GRIFFITH UNIVERSITY

Financial & Transactional
Analytic Solutions
Australia

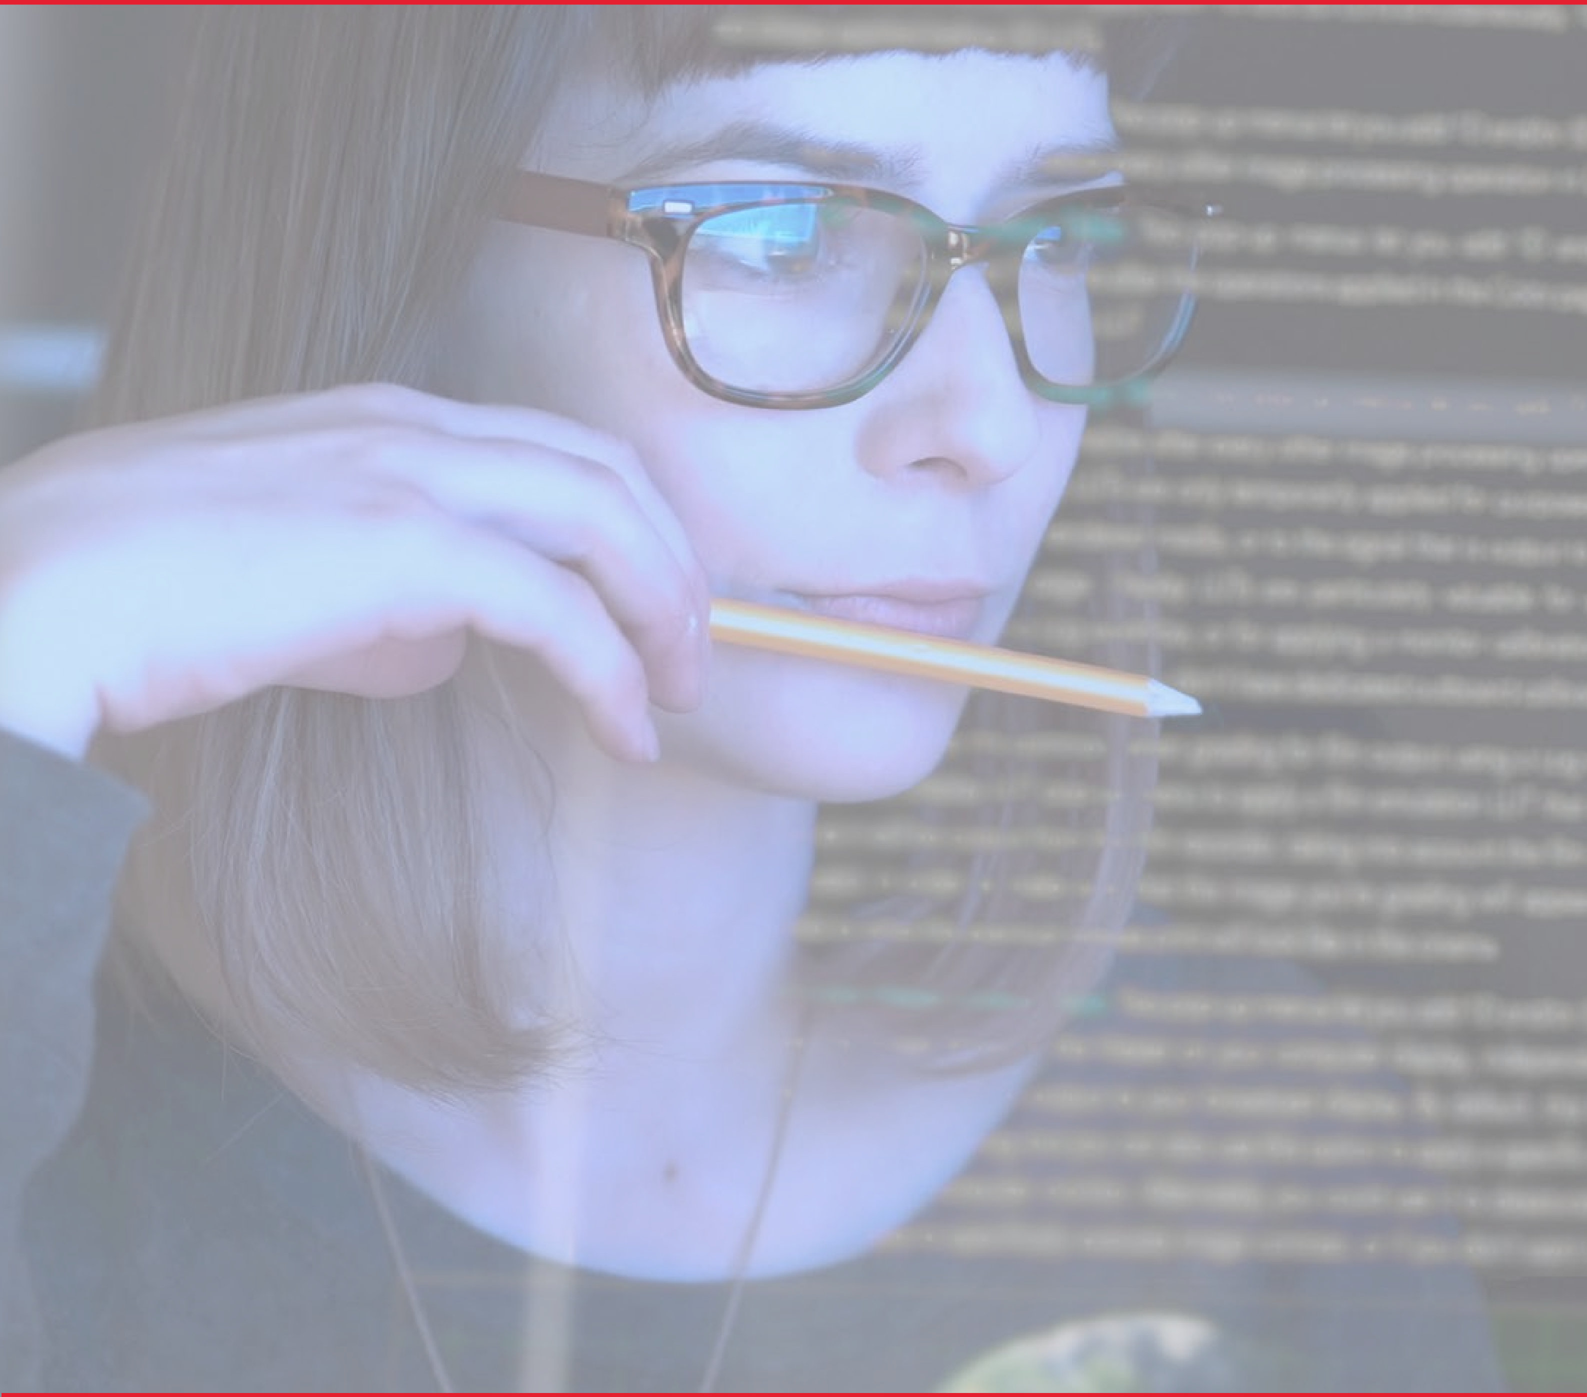

White Paper Series: No. 01.2025

Typological Risk Traceability

Mapping Typological Risk to Rules: A Foundational Transaction Monitoring Requirement

Academy of Excellence in Financial Crime
Investigation and Compliance

Typological Risk Traceability
Mapping Typological Risk to Rules: A Foundational Transaction Monitoring Requirement

David Coppin - Principal Consultant,
Financial & Transactional Analytic Solutions Australia

Academy of Excellence in Financial Crime Investigation and Compliance
White Paper Series No. 01.2025
November 2025

The Academy of Excellence in Financial Crime Investigation and Compliance collaborates with industry partners to provide national thought leadership and enhance training programs addressing global financial challenges, our initiatives reflect a shared commitment to integrity and resilience in combating evolving financial threats.

griffith.edu.au/academy-financial-crime

✉ *fcic@griffith.edu.au*

linkedin.com/…/academy-of-excellence-in-financial-crime-investigation-and-compliance

GRIFFITH UNIVERSITY

Contents

Abstract

Transaction Monitoring (TM) frameworks are critical in detecting financial crime, yet many Reporting Entities (REs) struggle to establish clear linage (traceability) between typological risk intelligence and implemented TM logic. This article explores the importance of contextualising risk signals for operational efficiency and effectiveness, the foundational distinction between rules and scenarios, and clarifies how typological risks should influence detection architecture. The article then describes an 8-layer Typological Risk Traceability Model designed to link source publications through to rules and scenarios in a structured, risk-aligned manner. By grounding TM design in typological clarity, the model supports defensibility, improves governance, and enhances explainability. Practical illustrations and architectural considerations are included to assist REs in aligning detection logic with assessed risk. This approach does not guarantee effective TM detection; it needs to operate in an environment with appropriate feedback mechanisms. The Typological Risk Traceability model serves as a foundational structure upon which more effective, risk-informed TM systems can be built.

The Imperative for Traceability: Mapping Risks to Rules

Transaction Monitoring (TM) requires more than a well-intentioned Anti-Money Laundering/Counter-Terrorism Financing (AML/CTF) Program or a feature-rich vendor solution. It requires transparency, structure, and traceability between clearly defined Financial Crime (FC) risks and the TM controls deployed to address them. For Reporting Entities (REs), the ability to demonstrate—step-by-step—how a risk described in a typology publication from sources such as AUSTRAC, FATF, or the Egmont Group leads to implemented detection logic (e.g., rules, scenarios) is not only an indicator of program maturity, but an established regulatory expectation [1].

This expectation plays out in the context of TM's operational reality: processing immense data volumes—often millions of transactions, product/account, relationship, customer records, and behavioural events every day to identify rare events - specific FC risks. Without a deliberate, risk-aligned design, this scale can overwhelm investigative capacity, dilute analyst focus, and degrade decision quality. This paper is written primarily for managers and leaders accountable for TM design and oversight, but the concepts will also be relevant to operational teams, compliance professionals, and senior executives seeking assurance that their TM frameworks are both effective and defensible.

