'AC-02_VERIFIED_PAYROLL'

-- PROHIBITED auto-close conditions (hard-blocked in Policy Manager):
--   * any condition referencing FIU capacity, backlog or queue depth
--   * any condition on a focal entity with a prior SAR
--   * any condition on a focal entity with a watch-list or 314(a) nexus
--   * any condition on an employee-nexus alert
--   * any condition on a Very High or High CRR entity
--   * any condition on scenarios SAM-HT-01, SAM-EE-01, SAM-WL-01, SAM-CB-01
```

10.3 Why the Alert Is Still Generated

The record that saves the bank

A suppressed alert is **generated, scored, persisted, actioned with a reason code, and retained for five years**. It is never prevented from existing. The distinction is not academic: if the subject of a suppressed alert is later indicted, the bank must be able to show that its system *detected* the behaviour, that a named person *considered* it, and that the suppression was *approved on stated evidence at the time*. A bank that instead configured the scenario not to fire has no such record — it has only silence, and silence is indistinguishable from a control failure.

11. Sanctions Interface — WLF and the AML Referral

Sanctions screening is governed by its own standard; what belongs here is the interface — how a WLF outcome becomes an AML alert, and why the two disciplines must not be conflated. Sanctions is a *strict-liability, real-time, transaction-blocking* control. AML is a *risk-based, retrospective, behaviour-reporting* control. The same customer can be perfectly clear of sanctions and thoroughly suspicious, and vice versa.

Figure 12 — Real-time watch-list filtering, interdiction, disposition, and the cross-referral into AML monitoring.

Event in WLF	Sanctions action	AML referral action
Confirmed true match — SDN	Block the payment; report to OFAC within 10 business days; annual blocked-property report	Force CRR to Very High; auto-create case; 12-month retro-review of all activity; SAR consideration mandatory
Confirmed true match — sectoral / non-SDN	Reject or restrict per the applicable programme	Same as above, with the programme-specific scope
False positive released	Release; consider a Good Guy list entry to reduce recurrence	No AML referral — but repeated near-matches on the same party are themselves a monitored KRI
PEP identified	No blocking (PEP status is not a sanction)	CRR uplift; EDD; source-of-wealth verification; ongoing enhanced monitoring
Adverse media — financial crime	None	CRR uplift; case created if the allegation concerns ML/TF/fraud
314(a) subject match	None (314(a) is not a sanction)	Mandatory report to FinCEN within 14 days; case created; retro-review

The error to avoid

“The customer cleared sanctions screening” is not, and never has been, an answer to an AML alert. It appears in closure narratives more often than any other single irrelevance. Screening tells the bank the customer is not a designated person. It says nothing whatever about whether their cash deposits make sense.

12. Model Governance, Tuning and Validation

Under SR 11-7 and OCC 2011-12, a detection scenario is a model: it uses statistical and business assumptions to produce a quantitative output used in decision-making. So is the segmentation logic, the CRR, the scoring rules and the AI prioritisation layer. All of them sit in the model inventory, and all of them are validated.

12.1 Threshold Tuning Lifecycle

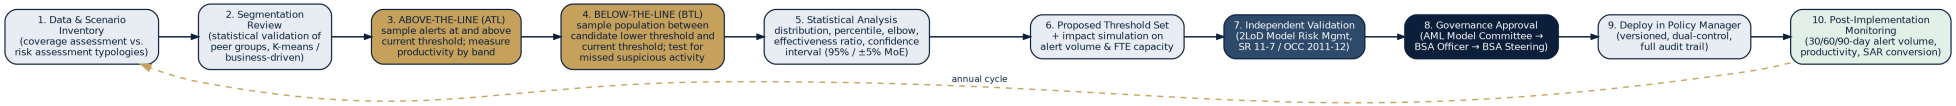

Figure 13 — The annual threshold tuning lifecycle: coverage, segmentation, ATL/BTL testing, statistical analysis, validation, approval, deployment and post-implementation monitoring.

12.2 Above-the-Line and Below-the-Line Testing

ATL and BTL are the two halves of a single question: *is the threshold in the right place?* ATL asks whether the alerts the bank is generating are worth generating. BTL asks — and this is the limb regulators scrutinise — **whether the bank is missing suspicious activity just below the line**.

	Above-the-Line (ATL)	Below-the-Line (BTL)
Question	Are alerts at and above the current threshold productive?	Is there suspicious activity between a lower candidate threshold and the current one?
Population	Alerts generated in the last 12 months, stratified by score band and segment	The *non-alerting* population between the candidate lower threshold and the current threshold — a population that, by definition, no one has ever looked at
Method	Sample by band; measure the escalation rate and the SAR conversion rate per band	Sample; simulate the alert; investigate each to full L1/L2 standard as if it had alerted
Success criterion	A band whose productivity is materially lower than adjacent bands is a candidate for threshold increase — *only if* BTL confirms nothing is missed at that level	**Zero** would-be SARs in the BTL sample. Any escalation in the BTL sample is decisive evidence that the threshold is too high and it must be lowered.
Sample size	Statistically determined; 95% confidence, ±5% margin of error	Same; typically 350–400 items per scenario per segment
Who performs it	AML Model Governance, with investigations resourced from the FIU	Same — and investigated by analysts who are **blind to the fact that the item is a BTL sample**, to prevent a foregone conclusion

BTL sample sizing and decision rule

