Financial & Transactional
Analytic Solutions
Australia

Resolving Risk Assessment Dimensions into Sub-Typology Exposure Profiles

David Coppin, CAMS | Principal Consultant
Financial & Transactional Analytic Solutions Australia
david.coppin@fintas.com.au
linkedin.com/in/david-coppin

March 2026

Abstract

Reporting Entities (REs) typically assess Financial Crime (FC) risk through dimensioned Risk Assessments (RAs) that rate exposure across jurisdictions, products, channels, and customer types, while implementing Transaction Monitoring (TM) detection through behavioural rules and scenarios. Both frameworks are necessary, yet they are rarely designed to interoperate. As a result, the linkage between RA outputs and TM detection design priorities is often reconstructed through narrative justification rather than governed resolution logic.

This paper proposes a structured method for resolving RA dimensions into sub-typology exposure profiles. The approach operates in two steps. First, each sub-typology is mapped to the RA dimension values through which it is enabled within the operating model. Second, per-dimension exposure profiles are calculated using RA-owned risk ratings and exposure proportions, producing a comparative view of exposure concentration without collapsing unlike dimensions into a single aggregate score.

The method does not attempt to estimate occurrence probability, detection effectiveness, or overall control performance. Instead, it provides a governed and repeatable resolution layer that supports scoping, prioritisation, and treatment discussion at sub-typology level. By formalising the intersection between dimensioned risk assessment and TM detection design, the approach strengthens traceability, improves explainability, and reduces reliance on narrative reconstruction in TM governance.

Table of Contents

The Operating Problem: Risk Assessment and Detection Use Different Languages 5
Why Sub-Typologies Are the Unit of Linkage 5
Scope and Limits of the Method 6
Method Overview 6
Implementation Prerequisites 7
Step 1: Mapping Risk Assessment Dimension Interactions 7
The mapping standard 8
A practical mapping register 8
Scale as a design constraint 9
What unmapped means 9
Treatment States After Mapping 9
Step 2: Per-Dimension Exposure Profiling 10
Rules for proportion 10
Score calculation 11
Interpretation discipline 11
Why dimensions remain separate 11
A Worked Example 11
Step 1: Map the sub-typology 12
Step 2: Apply Risk Assessment-owned inputs 12
What the Output Actually Produces 13
How the Method Informs TM Detection Design 14
Scenario selection 14
Design anchoring 14
Treatment posture 14
Governance: The Core Artefact Is the Register 14
Limitations and Future Extensions 15
Closing Thoughts 16
References 17

A Practical, Data-Driven Link Between Risk Assessment Dimensions and Transaction Monitoring Detection Priorities

Within Australia's Financial Crime (FC) prevention regime, Reporting Entities (REs) generally maintain two separate but related frameworks. The first is a Risk Assessment (RA), which expresses the RE's money laundering and terrorism financing exposure across dimensions such as jurisdiction, product, channel, and customer type [1]. The second is a detection estate, expressed through rules, scenarios, thresholds, segmentation, and other control mechanics that identify potentially suspicious behaviour.

Many REs also maintain typology and indicator libraries to support governance, communication, and control traceability. That is useful, but it does not by itself solve a persistent operating problem, because the RE still needs a repeatable way to translate dimensioned RA outputs into defensible TM detection design priorities.

In practice, that translation layer is often weak. The linkage between the RA and the TM detection estate is frequently reconstructed through workshops, narrative justification, inherited control libraries, or issue-driven responses. A typology is identified as relevant, a scenario exists or is built, thresholds are tuned, and a governance paper is written to explain the result. Even where those activities are reasonable in isolation, the mechanical path between RA inputs and TM detection design decisions is often unclear.

This paper proposes a deliberately simple method to address that problem. It introduces a governed, auditable approach for resolving RA dimensions against sub-typology risks. The method maps each sub-typology to the RA dimension values through which it can express in the RE's operating model, then quantifies per-dimension exposure using RA-owned risk ratings and exposure proportions. The result is a repeatable profile showing where exposure concentrates for each sub-typology, so scoping, prioritisation, and treatment discussions can be conducted on a more defensible basis.