Mapping typological risks to TM logic is a key enabler in transparency, structure, and traceability. It allows FC teams to demonstrate alignment between detection logic and their assessed risk landscape, enables validators to assess typological risk coverage[1], and supports governance mechanisms in determining whether deployed monitoring mechanisms are fit for purpose. Despite its importance, many REs either struggle to implement this foundational component or treat it as a proxy for TM effectiveness—overlooking the fact that design coverage alone does not guarantee operational performance.

This article builds upon foundational ideas from my TM effectiveness article series [2][3][4], advancing the conversation in two key areas. First, it provides a distinction between rules and scenarios and explores how recognising this difference can inform more structured and risk-aligned TM systems. Second, it describes a conceptual framework for typological traceability—to map financial crime intelligence from source publications through to deployed detection logic, capturing relevance to the RE's risk assessment along the way.

This article addresses two critical dimensions in TM framework design: typological risk traceability (demonstrating alignment from risk intelligence to controls—"risk-to-rules mapping") and an element of TM detection architecture (the structures through which detection logic is executed and risks are surfaced and acted upon). These dimensions intersect in the later stages of the traceability model, as it is there that architectural concerns begin to emerge. Regarding TM detection architecture specifically, this article makes only a light pass—focusing more on how typological alignment can inform and shape detection strategy, rather than detailing the technical or engineering layers involved.

Further, I'd note that throughout this article, there is a focus on TM – detective controls typically executed post transactional event. These are normally distinct from the capabilities considered for preventative controls, such as real time monitoring or controls passively implemented within core systems. Preventative controls are a critical component of the FC risk management framework and should be considered for relevant FC risks.

The goal of this article is to provide FC professionals—particularly those without day-to-day involvement in TM design—with practical guidance for bridging the gap between risk understanding

[1] This view of typological coverage does not imply TM control effectiveness.

and control implementation. In doing so, this article offers a pathway to help REs & FC practitioners not only meet regulatory expectations but also build towards explainable, effective TM systems.

The Starting Point: Technical Compliance

For many REs relatively closer to the start of the FC prevention journey than others, the question may be asked – what rules do we need? I've fielded this question on several occasions in recent years. For context, one such occasion involved a prospective RE expanding into the Australian market. My reply was a countering-question, in effect, what risks are relevant to your business?