```
-- Sample size for a proportion, at 95% confidence with a 5% margin of error.
-- p is set to 0.5 (maximum variance) where the true escalation rate is unknown.

     n0  =  z^2 * p * (1 - p) / e^2
         =  (1.96)^2 * 0.5 * 0.5 / (0.05)^2
         =  384.16   ->  385

-- Finite population correction, where N is the size of the BTL population:

     n   =  n0 / ( 1 + (n0 - 1) / N )

-- e.g. N = 4,200 items in the BTL band for SAM-CS-01 / segment RET-IND-ME:
--      n = 385 / (1 + 384/4200) = 353 items to be investigated.

-- DECISION RULE:
--   0 escalations in 353  ->  95% confident the true escalation rate is < 0.85%.
--                             Threshold increase MAY proceed to committee.
--   >= 1 escalation       ->  the threshold is DEMONSTRABLY too high.
--                             Lower it. There is no discretion in this step, and no
--                             appeal to alert volume, cost or capacity is admissible.
```

12.3 Scenario Coverage Assessment

Performed annually against the EWRA. Every inherent-risk typology the bank has identified must map to at least one production scenario, or be recorded as an accepted, documented gap with a named owner and a remediation date. This mapping — not the alert volume, not the SAR count — is the first thing a competent examiner asks to see, because it tests whether the monitoring system was designed against the bank's risk or merely inherited from a vendor.

EWRA typology (extract)	Covering scenario(s)	Status
Cash placement / structuring	SAM-CS-01, SAM-CS-02, SAM-CS-03	Covered

EWRA typology (extract)	Covering scenario(s)	Status
Funnel accounts / money mules	SAM-RM-01, SAM-RM-02, SAM-RM-03	Covered
Correspondent nesting	SAM-CB-01	Covered
Trade-based money laundering	SAM-TB-01	Covered — *partial*: structured HS-code and unit-price capture is incomplete on open-account trade. Remediation Q4 2026, owner: Head of Trade Technology.
Human trafficking	SAM-HT-01	Covered
Elder financial exploitation	SAM-EE-01	Covered
Virtual asset exposure / ransomware	SAM-VA-01	Covered
Account takeover / sold accounts	SAM-BD-03, IFM-X referral	Covered
Insider facilitation	SAM-IN-01	Covered
Real-estate laundering (all-cash purchases)	—	**GAP.** No dedicated scenario. Partially covered by SAM-RM-01 and SAM-CS-02. Accepted by the Steering Committee 2026-05-14; scenario in development, target Q1 2027.

The real-estate line is deliberately included. **A coverage assessment with no gaps in it is not credible** — it indicates the assessment was performed to produce a clean result rather than an honest one. Examiners know this. An honestly-declared, owned, dated gap is a sign of a functioning programme; a flawless matrix is a sign of one that has stopped looking.

12.4 Independent Validation

Validation element	Scope
Conceptual soundness	Does the scenario's logic actually detect the typology it claims to? Are the assumptions documented and defensible? Would a determined launderer trivially evade it?
Data adequacy	Are the fields the scenario depends on complete, accurate and timely? A scenario keyed on a field that is 60% null is not a control.
Outcomes analysis	Back-testing against known SARs and known enforcement typologies. Would this scenario have caught the last five cases the bank actually filed?