The proposition is not that the RA should replace typological risk traceability, nor that arithmetic can replace judgement. The proposition is that REs need an explicit resolution layer between dimensioned RA outputs and TM detection design if they want that linkage to remain intelligible, repeatable, and governable over time.

This paper extends my earlier work, *Typological Risk Traceability. Mapping Typological Risk to Rules: A Foundational Transaction Monitoring Requirement* [2]. In that paper, I argue that typologies are necessary governance objects, but too broad to serve reliably as the operational unit of TM detection design. They are useful for ownership, communication, and reporting, but they do not show where a sub-typology manifests in the operating model. The more natural unit of detective concern is the sub-typology. That distinction matters here because typologies remain the governance wrapper, while sub-typologies are the unit through which RA dimensions can be resolved into something operationally useful.

The Operating Problem: Risk Assessment and Detection Use Different Languages

Most REs do not lack risk information. What they lack is interoperability between the way vulnerability to FC risk is assessed and the way TM detection (controls) are built. The RA typically expresses exposure across dimensions such as jurisdiction, product, channel, and customer type. TM detection, by contrast, is implemented through behavioural logic expressed as rules, scenarios, features, and thresholds. These are not the same language, and in many REs, they are not produced through the same methodology or by the same teams.

The consequence is predictable – routine design and governance questions become workshop questions: why a scenario exists at all, why it is applied to this cohort but not another, why sensitivity is higher here than there, which RA inputs justified the build effort, which sub-typologies are genuinely relevant to the operating model, and which risks are excluded and on what basis.

Without a mechanical translation layer, prioritisation becomes vulnerable to narrative competition involving recent issues, stakeholder preference, inherited scenario libraries, vendor defaults, or simple path dependence (building similar to what is already there). The detection estate then drifts in familiar ways. Controls are applied too broadly because exposure cannot be distinguished cleanly. Material sub-typologies are missed because high-exposure intersections are not resolved explicitly. Tuning decisions vary across products, channels, and types because there is no stable basis for what proportionate coverage should mean.

The problem is therefore not that either artefact is necessarily incomplete. The problem is that, in many REs, they are not interoperable by design.

Why Sub-Typologies Are the Unit of Linkage

Typologies remain important because they are how REs usually communicate FC risks. They support governance, ownership, reporting, and policy alignment, and for those purposes typology-level views remain useful.

Typologies are, however, too broad to function as the primary unit for translating RA outputs into detection priorities. A typology usually contains multiple sub-typological behavioural expressions. It may cover a family of mechanisms, channels, products, geographies, and customer behaviours under a single thematic label. That is helpful at governance level, but it does not show where within the RE's operating model the risk actually intersects. It does not identify which products, which corridors, which channels, or which customer types matter most, and it therefore does not provide a sufficiently precise basis for scenario selection, segmentation, or treatment prioritisation.

Sub-typologies are better suited to that task because, at a sub-typology level, the behavioural expression is sufficiently specific to be mapped to operating-model intersections and translated into control logic. That is the level at which one can say that a behavioural expression is enabled through these products, these channels, these jurisdictions, and these customer types. It is also the level at which one can compare relative concentration without collapsing unlike behaviours into a single thematic label.

In this context, **enabled** does not mean merely that the RE operates in a given jurisdiction, offers a given product, or services a given customer type in general terms. It means that the specific behavioural expression described by the sub-typology is contextually relevant and technically possible through that particular dimension value within the RE's actual operating model. A dimension value is enabled only where the mechanics required for the sub-typology to occur are genuinely present. For example, if a sub-typology depends on cross-border outward payments initiated through online banking by SME customers, then it is enabled only in those jurisdictions where the RE supports the relevant corridor, only in those products that permit that transfer capability, only through channels that allow customer-initiated execution in that form, and only for customer types that can realistically access and use those capabilities. Where any of those operating-model conditions are absent, the dimension value is not enabled for that sub-typology, even if the RE has exposure in that dimension more broadly. **Enabled is a statement of feasibility at sub-typology level, not a general statement of enterprise presence.**

The core proposition of this paper is therefore straightforward: typologies should remain the governance layer, but resolution and prioritisation should occur at the sub-typology layer, because that is the level at which RA dimensions can be mapped, quantified, and defended.