The mapping of relevant typological risks to rules may seem like an obvious and sensible starting point, when the suggestion is made. However, my experience hasn't seen this consistently demonstrated. It's an anecdote worth sharing, as it highlights that even experienced REs can overlook the foundational step of aligning controls to risk.

While mapping typological risks to rules is a critical starting point, it should not be mistaken for a complete solution. This exercise can clarify whether your TM rules reflect the threats you're actually exposed to—but it is only one part of a broader detection strategy. As I've covered in my article series on TM effectiveness, risk-to-rule mapping is necessary, but not entirely sufficient. It helps demonstrate alignment and coverage on paper (technical compliance) but doesn't guarantee that detective controls are effective in practice [2][3][4].

Demonstrating technical compliance is where some REs run into trouble—when the exercise becomes one of filling gaps in a control register, rather than anchoring those controls in a deep understanding of risk and effectiveness testing. The result is often superficial compliance rather than meaningful risk detection.

In practice, poorly designed TM detection frameworks often share common traits, such as shallow or unvalidated segmentation models[2], overly simple rules with fixed thresholds that are inconsiderate of risk assessment dimensions[3], and absent governance structures [4]. These traits may be reactive products of audit or regulatory findings, without consideration for how they fit into the broader detection strategy. Absent systematic testing to confirm that detection events represent credible, relevant risks, such systems may meet the letter of compliance while offering little real protection against financial crime.

Why This Matters: Aligning to your Risks & Regulatory Expectations

Supervisory bodies and international standard-setters have explicitly called for stronger alignment between risks and deployed AML/CTF controls (e.g. rules). In Australia, AUSTRAC's 2024 regulatory priorities emphasise that REs must demonstrate how their TM systems are aligned with the Risk Assessments (RAs) that underpin their programs [1]. This is not a superficial exercise—regulators expect the ability to trace detection logic to specific risks and sources. REs must

[2] Consider customer segmentation models driven solely by domain expertise without calibration and statistical validation.

[3] Consider:

- Catch-all product rules that trigger on generic activity without context against a defined typology;
- Fixed absolute transactional or aggregate thresholds, applied uniformly across all customers/ customer cohorts, irrespective of customer risk or length of relationship with the RE; or
- Jurisdiction-based triggers that treat geography as a standalone risk rather than as one indicator among many

[4] Consider validation & benchmarking activities, and monitoring frameworks for rule performance.

demonstrate that their TM environments are aligned to the risks they face, and that rules and scenarios are grounded in a sound understanding of typological risks.

The Financial Action Task Force (FATF), the global standard-setter for AML/CTF, reiterates this through its guidance on risk-based approaches, urging REs to use national and sectoral risk assessments, typology reports, and emerging threats as inputs to control design. In particular:

- FATF Recommendation 1 and its Interpretive Notes recommend countries and REs to identify, assess, and understand the ML/TF risks they face, and to apply mitigation measures that are proportionate to those risks. The guidance emphasises that where risks are higher, enhanced measures are required, and where risks are lower, simplified measures may be applied—but only with a strong understanding of those risks and subject to appropriate safeguards. This approach requires not just a documented RA, but tangible evidence that detection controls—such as TM rules and scenarios—are implemented based on assessed typologies and are actively managed through internal governance, monitoring, and review processes [5].
- FATF's 2021 Guidance for a Risk-Based Approach to AML/CFT Supervision goes further by instructing supervisors to evaluate how well REs operationalise their risk understanding. Supervisors are expected to assess not only whether an RE has a risk-based TM framework, but whether that understanding is reflected in its control design (e.g. TM rules). The guidance encourages the use of national RAs, sectoral insights, and typology reports as direct inputs into TM logic development. It further highlights the importance of evaluating the effectiveness of these controls, not just their existence, underscoring that REs must be able to demonstrate how their monitoring strategy mitigates risk [6].

While FATF's guidance provides an important global foundation, its direction often lacks sufficient clarity on how to apply proportionality in practice. As noted in a recent policy critique [7], FATF guidance leaves uncertainty around how to factor likelihood into risk assessments, what levels of risk may be tolerated as "low," and how to embed formal concepts of risk appetite and tolerance into AML/CTF program design. This creates challenges for REs seeking to design TM frameworks that are both effective and practical reflections of relevant FC risks. Per the suggestion of Hunter, et. al (2007) [7], established enterprise risk management frameworks such as ISO 31000 can help bridge this gap—offering structured methods to define, document, and operationalise risk thresholds in a way that aligns detection strategy with organisational risk posture and supports defensible proportionality.

Refocusing on the Australian context, AUSTRAC's guidance follows FATF's risk-based expectations, reinforcing that TM must be grounded in each RE's specific RA. TM logic should address identified ML/TF risks, adapt to emerging typologies, and be demonstrably aligned to RAs and typology reports [8]. Crucially, this alignment must exist not only on paper—REs are expected to clearly show how risks surface through rules, scenarios, and detection strategies within their systems.

This expectation is underscored by (somewhat) recent enforcement action. In July 2023, the Federal Court ordered Crown Melbourne and Crown Perth to pay a $450 million penalty for breaches of the AML/CTF Act 2006. The court found that Crown's AML/CTF programs were not based on appropriate RAs and lacked adequate systems and controls to manage their risks [9].

Taken together, these international standards, local guidance and enforcement action make one thing clear: alignment between risks and detection logic is not optional—it is an expectation. Effective risk-to-rule mapping serves as the cornerstone for ensuring transparency and coherence within TM frameworks. But importantly, a cornerstone, does not a wall make (again, consider what feedback mechanisms are required to support effectiveness). The ability to clearly show how an external publication—such as an AUSTRAC typology report—results in a scenario within the TM system, allows REs to demonstrate coverage through controls of relevant financial crime risks and alignment between risk assessments and TM detection logic.

Conversely, without clear traceability, REs may struggle to identify coverage gaps or explain the rationale behind deployed controls—undermining explainability and oversight.

Critically, risk-to-rules mapping must carry through to how risks are operationalised—how they are logically surfaced, interpreted, and acted upon in practice. That means considering the structure and clarity of the detection logic itself, and how effectively it supports analysts in making risk-informed decisions. The next section explores the operational implications of detection design.

The Operational Implications of Detection Design

Why does it matter how risk is surfaced? The means through which risk is presented influences whether it will be meaningfully understood, consistently acted upon, and efficiently resolved. Even before considering the technical distinction between types of detection logic (which we'll explore shortly), it's important to establish that poor presentation of risk can dilute its interpretability and impact.