Benchmarking	Alert productivity and SAR conversion versus peer institutions and prior periods, per scenario per segment.
Threshold appropriateness	Independent replication of the ATL/BTL analysis on an independent sample.
Segmentation validity	Statistical re-testing of cluster stability and within-segment homogeneity.
Implementation verification	Does the code in production actually implement the approved logic? Independent replication of a sample of alerts from raw data — the step that catches the defect where the approved specification says one thing and the deployed model does another.
Ongoing monitoring	Are performance metrics defined, tracked, and acted upon? Are breaches escalated?

Frequency: full validation on implementation and every 3 years; annual review in the intervening years; immediate re-validation on any material change to logic, data, thresholds or segmentation.

12.5 Key Performance and Risk Indicators

Metric	Definition	Target / trigger	Reported to
Alert productivity	Escalated alerts ÷ alerts worked, **per scenario per segment**	< 2% or > 60% → tuning review	AML Model Committee, monthly
SAR conversion rate	SARs filed ÷ alerts worked, per scenario	Trend monitored; a sustained zero is a coverage question, not a success	Steering Committee, monthly

Metric	Definition	Target / trigger	Reported to
BTL escalation rate	Escalations found in the below-the-line sample	**Zero.** Any hit is a mandatory threshold reduction	AML Model Committee + MRM
Auto-close QA hit rate	Escalations found in the 5% sample of auto-closed alerts	< 0.5%; any breach suspends the offending rule	Compliance Testing, monthly
SAR timeliness	SARs filed within 30 days of initial detection	100%. Any miss is individually reported	Board BSA Committee, quarterly
Alert ageing	Alerts open beyond disposition SLA	< 5% of open inventory	Head of FIU, daily
Data completeness	DQ-06 mandatory field population	≥ 99.9%	BSA Officer, daily
Feed timeliness	Feeds arriving within the SLA window	≥ 99.5%	BSA Officer, daily
Suppression inventory	Active suppressions and trusted pairs; count and age	Trend; any growth > 10% QoQ is investigated	Steering Committee, quarterly
Model drift	AI prioritisation AUC and feature-distribution drift	AUC decline > 5% → retrain	MRM, monthly
Decision-quality score	QA sample pass rate on closure narratives	≥ 95%	Head of FIU, monthly
Reopened alerts	Alerts reopened after NFA closure	Trend; a spike indicates a triage-quality problem	Head of FIU, monthly

12.6 Quality Assurance Sampling

Population	Sample	Frequency	Performed by
Alerts closed NFA at L1	5% stratified random	Monthly	Compliance Testing (2LoD)
Auto-closed alerts	5% stratified random, worked to full L1 standard	Monthly	Compliance Testing (2LoD)
Suppressed alerts	10% in-life; 100% at recertification	Monthly / at expiry	Compliance Testing (2LoD)
Cases closed without a SAR	10%	Monthly	Compliance Testing (2LoD)
Cases closed without a SAR *where the investigator recommended filing*	**100%**	Continuous	BSA Officer + Compliance Testing
SARs filed	10% for narrative quality and timeliness	Monthly	FIU QA (1LoD) + Compliance Testing (2LoD)
CRR downward overrides	**100%**	Continuous	Compliance Testing (2LoD)
Threshold changes	**100%**	Per change	MRM + Internal Audit review

13. Governance, Roles and Access