Scope and Limits of the Method

Before setting out the method, its boundaries should be explicit. This approach does not estimate the probability that a given sub-typology will occur. It does not estimate the probability that a scenario will detect that sub-typology effectively. It does not measure control effectiveness, and it does not replace design judgement, validation, quality testing, or feedback loops.

Instead, the method does something narrower and more practical. It identifies where a sub-typology is enabled within the RE's operating model, and where the RA indicates concentrated exposure across the RE's dimensions. In other words, it creates a governed exposure-and-relevance profile that can support scoping, prioritisation, and treatment discussion.

That narrower claim matters for technical defensibility. A sub-typology with a concentrated exposure profile is not automatically the highest-priority TM detection design candidate. Other considerations still matter, including data availability, feature feasibility, control alternatives, operational capacity, legal obligations, and known control performance. What this method contributes is a more structured starting point for those discussions.

Method Overview

At a high level, the method takes a defined sub-typology library, a structured RA dimension set, and a governed mapping standard as inputs. It then produces, for each sub-typology, a mapped set of dimension values and a per-dimension exposure profile. In effect, the method turns RA dimensionality into a repeatable sub-typology-level decision support view without collapsing unlike dimensions into a single aggregate score.

Step 1: Mapping

Each sub-typology is mapped to the RA dimension values through which it can express in the RE's operating model. The dimensions used here are jurisdiction (***J***), product (***P***), channel (***Ch***), and customer type (***C***). This step records where the sub-typology is enabled.

Step 2: Per-Dimension Exposure Profiling

For each mapped sub-typology, a per-dimension score is calculated using RA-owned inputs: a numeric encoding of the RA risk rating for each mapped value, and the within-dimension exposure proportion for that value. The outputs remain separate by dimension and are not blended into a single cross-dimension total.

The result is a profile of the form:

profile = {J, P, Ch, C}

This profile shows where exposure concentration sits for that sub-typology and which dimension is the dominant driver.

Implementation Prerequisites

The method is easiest to apply where the RE already maintains the RA as a structured dataset rather than only as narrative text. At minimum, effective implementation requires a stable sub-typology library, RA ratings at dimension-value level, a RA-owned exposure basis for each dimension, a governed mapping standard, a defined RA population cut such as enterprise, business unit, or product line, and version control over mappings, inputs, and outputs.

Many REs will have some, but not all, of these elements in place. Where the RA is primarily narrative, where dimension granularity is coarse, or where exposure data is inconsistent, additional data treatment may be required before the method can be applied reliably.

Step 1: Mapping Risk Assessment Dimension Interactions

Mapping is the first mechanical step because it forces the RE to make explicit where a sub-typology can actually express. This is not a judgement about likelihood. It is not a weighted relevance score. It is a statement that a sub-typological risk intersects a real operating-model dimension value that the RE has exposure to.

The mapping exercise is not a theoretical catalogue of every way the sub-typology could manifest in the abstract; it is a controlled record of the RA dimension values through which that sub-typology is enabled in the RE's current operating model.

The basic rule is simple: for each sub-typology, record the RA dimension values through which that expression is genuinely enabled in the operating model. That includes, for example, jurisdictions in which the relevant pathway exists, products that provide the relevant capability, channels through

which the behaviour can be executed, and customer types through which the expression can realistically manifest.

The mapping standard

Because the mapping step is decisive, it needs discipline. A weak mapping standard will simply relocate subjectivity into Step 2 and then disguise it with calculation.