To understand why structure and presentation matter, consider TM from the perspective of operational teams. In high-volume investigative environments, efficiency isn't a luxury—it's a prerequisite for scalability and sustained effectiveness. Surfacing risk is only part of the task; presenting it in a form that is quickly comprehensible and relevant to an underlying FC typology is equally critical.

When risk is surfaced through isolated indicators, generic or loosely defined conditions, analysts are left to piece together the context themselves—essentially weaving a mental tapestry from fragmented signals. Consider:

- If a velocity rule triggers due to multiple deposits and withdrawals in a short space of time, how best should an analyst determine what typological risk the activity relates to?
- Should analysts be expected to assess individual indicators without typological context from the detection platform?
- Are they relying purely on training and experience to interpret the risk? Ask yourself – how confident are you in your training or internal accreditation processes?
- What is the operational value of flagging an indicator, if it doesn't meaningfully represent a typological risk on its own?

Poorly conceived TM logic—especially logic that represents isolated indicators or lacks direct alignment to known risk typologies—adds friction to operational processes:

- Consuming unnecessary investigative capacity, extending alert handling times
- Introducing inconsistency in decisioning and potential outcomes
- Eroding confidence in alerts over time through de-sensitising analysts to false positive volumes, diluting the precision of risk detection

The challenge with isolated indicator rules is that they lack the specificity needed to indicate credible risk on their own. While they may technically capture a suspicious behaviour, the absence of targeted context leaves analysts to interpret the typological significance—an inefficient and inconsistent process.

Some might argue that filtering noise is expressly the role of triage—to sort the wheat from the chaff. And while that's partly true, the more important question is how best to allocate operational effort. Should analysts be interpreting isolated indicators and stitching together context? Or should their focus be reserved for evaluation within a contextualised view of typological risks where the detection logic has already done the heavy lifting?

From Signal to Story: Structuring Risk Detection for Analysts

Effective TM detection design is not just about surfacing anomalies—it's about surfacing risk in a way that is readable, actionable, and aligned to investigative intent. TM logic should do more than raise flags; it should frame those flags in context. This framing requires decisions about when to surface simple signals and when to require convergence across indicators before triggering. The key questions for REs to consider are:

- When should risk surface as a single indicator, versus as part of a broader set of indicators?
- Is a single indicator ever sufficient to represent a typological risk?
- What is the relevance of single indicators to risk management - what investment of analyst time and resource is required to resolve?
- What capabilities are required to ensure that detection logic not only surfaces risk but contextualises it meaningfully?

These questions also require some consideration around detection architecture:

- Are there data products, such as segmentation models, classification logic or aggregate features, that are prerequisites for building robust TM detection logic?
- Can TM capabilities be designed to support both granular, indicator-level risk identification and more holistic, typological risk-based analysis?

The Cost of Ambiguity: Why Structure Matters

If TM logic is not considerately designed and aligned to defined risks, REs may encounter:

- Strategic risks: Misalignment with the RE's risk assessments, missed typologies, failure to meet regulatory expectations
- Operational inefficiencies: Redundant TM detective logic and inconsistent triage

These operational considerations point to a deeper structural challenge: when TM detection frameworks lack strategic consideration of how risk surfaces, even well-intentioned controls can fall short of their purpose. In the next section, we examine a common source of this challenge in the context of typological risk traceability: the blurred distinction between rules and scenarios. Clarifying this difference is necessary—not just for building scalable and explainable TM systems, but for maintaining a connection between identified risks and the logic deployed to detect them.

A Common Challenge: Inconsistent Terminology and Detection Design

The operational considerations discussed in the prior section often stem from how TM environments are built and evolve over time. Most REs don’t start with a blank slate—instead, operations teams inherit a patchwork of legacy logic, vendor rulesets, and responses to regulatory findings. Over time, this piecemeal development results in fragmented TM detection frameworks—collections of rules, detection components, and logic artefacts that lack cohesion. Without a guiding design principle or clear conceptual foundation, these systems evolve reactively rather than strategically. The result is often inconsistent structure, redundant logic, and unclear links between controls and the risks they are meant to address—particularly where definitions of rules and scenarios are blurred. As these frameworks expand, assessing typological coverage or justifying individual rules becomes increasingly difficult, undermining both effectiveness and traceability.

Oftentimes, fundamental to this experience is the question of “what is a rule”. It is not a trivial question; the answer has implications for TM frameworks, the TM detection solution architecture for surfacing risks, and operations. It’s also fairly common to see interchangeable terminology here – “rules” “scenarios”. In fact, I’ve used the term “rules” rather loosely to this point, to refer to detection controls broadly prior to establishing this point[5].

I encourage distinguishing between rules and scenarios, as they can represent different capabilities and approaches to financial crime risk management. Recognising their distinct purposes can improve TM detection design and enhance workflow strategies. What follows is my suggested interpretation.

Indicators and Rules: Building Blocks of Detection

Rules can be understood as the fundamental building blocks of detection logic. They are individual logical constructs designed to detect isolated behaviours or indicators (red flags). Typically, rules evaluate a single condition or a narrow set of conditions, providing a straightforward, explainable basis for generating a output (such as a detection event). For example, a rule might specify "more than 5 cash deposits under $10,000 within 7 days" or "international transfers over $x to a high-risk jurisdiction."

I'd hazard a guess that most readers will be familiar with rules lacking explicit typological focus—such as velocity rules (e.g., detecting rapid movement of funds) or cash threshold avoidance rules (e.g., flagging cash deposits just below reporting limits). These types of rules are often implemented because they are relatively easy to configure and can be directly aligned with regulatory thresholds or generalised ML/TF indicators. However, their simplicity can be a weakness—alone they lack context or fail to clearly identify nuanced patterns of suspicious behaviour. Further, the broad nature of rules lacking typological focus contributes to, for many REs, the ever-present challenge of false positives.

When rules are used independently, they often capture isolated behaviours rather than complex financial crime patterns. While rules are essential for flagging indicators, they lack the depth needed to fully represent how a typological risk may manifest. Therefore, rules are necessary but insufficient on their own to efficiently surface risk. With the right architectural view, they can be more effectively employed in the pursuit of targeted risk detection.