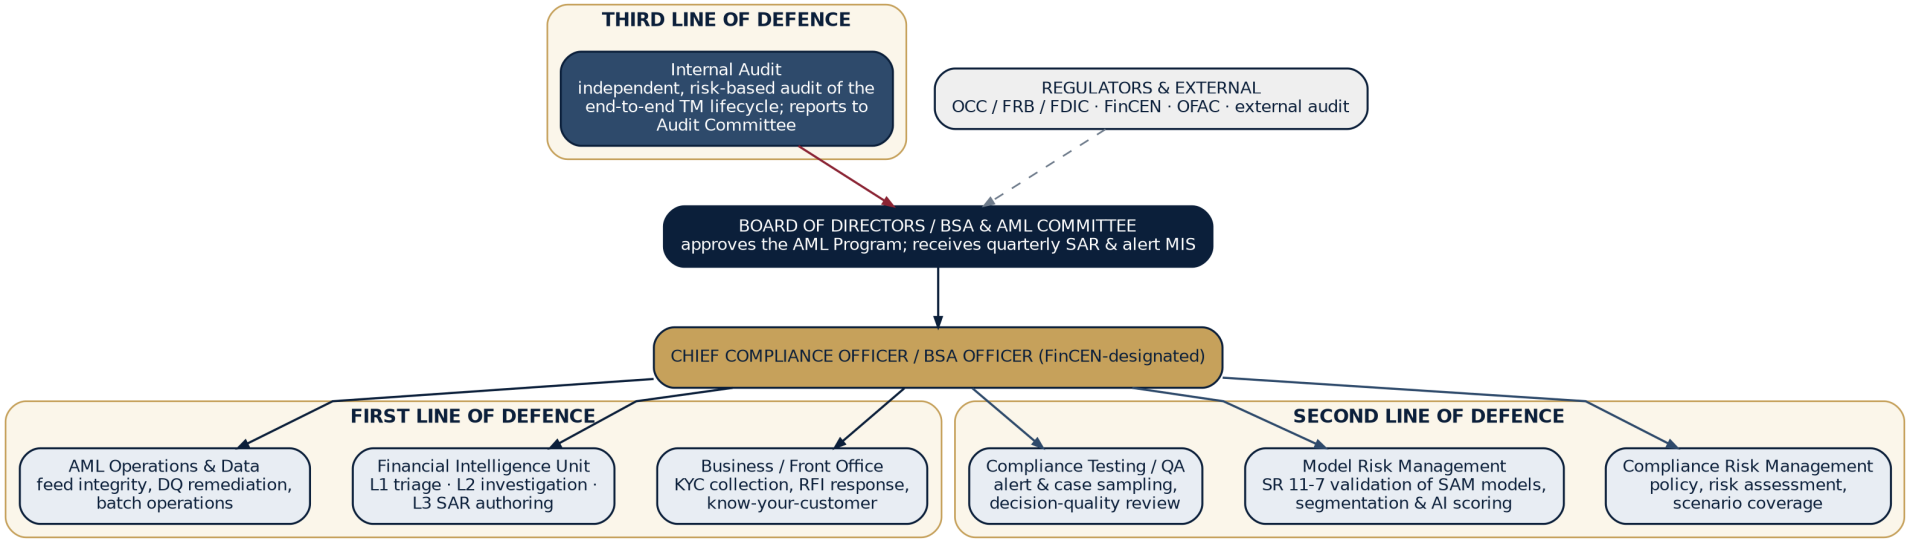

Figure 14 — Three lines of defence and the governance structure for the transaction monitoring programme.

13.1 RACI — Alert and Case Lifecycle

Activity	L1 Analyst	L2 Inv.	SIU	FIU Lead	BSA Officer	MRM	Audit
Alert triage and closure	R	C	—	A	I	—	—
Case creation and investigation	C	R	C	A	I	—	—
Employee-nexus investigation	—	—	R	I	A	—	I
RFI issuance	R	R	R	A	I	—	—
SAR recommendation	C	R	R	A	I	—	—
SAR decision and filing	—	C	C	C	**A/R**	—	—
Decision *not* to file (against recommendation)	—	C	C	C	**A/R**	—	I
Suppression / trusted pair request	—	R	R	C	**A**	—	—
Threshold change proposal	—	C	—	C	A	R	I
Threshold change validation	—	—	—	—	C	**R**	I
Scenario coverage assessment	—	C	C	C	A	R	I
Auto-close rule design	—	C	—	C	A	R	I
AI model approval	—	—	—	—	A	R	I
QA sampling	—	—	—	I	A	—	I
Independent audit of the lifecycle	—	—	—	I	I	I	**R**

R = Responsible · A = Accountable · C = Consulted · I = Informed

13.2 System Access and Segregation of Duties

Privilege	L1	L2	FIU Lead	BSA Officer	AML Model Gov.	Actimize Admin
View assigned alerts	X	X	X	X	—	—
View all alerts (portfolio)	—	—	X	X	X	—
Add narrative / attachments / comments	X	X	X	X	—	—
Close alert (NFA)	X	X	X	X	—	—
Reassign alert	—	—	X	X	—	—
Create / manage case	—	X	X	X	—	—
Recommend SAR	—	X	X	X	—	—
Approve / file SAR	—	—	—	**X**	—	—
Request suppression / trusted pair	—	X	X	—	—	—
Approve suppression / trusted pair	—	—	—	**X**	—	—
View Policy Manager (read-only)	—	—	X	X	X	X
Edit thresholds / segments / auto-close rules	—	—	—	—	**X**	—
Approve a threshold change in production	—	—	—	**X**	—	—
Manage watch lists	—	—	—	X	—	X
Manage users and roles	—	—	—	—	—	X
Execute / restart a detection run	—	—	—	—	—	X

Privilege	L1	L2	FIU Lead	BSA Officer	AML Model Gov.	Actimize Admin
View SIU-restricted queue	—	—	—	X	—	—
View the audit trail	—	—	X	X	X	X
Delete or amend an audit record	**—**	**—**	**—**	**—**	**—**	**—**

Three hard separations
(1) The person who *edits* a threshold in Policy Manager may not be the person who *approves* it. **(2)** The person who *investigates* a case may not be the person who *approves the SAR decision* — in either direction. **(3)** Nobody, at any level, including the Actimize administrator and the BSA Officer, can delete or amend an audit record. The last row of that table is empty by design, and it is the row an auditor checks first.

13.3 Committee Structure

Forum	Frequency	Mandate
Board BSA & AML Committee	Quarterly	Approves the AML programme; receives SAR volumes, timeliness, coverage gaps, audit findings, and every decision not to file against an investigator's recommendation.
BSA/AML Steering Committee	Monthly	Chaired by the CCO. Owns the framework, the risk assessment, resourcing, and the alert-inventory position.
AML Model Committee	Monthly	Approves scenarios, thresholds, segmentation, auto-close rules and AI models on the strength of MRM's opinion. Cannot approve a change over an unresolved MRM objection — such changes escalate to the Steering Committee with the objection on the record.
FIU Operations Forum	Weekly	Queue health, ageing, capacity, escalations, quality themes.
SAR Review Committee	Twice weekly	Chaired by the BSA Officer. Decides filings.

13.4 Training

Audience	Content	Frequency