A conservative mapping standard is therefore required. Mapping is limited to those dimension values at which the sub-typology is genuinely enabled, rather than extended by broad association. Where capability is conditional or constrained, that context is noted in the evidence field so the basis for inclusion is clear. The aim is simply to ensure that mapping reflects operating-model reality rather than thematic proximity.

The mapping output is binary by design, so a dimension value is either mapped or not mapped. That binary treatment is intentional because it avoids pseudo-scoring at the mapping stage and preserves a clear separation between enablement and quantification.

A practical mapping register

The method requires a governed register because mapping decisions are exactly where narrative drift usually returns.

The register is the explicit intersection layer between the RA and typological risk traceability. Each row records one testable claim: that a sub-typology intersects a particular dimension value, and why that claim is being made.

A minimal schema is sufficient:

- sub_typology_id
- dimension_type
- dimension_value
- mapping_basis_note

Only positive intersections are recorded. The presence of a row indicates that the sub-typology is enabled at that specific dimension-value intersection. Non-mapped dimension values are not recorded individually. Absence therefore reflects non-enablement at that intersection rather than an incomplete negative catalogue.

The ***mapping_basis_note*** field is the lineage and evidence anchor and should explain the operating-model basis for the mapping in a way that can be reviewed later without re-litigating the original judgement.

Typical mapping basis notes might refer to corridor availability, product capability, channel initiation mechanics, customer cohort definitions, policy or onboarding constraints, and broader operating-model limits.

Scale as a design constraint

This approach is multiplicative in scale. Under the per-dimension schema used here, the theoretical upper bound is:

rows <= (# sub-typologies) × (# jurisdictions + # products + # channels + # customer types)

For example, with 50 sub-typologies, 200 jurisdictions, 10 products, 10 channels, and 10 customer types, the complete upper bound is approximately 11,500 mapping records.

In practice, mapping is usually sparser than that because most sub-typologies do not intersect every value. Even so, the scaling behaviour should be treated as a deliberate design trade-off. The method buys granularity without requiring the RE to maintain full indicator-level mapping, but it does require a governed intersection register of meaningful size.

What unmapped means

A sub-typology may be unmapped because the RE does not have the enabling operating-model intersections. In practice, that usually means one or more of the following conditions applies. The RE does not:

- operate in the relevant jurisdictions or corridors,
- offer the required product capability,
- provide the required channel capability or provides it only in a materially constrained form, or
- have the enabling customer cohort at meaningful scale.

Where a sub-typology has no mapped dimension values at all, that overall determination should be recorded explicitly at sub-typology level for the relevant RA version. This does not require governance of every non-mapped dimension value; it requires only that whole-sub-typology non-enablement is deliberate rather than silent.

Treatment States After Mapping

To keep the method logically clean, it is helpful to distinguish between several different states that are often conflated. A sub-typology may be:

- **Not enabled** - the required operating-model intersections are absent and the sub-typology is mechanically out of scope in the current operating model.
- **Enabled but low relevance** - the sub-typology is possible but the resulting exposure profile is low or diffuse relative to other in-scope sub-typologies.
- **In scope for detective treatment** - the sub-typology is enabled and sufficiently relevant to support monitoring design or coverage review.
- **In scope but better addressed through non-detective treatment** - where prevention, gating, friction, or other non-TM controls are the more appropriate primary response.

- **Relevant but currently constrained** - the sub-typology is enabled and relevant but current data, architecture, control capability, or operating conditions prevent immediate implementation.

These distinctions matter because enablement, relevance, and treatment are not the same thing.

Step 2: Per-Dimension Exposure Profiling

Once sub-typologies have been mapped, the RE still needs a way to distinguish between them because many sub-typologies will be in scope and not all can be treated equally.

Step 2 provides a constrained form of quantification. The purpose is not to create a risk model, but to locate exposure concentration using inputs the RA already owns or can defensibly derive. Each dimension score is calculated from two ingredients.