[5] Ultimately, the meaning of terms like "rules" and "scenarios" within a TM detection framework depends on how the RE chooses to define and use them. If the RE’s framework makes no functional distinction between the two, they effectively become synonyms—often influenced by prevailing vendor terminology.

Scenarios: Contextualising Typological Risks

Scenarios can be understood as detective constructs that combine multiple rules and features to reflect real-world behaviours. They are where typological risk understanding is operationalised.

Depending on an RE's detection solution architecture, scenarios can be built from various inputs: individual indicator-rules, rule outputs, or other derived features. These features might include native transactional or customer attributes or be pre-processed elements—such as aggregates or the outputs of non-detective models. In some cases, existing rules themselves may be reframed as non-detective models that output features, contributing to a broader detection construct rather than operating independently.

Scenarios capture risk in more contextualised and operationally relevant manner. Rather than generating alerts on isolated conditions, scenarios detect convergence—where multiple risk factors, over time or across relationships, collectively point to suspicious behaviour. This makes scenarios especially powerful for detecting complex typologies, such as phoenixing, layering, or trafficking.

Importantly, this is where the discussion of typological risk traceability begins to intersect with detection architecture. How features are constructed, managed, and orchestrated in the execution environment is an architectural consideration that underpins scenario performance and flexibility[6].

This conceptual structure becomes more tangible when applied to a specific typological risk. To illustrate this, consider a scenario designed to detect phoenixing, as outlined in the AUSTRAC's Fintel Alliance publication on phoenix activity (relevant to labour hire and payroll services) [9]. Such a scenario might incorporate several rules across customer profiles (personal and non-personal), account and behavioural contexts. Example indicators include frequent changes in company directors and beneficial ownership across related entities, repeated patterns of business deregistration followed by new company formation using similar trading names, and payroll payments made to the same employees across multiple entities (see Figure 1 for the full list of indicators). Individually, these behaviours may not be sufficient to trigger suspicion—but when combined within a scenario, they present a coherent risk profile aligned with known phoenix typologies and abuse patterns.

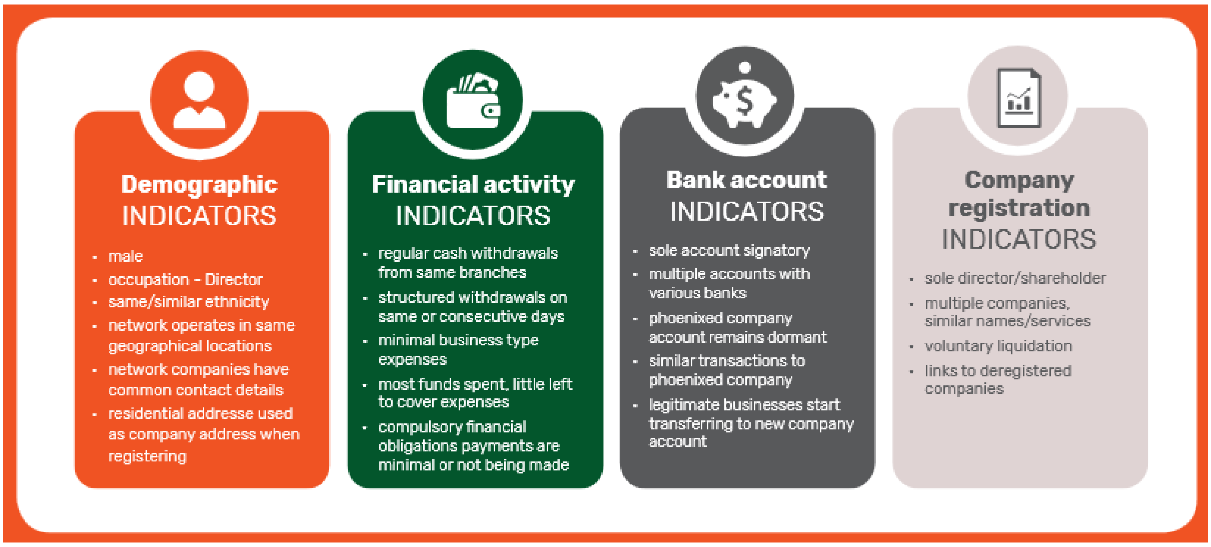

Demographic INDICATORS

- male
- occupation - Director
- same/similar ethnicity
- network operates in same geographical locations
- network companies have common contact details
- residential addresse used as company address when registering

Financial activity INDICATORS

- regular cash withdrawals from same branches
- structured withdrawals on same or consecutive days
- minimal business type expenses
- most funds spent, little left to cover expenses
- compulsory financial obligations payments are minimal or not being made

Bank account INDICATORS

- sole account signatory
- multiple accounts with various banks
- phoenixed company account remains dormant
- similar transactions to phoenixed company
- legitimate businesses start transferring to new company account

Company registration INDICATORS

- sole director/shareholder
- multiple companies, similar names/services
- voluntary liquidation
- links to deregistered companies

Figure 1: Phoenixing Indicators [10]

[6] While this article only touches lightly on those technical aspects, recognising this crossover is essential. The groundwork laid here—defining scenarios as structured expressions of sub-typological risk and composites of indicators/rules—provides a foundation for exploring architectural implications in future articles.

By integrating multiple indicators/rules into a scenario, scenarios offer the prospect of enhanced detection efficiency (reducing false positives) by focusing investigative resources on the convergence of indicators that are more likely to indicate financial crime.

This illustration underscores key implications for TM strategy. As financial crime typologies become more complex—and regulatory expectations continue to evolve—REs must move beyond isolated controls to cohesive, traceable systems. Designing or reviewing a TM detection framework should involve more than evaluating individual rules or scenarios; it demands a clear rationale for how typological risks are interpreted, prioritised, and operationalised through detection logic. Without this approach, fragmentation, ambiguity, and misalignment can quickly erode both operational effectiveness and explainability.

To address this challenge, the next section describes an approach for bridging typological intelligence and TM detection logic—offering a practical blueprint for traceability, logical modularity, and alignment with risk.

The Proposition: the 8-Layer Typological Risk Traceability Model

This section describes a structured approach: the 8-Layer Typological Risk Traceability Model. The model provides a practical method for linking typological intelligence to TM detection logic, ensuring traceability from external publications to deployed controls. It serves as a bridge between risk and operational design, while also informing the architectural principles needed to build efficient and effective detection systems.