All staff	BSA/AML awareness; red flags; how and to whom to escalate; the absolute prohibition on tipping off	Annual
Branch and front-office staff	Cash red flags; conductor capture; the CTR conversation; recognising a customer who is being coached or controlled	Annual + on hire
L1 analysts	Scenario intent (Ch. 6); evidence standards; closure narrative quality; system use	On hire + semi-annual + on any scenario change
L2 / SIU investigators	EDD; public-records and open-source research; flow-of-funds analysis; SAR narrative authoring; 314(b); typology deep-dives	On hire + quarterly
Relationship managers	Why the RFI matters; the two-day SLA; what a useful answer looks like; no tipping off	Annual
Board and senior management	Programme performance; enforcement landscape; personal liability under the AMLA 2020	Annual

Annexures

Annex A — Scenario, Frequency, Lookback and Threshold Matrix

Thresholds by risk segment. High-risk segments carry stricter (lower) thresholds; low-risk segments carry looser ones. This table is the single most sensitive artefact in the framework and its distribution is restricted.

ID	Scenario	Focus	Freq.	Lookback	Key threshold — Low	Medium	High	Very High
SAM-CS-01	Structuring — threshold avoidance	CU	Daily	7d	4 txns / $25k agg	3 txns / $18k	3 txns / $12k	2 txns / $9k
SAM-CS-02	Large cash activity	CU	Daily	1d	$60k agg	$40k	$25k	$12k
SAM-CS-03	Cash inconsistent with profile	CU	Weekly	90d	ratio 0.75 / z 3.0	0.60 / z 2.5	0.45 / z 2.0	0.30 / z 1.5
SAM-RT-01	Possible CTR (completeness control)	CU	Daily	2d	$10,000 (statutory — not segmented)	$10,000	$10,000	$10,000
SAM-RM-01	Rapid movement / pass-through	AC	Daily	10d	$50k / 90%	$25k / 80%	$15k / 70%	$10k / 60%
SAM-RM-02	Deposits followed by rapid wires out	CU	Daily	10d	$30k / 80%	$15k / 70%	$10k / 60%	$5k / 50%
SAM-RM-03	Transfers between unrelated parties	AC	Daily	30d	$40k / 5 cpty	$20k / 3	$12k / 3	$8k / 2
SAM-BD-01	Significant change from average	AC	Month-start	12m	500% / z 4.0	300% / z 3.0	200% / z 2.5	150% / z 2.0
SAM-BD-02	Activity exceeds anticipated (CDD)	CU	Month-end	1m	300% / $40k	200% / $20k	150% / $10k	100% / $5k
SAM-BD-03	Dormant account escalation	AC	Daily	1d / 180d	$25k in / $12.5k out	$10k / $5k	$7.5k / $3.75k	$5k / $2.5k
SAM-HR-01	Focal high-risk entity	AC	Daily	1d	n/a (entity risk ≥ 7 only)	$10k	$5k	any
SAM-HR-02	High-risk counterparty / corridor	AC	Daily	14d	$40k / 3 txn	$15k / 2	$10k / 2	$5k / 1
SAM-CB-01	Correspondent nesting / volume	CB	Weekly	30d	400% dev	250% dev	150% dev	100% dev
SAM-IN-01	Internal / employee-nexus patterns	AC	Weekly	7d	Not segmented — any match routes to SIU	—	—	—
SAM-IN-02	Round-dollar / repeating amounts	CU	Weekly	30d	10 txn / $100k	6 / $50k	5 / $30k	4 / $20k

ID	Scenario	Focus	Freq.	Lookback	Key threshold — Low	Medium	High	Very High
SAM-HT-01	Human trafficking indicators	CU	Weekly	90d	4 indicators	3	3	2
SAM-EE-01	Elder financial exploitation	CU	Weekly	60d	3 indicators	2	2	2
SAM-VA-01	Virtual asset exposure	CU	Weekly	30d	$25k	$10k	$5k	any
SAM-TB-01	Trade-based ML indicators	CU	Monthly	180d	3 flags	2	2	1
SAM-WL-01	Watch-list / 314(a) nexus	CU	Real time	n/a	Any true match (not segmented)	—	—	—

Four production scenarios (SAM-CS-04 seasonal-business variance, SAM-RM-04 ATM cash clustering, SAM-HR-03 MSB customer activity, SAM-IN-03 third-party payment processor concentration) are omitted from this extract for brevity; they follow the same specification format.