The first is a **risk rating index**, being a numeric encoding of the RA risk rating for the relevant dimension value. A simple example would be:
High = 3
Medium = 2
Low = 1

The second is **exposure proportion**, being the within-dimension share of the assessed population cut represented by that value.

Examples include share of relevant value or volume by country or corridor for jurisdiction, share of exposure or value by product type for product, share of initiation volume or value by channel for channel, and share of customers, value, or exposure by customer type.

Rules for proportion

For the method to remain defensible, the exposure basis needs discipline. The basis may differ by dimension, but it should not be selected opportunistically by sub-typology. Within a given RA version and population cut, the basis for each dimension should be defined once and held constant across all sub-typologies in that run, proportions within each dimension should sum to 1, the basis should be RA-owned or derived from the same underlying population logic used by the RA, and where different bases materially change interpretation, parallel profiles are preferable to blended measures.

This means, for example, that an RE may choose value share for jurisdiction, product balance share for products, initiation volume share for channels, and customer-count share for types, provided those choices are made deliberately, documented clearly, and applied consistently within the relevant run.

Score calculation

For each dimension, the score is the sum of ***risk_index(value) × proportion(value)*** over mapped values only.

So for jurisdiction:
J = Σ (risk_index(jurisdiction) × proportion(jurisdiction))

over all mapped jurisdictions for the sub-typology.

The same structure applies to products, channels, and customer types:
P = Σ (risk_index(product) × proportion(product))
Ch = Σ (risk_index(channel) × proportion(channel))
C = Σ (risk_index(customer type) × proportion(customer type))

The resulting profile is:
profile = {J, P, Ch, C}

Interpretation discipline

These scores should be treated as relative prioritisation aids rather than as absolute measures of risk. They are most safely interpreted as useful for comparison within a consistent RA version and population cut, for ranking sub-typologies within a dimension, for identifying which dimension is the dominant driver for a sub-typology, and for highlighting concentration patterns that warrant design or treatment discussion.

They should not be interpreted as a direct measure of suspiciousness, a probability of occurrence, a proxy for scenario effectiveness, a cardinal scale on which one score can be read as a clean multiple of another, or a number that can be compared cleanly across REs with different RA methods.

Why dimensions remain separate

The method intentionally avoids blended cross-dimension totals. This is partly to avoid arbitrary weighting schemes, but it is also a recognition of the method's limits. Jurisdiction, product, channel, and customer type are not independent in reality. A high ***Ch*** score and a high ***P*** score do not prove that the same underlying cohort is driving both. The method therefore provides per-dimension marginal views rather than a full joint interaction model.

It is better understood as a structured first-order approximation that identifies concentration drivers, not as a replacement for deeper intersectional analysis where the RE's data and operating-model maturity support it.

A Worked Example

Assume a RE is assessing sub-typology ***ST-017***, which has a cross-border element and a channel dependency.

Step 1: Map the sub-typology

sub_typology_id	dimension_type	dimension_value	mapping_basis_note
ST-017	jurisdiction	AU	Relevant payment flows include AU corridor exposure
ST-017	jurisdiction	SG	Relevant payment flows include SG corridor exposure
ST-017	product	business_account	Business accounts provide the relevant transfer capability
ST-017	channel	online_banking	Online banking initiates the relevant transaction behaviour
ST-017	customer type	SME	SME customers represent an enabling cohort for the expression

At this illustrative point, the RE has established that the sub-typology is enabled across specific operating-model RA dimension intersections. In practice, a register would typically contain a materially larger number of dimension-value intersections, reflecting the full set of jurisdictions, products, channels, and customer types relevant to the operating model. The table above is intentionally simplified for clarity of illustration; the method itself scales with the granularity of the RA dimension set and the breadth of the RE's operating footprint.

Step 2: Apply Risk Assessment-owned inputs

Assume the RA provides the following values.