The intent is not to prescribe a single standard, but to provide an illustration that can be tailored to a RE's unique operational context, regulatory obligations, and risk appetite. Depending on RE maturity, elements of what follow may seem familiar to colleagues who have undertaken typological traceability exercises.

By establishing lineage between risk assessment, external sources/ intelligence and TM detective logic into distinct layers—from source publications to deployed scenarios—this model offers a way to organise, rationalise, and explain detection logic. It supports both the design and validation of TM systems and can serve as a traceability register that links regulatory typologies directly to implemented controls.

Each layer is illustrated using an example: phoenixing, as outlined in AUSTRAC's (Fintel Alliance) Illegal Phoenix Activity Indicators: Labour Hire/Payroll Services (2019). Phoenixing involves the deliberate liquidation of a company to avoid liabilities (such as taxes and employee entitlements), followed by the re-emergence of a similar business through a different legal entity—often with the same personnel, trading names, or operations.

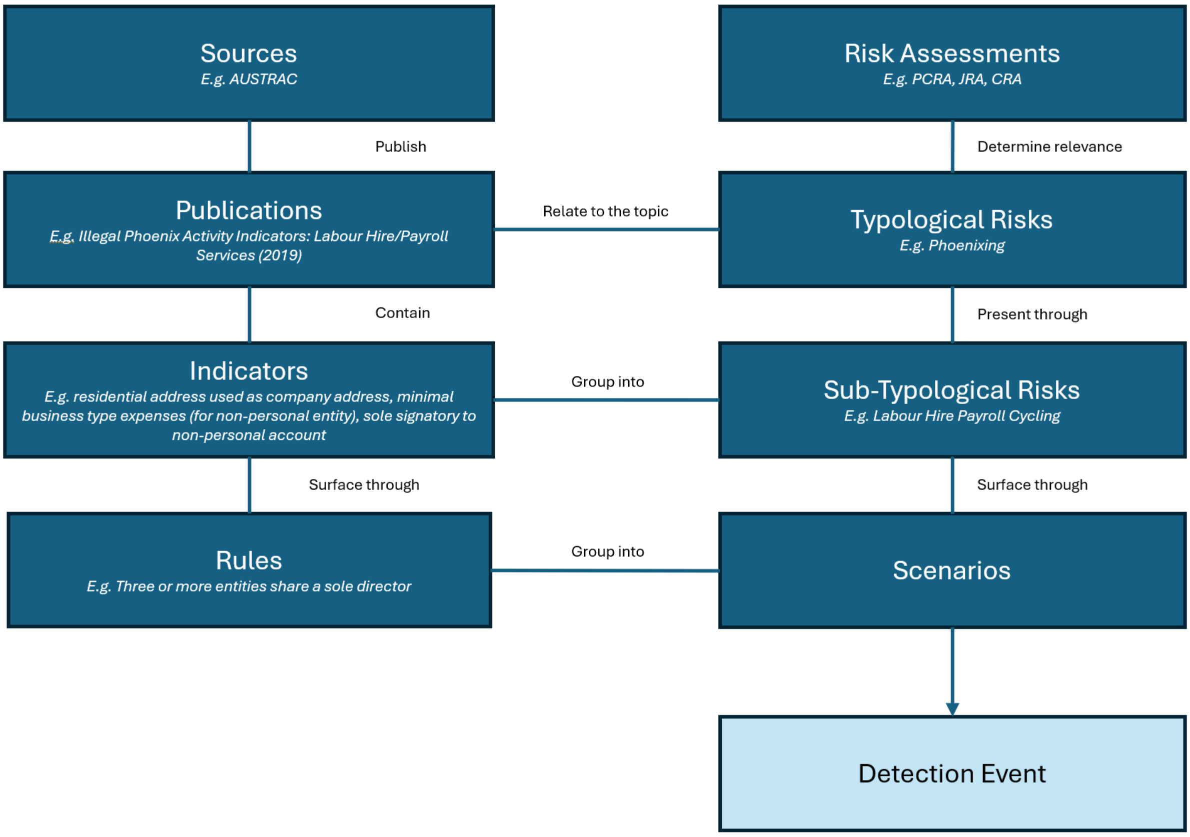
Sources
E.g. AUSTRAC
Publish
Publications
E.g. Illegal Phoenix Activity Indicators: Labour Hire/Payroll Services (2019)
Contain
Indicators
E.g. residential address used as company address, minimal business type expenses (for non-personal entity), sole signatory to non-personal account
Surface through
Rules
E.g. Three or more entities share a sole director
Risk Assessments
E.g. PCRA, JRA, CRA
Determine relevance
Relate to the topic
Typological Risks
E.g. Phoenixing
Present through
Group into
Sub-Typological Risks
E.g. Labour Hire Payroll Cycling
Surface through
Group into
Scenarios
Detection Event

***Figure 2: Illustration of 8-Layer Typological Risk Traceability Model*[7]**

Sources: Credible Authors

These are recognised authorities that produce information relevant to money laundering (ML), terrorism financing (TF), and broader financial crimes. Examples include AUSTRAC, FATF, Egmont Group, Wolfsberg, Interpol, and UNODC. These sources should be credible and legitimate. In mature REs, a formal process should be in place to review publications from such entities for relevance. We start our illustration with AUSTRAC as the monitored source.

Publications: Sources of Indicators

These are the actual documents—such as typology reports, risk assessments, or intelligence briefs. It is recommended that REs should have a formal process for reviewing and tagging publications for relevance. The example considered here is AUSTRAC's (Fintel Alliance) Illegal Phoenix Activity Indicators: Labour Hire/Payroll Services (2019)[8].

In addition to external sources, REs may want to recognise the value of internally generated intelligence. Investigative findings, internal risk reviews, and insights derived from Suspicious Matter/Activity Reports (SMRs/SARs) can provide rich, context-specific information on emerging

[7] Regarding the outputs of rules, scenarios, and TM systems, it's important to distinguish between detection events and alerts. A detection event is the system-level signal that a defined risk condition—such as a scenario—has been met. It is distinct from an alert, which is the workflow artefact routed to analysts for triage, investigation, and resolution. While often used interchangeably in conversation, the distinction matters: treating detection events as intermediate outputs enables aggregation or prioritisation before promotion to alert status. This supports scalable design, reduces noise, manages investigative workload, and preserves typological clarity in how alerts are presented to operations.