Annex B — Alert Disposition and Closure Code Set

Code	Disposition	Requires	Counts as
NFA-01	Activity consistent with the known customer profile	Explicit profile comparison in the narrative	Closed — non-productive
NFA-02	Activity explained by verified documentation (invoices, contracts, sale proceeds)	Documents attached	Closed — non-productive
NFA-03	Activity explained by a verified RFI response	RFI attached; RM's explanation tested against the transaction record	Closed — non-productive
NFA-04	One-off event with an established legitimate purpose (property sale, inheritance, settlement)	Corroborating evidence	Closed — non-productive
NFA-05	Duplicate of an alert already investigated	Reference to the primary alert	Closed — duplicate
NFA-06	Data quality defect — the alert is an artefact of bad data	DQ incident raised and referenced	Closed — DQ (tracked separately; a rising trend is a control failure)
SUP-01	Auto-closed under an approved suppression rule	Rule reference and expiry date	Suppressed
TP-01	Auto-closed under an approved trusted pair	Pair reference and expiry date	Suppressed
AC-xx	Auto-closed under an approved auto-close rule	Rule reference	Auto-closed (5% QA sampled)
ESC-01	Escalated to L2 — further investigation required	Escalation rationale	Escalated
ESC-02	Escalated to SIU — employee nexus or sensitive	Escalation rationale	Escalated (restricted)
CASE-01	Promoted to case	Case ID	Escalated
SAR-01	SAR filed	DCN captured; post-filing actions completed	**Productive**
SAR-02	Continuing-activity SAR filed	Prior DCN referenced	**Productive**
NOSAR-01	Case closed — no SAR, investigator agreed	Four-eyes approval	Closed
NOSAR-02	Case closed — no SAR **against the investigator's recommendation**	**BSA Officer written rationale; reported to the Board**	Closed — **100% QA**
REOPEN-01	Reopened — new information	Trigger documented	Reopened
REOPEN-02	Reopened — QA rejection	QA finding referenced	Reopened — quality failure

Annex C — Suppression / Trusted Pair Request Form (fields)

Field	Requirement
Request type	Suppression (scenario × entity) or Trusted Pair (party ↔ party)
Scenario / pair	Scenario ID, or Party A and Party B with direction, channel and amount ceiling
Focal entity	Party ID, name, CRR, LoB, segment
Alerts relied upon	Alert IDs previously investigated and closed, with dispositions
Investigation summary	What the activity is, why it recurs, why it is not suspicious
Evidence attached	Documents, RFI responses, CDD records relied upon

Field	Requirement
Screening confirmation	Confirmation that no watch-list, PEP, 314(a) or prior-SAR nexus exists as at the request date
Employee-nexus attestation	Explicit confirmation that no party in scope has an employee nexus
Requested duration	Maximum 6 months; a shorter period must be justified upward, not downward
Expected volume impact	Alerts expected to be suppressed per month (from the UAT simulation)
Requester	L2+ investigator, with date
Approver 1	LoB Compliance Head
Approver 2	BSA Officer
Auto-expiry date	System-enforced; no auto-renewal is possible
Recertification evidence	Required in full before any renewal; a renewal is a new request, not an extension

Annex D — SAR Narrative Template

SAR narrative structure