Jurisdiction Inputs

jurisdiction	risk_index	proportion
AU	3	0.15
SG	2	0.05

Then:
J = (3 × 0.15) + (2 × 0.05) = 0.55

Product Inputs

product	risk_index	proportion
business_account	3	0.40

Then:
P = (3 × 0.40) = 1.20

Channel Inputs

channel	risk_index	proportion
online_banking	2	0.70

Then:
Ch = (2 × 0.70) = 1.40

Customer Type Inputs

customer type	risk_index	proportion
SME	3	0.25

Then:
C = (3 × 0.25) = 0.75

The resulting profile is:
profile = {J=0.55, P=1.20, Ch=1.40, C=0.75}

The dominant driver is therefore:
argmax(J, P, Ch, C) = Ch

The interpretation is not that this sub-typology is "1.40 risky". The interpretation is that, within the defined RA version and population cut, channel concentration is the strongest dimension-level driver of exposure for this sub-typological risk. That conclusion is operationally useful because it suggests that channel-specific design treatment should be examined first, while still considering the product and customer type context.

What the Output Actually Produces

The missing artefact in many REs is not a typology list, and not merely a scenario inventory. It is a governed, durable record showing how sub-typologies intersect the RA and how exposure concentrates by dimension.

For each sub-typology, the output should include the mapped dimension values and basis notes, the per-dimension profile ***{J, P, Ch, C}***, the dominant dimension, the treatment state, the treatment rationale, and the relevant RA version and population cut.

This gives the RE a repeatable, reviewable basis for asking and answering practical questions such as whether the expression is genuinely enabled in the operating model, which dimensions are driving its exposure profile, whether it is relatively more important than other in-scope sub-typologies, whether it should be detected, prevented, excluded, or deferred, and, if it is to be detected, what dimension should anchor the first design discussion.

How the Method Informs TM Detection Design

Once a sub-typology profile exists, design discussion becomes more disciplined. This paper does not claim that the profile determines scenario logic or calibration automatically. TM detection design still requires judgement, feasibility analysis, and control architecture thinking. What the profile provides is a more defensible starting point.

Scenario selection

The method encourages a shift away from broad thematic logic such as "one scenario per typology" or indicator-based rule detection, and toward scenario selection based on prioritised sub-typologies. That matters because typologies are governance categories rather than control specifications, and indicators are often too granular and contextually limited[1]. A sub-typology profile provides a clearer basis for asking whether a specific sub-typology warrants dedicated logic, inclusion within an existing control pattern, or no detective treatment at all.

Design anchoring

The dominant dimension can provide a practical anchor for initial design thinking. Jurisdiction-driven sub-typologies may justify corridor overlays, country groupings, or cross-border feature emphasis. Product-driven sub-typologies may justify product-specific logic, product-tuned thresholds, or lifecycle-sensitive controls. Channel-driven sub-typologies may justify channel-specific behaviour logic, velocity analysis, or alignment to channel constraints. Customer type-driven sub-typologies may justify segmentation overlays, or differentiated thresholds. This does not mean the other dimensions are irrelevant. It means the method provides a more structured basis for deciding where to begin and how to explain the decision.

Treatment posture

The method also sharpens treatment discussion. Some sub-typologies may be relevant but poorly suited to post-event monitoring as the primary control response. In those cases, prevention, onboarding restrictions, product design, friction, or gating controls may be more appropriate. A structured profile gives that conclusion a clearer analytical basis, rather than leaving it to rest primarily on implicit judgement or anecdotal preference.

Governance: The Core Artefact Is the Register

If the method is performed informally, it will tend to decay into discussion and institutional memory; if it is governed properly, it becomes a more durable institutional asset whose logic can be re-run, reviewed, and defended over time. The central governed artefact is the register, because it preserves the linkage between the RA and sub-typology treatment decisions over time.

[1] Please read this foundational paper from which this paper extends - *Typological Risk Traceability. Mapping Typological Risk to Rules: A Foundational Transaction Monitoring Requirement*

At minimum, the governed record should store the sub-typology identifier and narrative, mapped dimension values and mapping basis notes, RA rating and exposure snapshot as at RA version ***X***, the population cut used for the run, the per-dimension outputs ***{J, P, Ch, C}***, the dominant dimension, the treatment state and rationale, and review cadence and change triggers.