[8] I'd note that not all publications have consistently clearly stated indicators – the AUSTRAC phoenixing publication is one of the better examples.

patterns of behaviour. These sources—though not formally published—may reveal typologies or indicators that are highly relevant to the institution's customer base and operating environment. As such, they can serve as critical complements to external publications. REs can consider formalising a process for extracting relevant insights from internal casework and integrating them into the typological risk library and detection framework.

Indicators: Red Flags

From the publication, indicators are extracted—observable data points, actions or patterns that may signal risk. Indicators should be extracted and stored in a standardised format serving as TM detection primitives. For the full list of phoenixing indicators, refer back to Figure 1.

Typologies: Thematic Risks

This layer is the highest-level classification of threats into a recognisable risk theme, used across governance and risk assessment processes. REs may draw from regulatory typology taxonomies or establish their own. Examples include Child Sexual Exploitation, Illicit Wildlife Trade, Corporate Phoenixing, and Terrorism Financing. These categories form the bridge between external threat intelligence and the internal RA.

Sub-Typological Risks: Bridging Risk with Behaviour

Sub-typological risks articulate how a broader typology may manifest through specific, observable behavioural patterns. These act as a critical intermediate layer between thematic risk categories and detection logic—translating abstract threats into concrete risk expressions that can be monitored. Sub-typologies support more targeted and contextual scenario design by breaking down high-level risks into distinct, monitorable behavioural patterns.

For example, under Phoenixing it's possible to consider Payroll Cycling and Business Identity Reuse, to name just a couple of possible sub-risks:

- Labour Hire Payroll Cycling
  - Repeated creation and liquidation of payroll companies within a labour hire network.
  - New entities assume business functions of the liquidated ones with minimal operational change.
- Business Identity Reuse
  - New companies created with similar trading names, business activities, or contact details to deregistered entities.
  - Business identity continuity signals risk of phoenixing for ongoing operations.

REs may develop their own internal sub-typology definitions based on observed behaviours and case experience or leverage the taxonomy divisions provided within regulatory or industry publications—where available—to guide detection logic design and maintain consistency across the TM detection framework.

Risk Assessment: Factoring Relevance

Risk Assessments (RAs) play a critical role in determining how typological intelligence is operationalised. At a high level, it can act as a scoping tool—establishing which typologies and sub-typologies are relevant to the RE based on its customer base, product suite, transacting channel, geographic exposure, and business model.

Importantly, not every identified threat from external sources warrants coverage within the TM detection framework. Typologies deemed low-priority or outside the RE's exposure profile may be ruled out as candidates for TM detection logic—provided this exclusion is justified, documented, and subject to periodic reassessment. This helps prevent overengineering and reduces noise, while maintaining defensibility through traceability.

Conversely, some risks may be assessed as prohibitively high—outside the RE's risk appetite and unmanageable through detective means. In such cases, TM logic isn't the appropriate treatment; preventive controls such as onboarding restrictions, product limitations, or real-time interventions may be required instead.

Returning to the AUSTRAC phoenixing typology example, this may hold high relevance for REs servicing payroll or labour hire clients. However, a digital-only neobank with no non-personal or business account offerings may assign it low priority. REs may assign a relevance tier (e.g. high/medium/low) or designate a risk as out of scope to inform decisions around detection coverage, design prioritisation, and validation depth. In the case of the hypothetical neobank, the low relevance assigned to phoenixing may justify its exclusion from TM detection logic—provided this decision is documented, risk-based, and periodically reviewed.

Where typological risks are deemed in-scope for TM detection, RAs can also shape risk-based treatments—not just whether risk are covered, but how they're operationalised. This goes beyond translating typological indicators into detection logic; it concerns how that logic is applied in a risk-based manner. This approach is most commonly applied through threshold adjustments based on Customer Risk Assessments (CRA), often in conjunction with segmentation models, but can also extend to product, channel, and jurisdictional risk domains.

Rules: The Logic of Indicators

Rules can be the atomic units of detection—individual conditions that test for indicators.

For Phoenixing an example rule may include a customer-level trigger where the individual is the sole director of three or more entities. Rules may be reused across multiple scenarios or typologies and are most effective when clearly linked to standardised indicators.

The precise utility of rules in this illustrative model depends on the TM solution architecture. Rules could be implemented as detective or non-detective models[9]. In my illustration, they are conceived as non-detective, contributing to stored features. It's conceivable that they could be removed from the typological traceability model altogether, with direct mapping from indicators to scenarios. For many REs it's more likely that they'd need to consider how their current state aligns to this model, where detection events are likely generated from their rules – this would be a trivial expansion to the illustrated model.

[9] A detective model is one whose firing directly triggers a detection event—serving as a frontline control that surfaces risk for investigation. By contrast, a non-detective model outputs a feature or signal that other models can consume. Non-detective models can support more complex typological representations without being detection-event-generating mechanisms in themselves.

Scenarios: Control Constructs, Typologies in Practice

Scenarios can combine multiple rules or indicators into a higher-order control structure, reflecting how a sub-typological risk actually manifests. A scenario may involve:

- multiple rules triggering within a defined window (note prior consideration regarding solution architecture)
- weighted components or scorecards
- sequences of behaviour over time, or
- the convergence of transactional and demographic variables

Scenarios are the logical unit of TM systems that are tuned, validated, and subject to governance. In practical terms, they often evaluate as composite logical elements built using boolean operators (AND, OR, etc) to capture the convergence of multiple indicators —for example, a high-value transfer AND a link to a high-risk jurisdiction OR multiple rapid cash deposits within a short period.

In more advanced implementations, scenarios may also be expressed through weighted scoring models, enabling greater flexibility in reflecting typological complexity. Here, each indicator is assigned a weight based on modelling, with the scenario triggering when the cumulative score exceeds a set threshold. This can allows for greater flexibility and permissiveness in how different behaviours contribute to a detection event—such that a group of moderate-risk behaviours may generate a detection event in the same way a singular or relatively narrower set of high-risk behaviours might in a rule driven by the aforementioned boolean operators.