```
1. INTRODUCTION
   This report is filed by Meridian National Bank, N.A. (RSSD ########) on
   [SUBJECT NAME], [DOB / EIN], [ADDRESS], a customer since [DATE], to report
   [TYPOLOGY] totalling $[AMOUNT] between [START DATE] and [END DATE].

2. SUBJECT AND ACCOUNT INFORMATION
   Subject:      name, identifiers, occupation / stated business, relationship since.
   Accounts:     numbers, types, opening dates, signatories, beneficial owners.
   Related parties: counterparties, conductors, associated entities, and the nature
                 of each relationship.

3. SUSPICIOUS ACTIVITY  (WHO / WHAT / WHEN / WHERE / WHY / HOW)
   Chronological. Specific. Quantified. State amounts, dates, branches, channels,
   counterparties and jurisdictions. Describe the PATTERN, not merely the individual
   transactions:

     'Between 06 and 09 July 2026, three cash deposits totalling $28,450 were made
      into account 0019-448220 at three separate branches (BR-0182, BR-0407,
      BR-0311). Each deposit fell between $9,200 and $9,750 - immediately below the
      $10,000 currency transaction reporting threshold. Two of the three deposits
      were made by a third-party conductor, [NAME], who is not a signatory on the
      account and whose relationship to the subject the bank has been unable to
      establish. On 10 July 2026, $27,900 was wired from the same account to
      [BENEFICIARY] at [BANK], [COUNTRY] - a beneficiary with no prior transaction
      history with the subject and no apparent nexus to the subject's stated
      business as a licensed general contractor...'

4. WHY THE ACTIVITY IS SUSPICIOUS
   State the reasoned basis. Reference the customer's expected activity, the
   inconsistency, and what the bank did to try to resolve it (RFI, EDD, 314(b)).
   Facts and reasoning only - no speculation about predicate offences the bank
   cannot evidence.

5. WHAT THE BANK DID
   Investigation steps, RFI issued and response received (or not), EDD performed,
   314(b) requests made, prior SARs on the subject, and any account action taken
   (restriction, exit) - or an explicit statement that the relationship continues
   and monitoring is enhanced.

6. FINCEN KEY TERM  (where an advisory applies)
   e.g.  'HUMAN TRAFFICKING FIN-2020-A008'
         'ELDER FINANCIAL EXPLOITATION FIN-2022-A002'

7. CONTACT
   BSA Officer, name, title, telephone, and confirmation that supporting
   documentation is retained and available to law enforcement on request.