Versioning is critical because, when the RA changes, the method should be re-runnable. The RE should be able to compare one run to the next and identify new mappings, removed mappings, changes in dimension scores, changes in dominant dimension, and changes in treatment decision. That allows the linkage between RA and TM detection design to remain live between major review cycles rather than being re-created ad hoc each time a challenge arises.

Limitations and Future Extensions

This method is intentionally simple, and that simplicity is both its value and its limit. Several limitations should be recognised explicitly:

- The method is only as good as the RA structure behind it. Where the RA is coarse, narrative-focused, stale, or inconsistent at dimension-value level, the method will inherit those weaknesses. It is strongest where the RA is already maintained as a structured, governable dataset.
- The mapping step remains judgement-based. The method constrains mapping judgement, but it does not eliminate it. Governance discipline is therefore essential, because poor mapping practice will produce poor outputs regardless of the arithmetic that follows.
- The method does not model full interaction effects. The profile preserves dimensional separation and does not model full joint intersections across jurisdiction, product, channel, and customer type. It is therefore a first-order approximation rather than a complete exposure model.
- The scores are relative rather than absolute. The outputs are useful for internal ranking and interpretation within a defined RA version and population cut, but they should not be mistaken for precise absolute measures of FC risk.
- The method does not measure control effectiveness. A strong exposure profile does not prove that an available scenario will detect the relevant behaviour well. Effectiveness still depends on data quality, logic design, calibration, alert quality, investigative capability, and feedback mechanisms. These limitations do not weaken the proposition so much as define it properly.

Future extensions are possible. A RE with sufficient maturity may choose to extend the approach through parallel value and volume profiles, more granular corridor treatment, explicit capability constraints, weighting overlays, or integration into broader control effectiveness and validation frameworks. Those extensions may be useful, but they should be treated as later maturity steps rather than prerequisites.

Closing Thoughts

REs already assess risk in dimensioned terms, while TM detection design is executed in behavioural terms. That difference is not merely semantic; it is one of the main reasons REs struggle to maintain a defensible linkage between RA outputs and the controls they build.

Against that background, this paper proposes a practical resolution layer for the problem, although the method itself is intentionally modest in scope. It does not claim to predict behaviour, to replace judgement, or to solve control effectiveness; rather, it provides a governed and repeatable way to resolve RA dimensions into sub-typology exposure profiles, so REs can make scoping and prioritisation decisions on a basis that is more explicit, more auditable, and less dependent on narrative reconstruction.

The practical contribution lies in disciplined interoperability between artefacts that are usually maintained separately. A governed mapping register makes operating-model intersections explicit. A per-dimension exposure profile shows where concentration sits for each sub-typology without collapsing unlike dimensions into a single blended total. Taken together, those outputs give typological risk traceability a more operational bridge into TM detection design.

What it offers is a more stable and explainable basis for answering questions that otherwise tend to become debates about preference, history, or inherited control logic. REs can therefore explain more clearly why a sub-typology is in scope, why another is excluded, why one cohort is treated differently from another, and why a given design effort is proportionate to the exposure the operating model actually presents. Those are not minor governance questions. They are central to whether a detection estate can be explained, maintained, and defended over time.

References

1. AUSTRAC. Money laundering/terrorism financing risk assessment. Accessed 5 February 2026. https://www.austrac.gov.au/business/core-guidance/amlctf-programs/money-launderingterrorism-financing-risk-assessment

2. Coppin, David. (2025) Typological Risk Traceability. Mapping Typological Risk to Rules: A Foundational Transaction Monitoring Requirement. Available at: https://www.griffith.edu.au/__data/assets/pdf_file/0024/2225364/Typological-Risk-Traceability-White-Paper-FINAL.pdf