Implementing a Traceability Register

The phoenixing example illustrates how typological intelligence can be translated into structured, risk-aligned detection logic. By anchoring scenarios in credible sources and clearly defining each intermediate layer, REs can build TM systems that are not only transparent and defensible, but operationally effective. A traceability model helps introduce structure and intent to TM design—supporting stronger validation, improved alert quality, and enhanced explainability. A clear lineage from source intelligence through to scenario logic supports both internal governance and external assurance.

To operationalise this model, each layer—from source intelligence to deployed scenario—should be documented in a typological traceability register. This register becomes the RE's blueprint for TM design and development, linking detection logic directly to identified risks and their behavioural patterns. It allows institutions to demonstrate not just risk awareness, but how that awareness translates into targeted and explainable controls.

In practice, multiple typologies can share the same indicators—for example, changes in cash deposit behaviour could be relevant to both money laundering and tax evasion typologies. Conversely, a single typological risk will likely have many indicators. In such cases, the register could accommodate each relationship between indicator and sub-typology. A relational model can support this, where Indicators and sub-typologies are stored in separate tables, linked via table that describes the table relationship.

Importantly, this model does not exclude or limit the use of advanced analytics, including machine learning and other data-driven approaches. These approaches are increasingly vital for identifying subtle behavioural anomalies, complex network patterns, and previously unseen typological risks—particularly in high-volume or fast-evolving environments. However, advanced models are not a substitute for foundational design discipline. Even the most sophisticated detection approaches require clear design provenance, risk alignment, and explainability to be operationally useful and regulatorily defensible. The structured traceability model outlined here can provide the scaffolding needed to integrate these advanced techniques meaningfully, ensuring that innovation does not come at the expense of transparency, auditability, or strategic alignment.

Lastly, to reiterate, a traceability model does not guarantee effective TM. Effectiveness needs to be supported through feedback mechanisms embedded throughout the entire TM framework [2][3][4].

Building a Bridge Between Risk and Response

A sound TM framework doesn't just surface risk—it ensures that what is surfaced is relevant and contextualised. This alignment is central to both operational performance and regulatory credibility.

The 8-Layer Typological Risk Traceability Model outlined in this article offers a practical structure for linking typological intelligence to detection logic (which remains solution-dependent). By distinguishing between rules and scenarios—and embedding traceability from source to control—REs can build TM environments that are not just compliant, but explainable, targeted, and risk-aligned.

This approach isn't about chasing every typology. It's about detecting the right risks—and being able to demonstrate how and why those risks are being monitored. In a landscape of rising expectations and growing complexity, traceability is foundational.

When paired with effective feedback mechanisms, this structured model can enable sharper alerts, stronger governance, and better audit outcomes. Most critically, it shifts detection away from reactive box-ticking toward a proactive, risk-informed capability.

Further research is needed into how TM frameworks can remain agile in the face of evolving FC threats while maintaining strong traceability and governance. There is often a perceived trade-off between the flexibility of a detection environment—capable of rapidly adapting to new threats—and the stability of a structured framework, which may offer superior traceability but require more effort to engineer and maintain. However, agility is not determined by governance overhead and typological risk traceability alone. In practice, it is shaped by a broader set of organisational and technical factors—such as data quality, analytical tooling, environment replication and access, resourcing and skillsets, and vendor constraints. Many of these factors can be shaped by detective strategy and architectural choices and must be factored deliberately into design if responsiveness is to be preserved.

References

[1] AUSTRAC. (2024). Regulatory Priorities 2024. Available at: https://www.austrac.gov.au/sites/default/files/2023-12/AUSTRAC%20Regulatory%20Priorities%202024.pdf [Accessed 2025-04-01].

[2] David Coppin. (2024). Transaction Monitoring Effectiveness — part 1: the necessity of feedback mechanisms. Available at: https://www.linkedin.com/pulse/transaction-monitoring-effectiveness-part-1-necessity-david-coppin-n5khc [Accessed 2025-05-23].

[3] David Coppin. (2025). Transaction Monitoring Effectiveness — part 2: risk assessment sensitivity and continuous rule monitoring. Available at: https://www.linkedin.com/pulse/transaction-monitoring-effectiveness-part-2-risk-rule-david-coppin-8xqzc [Accessed 2025-05-23].

[4] David Coppin. (2025). Transaction Monitoring Effectiveness — part 3: quality frameworks, controls testing and validation. Available at: https://www.linkedin.com/pulse/transaction-monitoring-effectiveness-part-3-quality-controls-coppin-7jz9e [Accessed 2025-05-23].

[5] Financial Action Task Force (FATF). (2012; updated Feb 2024). FATF Recommendation 1: Assessing Risks and Applying a Risk-Based Approach. Available at: https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/FATF%20Recommendations%202012.pdf.coredownload.inline.pdf [Accessed 2025-05-23].

[6] FATF. (2021). Guidance for a Risk-Based Approach to AML/CFT Supervision. Available at: https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Guidance-Risk-Based-Supervision.pdf.coredownload.inline.pdf [Accessed 2025-05-23].

[7] Hunter, S., Johnson, T., Chai, A., de Koker, L. (2025). Navigating the Nexus of Financial Crime Prevention and Financial Inclusion in the Age of Technology and FinTech. Available at: https://www.adb.org/sites/default/files/publication/1054521/adbi-navigating-nexus-financial-crime-prevention-and-financial-inclusion-age-technology-and-fintech.pdf [Accessed 2025-08-08].

[8] AUSTRAC. (2025). Transaction Monitoring. Retrieved from https://www.austrac.gov.au/business/core-guidance/amlctf-programs/transaction-monitoring [Accessed 2025-05-23].

[9] AUSTRAC. (2023). Federal Court makes ruling in Crown matter. Retrieved from https://www.austrac.gov.au/news-and-media/media-release/federal-court-makes-ruling-crown-matter [Accessed 2025-05-23].

[10] AUSTRAC - Fintel Alliance. (2019). Illegal Phoenix Activity Indicators Labour Hire/Payroll Services. Retrieved from https://www.austrac.gov.au/sites/default/files/2019-10/Fintel%20Alliance%20phoenix%20activity%20report.pdf [Accessed 2025-05-25].

GRIFFITH UNIVERSITY

CRICOS: 00233E | TEQSA: PRV12076

griffith.edu.au/academy-financial-crime

MAKE IT MATTER