NEVER INCLUDE:  alert scores  |  model outputs  |  internal scenario IDs
                speculation   |  the words 'possible' or 'may be' used to hedge a
                                 suspicion the bank in fact holds
```

Annex E — Request for Information (RFI) to the Business

#	Question	Response
1	State the exact nature of the customer's business or profession and the source of the funds. Generic descriptions ('self-employed', 'trading', 'consultancy', 'private service') are not acceptable and will be returned.	Free text
2	State the specific purpose of the transactions listed in the attached schedule.	Free text
3	Is the customer, a close associate, or a family member a PEP? If yes, confirm that the required senior approval is held.	Y/N + evidence
4	Are the counterparties in the attached schedule consistent with the customer's KYC profile, and do the transactions make commercial sense?	Y/N + explanation

#	Question	Response
5	Is the expected turnover on file consistent with the customer's financial standing and nature of business? If not, has the CDD record been updated?	Y/N + date of update
6	Are you satisfied with the conduct of the account in the light of the customer's stated business needs?	Y/N + explanation
7	Legal entity: has the natural person acting on behalf of the customer been identified and verified, and is the beneficial ownership record current? Natural person: is the account operated by a third party, mandate holder or nominee? If yes, has EDD been performed on that party?	Y/N + evidence
8	Is the customer associated with an NGO, NPO, trust or charitable association? If yes, what EDD has been performed?	Y/N + evidence
9	Has the customer been the subject of, or party to, any law-enforcement enquiry, subpoena or investigation known to the business?	Y/N + detail
10	Provide any other information relevant to the customer's due diligence, including anything that gives rise to concern.	Free text

Standing warning on every RFI

This request is confidential. It must not be discussed with the customer or with any person named in it. Contacting the customer about this request, or altering your handling of the relationship in a way the customer could notice, may constitute tipping off — a federal offence under 31 U.S.C. 5318(g)(2) carrying personal criminal liability. If you believe you need to contact the customer, you must speak to the FIU first.

Annex F — Batch Run-Book: Operational Checklist

Time	Check	Pass condition	On failure
00:20	All 13 feeds present and manifest-matched	13/13	Page feed owner; escalate at +30 min; batch halts at 01:00 if unresolved
00:45	DQ Critical rules DQ-01…05, 12, 14 all pass	Zero breaches	**Halt.** Notify BSA Officer. Do not proceed.
00:45	DQ High rules DQ-06…11	Within tolerance	Proceed; raise a compliance incident; report same day
01:30	ADM load record counts reconcile to feed manifests	±0	Rollback to checkpoint; re-run
02:30	Entity resolution: unresolved queue size	< 500	Proceed; escalate to the data steward
03:00	Watch-list version current	≤ 24 h since OFAC publication	**Halt.** DQ-14 is Critical.
03:15	Profile build complete; row counts within ±3σ of trailing mean	Pass	Re-run; if it fails twice, halt
04:00	Policy Manager parameter bind: all 24 scenarios have an active, approved threshold set	24/24	**Halt.** A scenario running on an unbound or expired threshold set is prohibited.
06:00	All scenario plans completed	5/5 plans	Restart failed plans from checkpoint; if the 07:30 sign-off cannot be met, notify the BSA Officer and run a catch-up batch the same day
06:45	Alert volume within ±3σ of the trailing 30-day mean for that weekday	Pass	**Investigate before release.** A sudden volume collapse is the classic signature of a silently broken feed or a mis-bound threshold — and it is far more dangerous than a volume spike.
07:15	All alerts routed; zero alerts in the unrouted state	0 unrouted	Manual release under dual control
07:30	Batch sign-off by the on-call Actimize administrator and the FIU duty lead	Both	Sign-off withheld; incident raised; reported in the daily BSA Officer pack

The check that matters most
The 06:45 volume check is the single most valuable control in the run-book. Every large-bank monitoring failure that has ended in an enforcement action shares one feature: **the alerts simply stopped, and nobody noticed.** A feed silently dropped a product. A threshold was bound to the wrong segment. A scenario was disabled during an upgrade and never re-enabled. In every case the batch completed successfully, the dashboards were green, and the volume fell — and no one was watching the volume. Watch the volume.

Annex G — Glossary of Detection Field Derivations

Derived attribute	Derivation
NEW_ACCT_FLAG	ACCOUNT.OPEN_DT ≥ (run_dt – 180). Tightens thresholds on unseasoned relationships.
DORMANCY_FLAG	No customer-initiated transaction for ≥ 180 days. Bank-initiated postings (interest, fees) do not reset dormancy — a common and consequential mapping error.
CASH_RATIO	Cash credits ÷ total credits over the window.
FLOW_THROUGH_PCT	Debits ÷ credits over the window.
HOLD_PERIOD	Median days between a credit and the debit that consumes it (FIFO attribution).
COUNTERPARTY_DIVERSITY	Distinct counterparties ÷ transaction count. Low diversity with high volume indicates an exclusive relationship; high diversity with small amounts indicates aggregation.
Z_SCORE	(observed – peer mean) ÷ peer standard deviation, within the focal entity's segment and window.
CORRIDOR_RISK	MAX(country risk) across the originating country, the beneficiary country *and every intermediary in the chain*.
ACTIVITY_RISK	MAX(product risk, channel risk, counterparty risk, corridor risk, cash indicator).
ENTITY_EFFECTIVE_RISK	Watch-list precedence (exempt > trusted > risk); where risk lists only, MAX of all matched list values; where no list match, MAX(KYC risk, geography risk, business risk).
PROXIMITY_TO_THRESHOLD	1 – (|amount – reporting threshold| ÷ reporting threshold). Values above 0.95 indicate deliberate placement just below the line.
NORMAL_LIVING_EXPENSE_RATIO	Spend on groceries, utilities, rent and healthcare ÷ total debits. A near-zero ratio on a personal account with substantial throughput is a trafficking and mule indicator.

Closing note — what this framework is really for
A transaction monitoring system does not exist to generate alerts. It exists so that, when the bank is asked why it did not see something, it can answer with a document, a threshold, an approval, a name and a date — or, better, so that it saw it. Every control in this framework is designed around that single test.

Prepared by AML Base LLC. Meridian National Bank, N.A. is a fictional institution; all thresholds, volumes and figures are illustrative and calibrated for instructional realism, not drawn from any actual institution.