Unit21.ai

Catalogue

The AI Agent Prompt Catalogue for AML Investigations

A companion to How to Write AI Agent Prompts for AML Investigations. 40+ ready-to-adapt prompts for Unit21's Agentic Task Builder, organized by typology and built on the same framework: the number, the time frame, the data fields, and what the write-up should contain.

Before you copy a prompt

Every prompt in this catalogue is a starting point, not a finished product. Swap in your own thresholds, your own lookback windows, your own field names.

A prompt that says "flag deposits over $9,000" is only useful if $9,000 is actually the number your program cares about.

Two builders, same underlying discipline:

Online Search

Online Search tasks go out to the open web: adverse media, sanctions lists, corporate registries, litigation records. There's no data-gathering layer to configure; the work is telling the agent exactly what to look for and what would count as a hit.

Data Analysis

Data Analysis tasks work on your own data inside Unit21, across three layers. Data gathering is deterministic, since your instructions become a SQL query, so this is where you can afford to be broad. Data summary is where a language model reads what came back and describes what matters, so this is where specificity earns its keep. The narrative stitches everything together across tasks into a filing-ready write-up, and it's the layer with the least room for ambiguity, since it's the one a regulator eventually reads.

Four things separate a prompt that works from one that produces something reasonable-sounding and useless:

1
A number, not a word like "large" or "suspicious." $9,000 is a number. "High risk" is a guess about what you meant.

2
A time frame, not "recent." Ninety days is a window; "recent" isn't.

3
The data fields, stated generously. An extra field is a column nobody reads. A missing one is a finding that can't exist.

4
What the write-up should contain: what it shows, grouped how, ordered by what. This is the instruction most people skip, and the one the builder has the hardest time inferring on its own.

The prompts on the following pages lead with these where they apply. Bracketed terms like [entity name] or [90 days] are yours to fill in.

Part 1

Online Search Prompts

Adverse media & reputational risk

Search for adverse media coverage of [entity name] and any parent or subsidiary companies from the past [5 years]; return a clear risk verdict with source links for each finding.

Search for negative news tied to [merchant name], including consumer complaints, publicly reported chargebacks, or FTC/BBB actions filed in the past [3 years].

Check whether [individual name] has been named in fraud, embezzlement, or financial crime reporting; summarize each allegation with the date, the source, and whether it was ever resolved or dismissed.

Search for adverse media on this counterparty's beneficial owners individually, not just the company name, and flag anything found under variant spellings or transliterations of the name.

Sanctions, PEP & watchlist exposure

Confirm whether [entity name] or its listed directors appear on OFAC, UN, EU, or UK sanctions lists; cite the exact list, the listing date, and the date the search was run.

Search for public reporting connecting [entity] to a sanctioned jurisdiction, including indirect ties through ownership, supply chain, or subsidiary relationships in the past [24 months].

Determine whether [individual] holds or has held a senior political, military, or state-owned enterprise position that would classify them as a PEP, and identify the jurisdiction and the dates they held it.

Check whether this wallet address or exchange account has been publicly linked to a sanctioned entity, ransomware operation, or darknet market; cite the source and the date of the report.

Digital footprint & identity verification

Check whether the phone number, email domain, or physical address tied to this account has been reported elsewhere in connection with fraud in the past [12 months].

Search for a verifiable digital footprint for [individual name]: professional history, social profiles, public records, and note whether it's consistent with the identity on file, or absent prior to [account opening date].

Search for signs this identity may be synthetic: no history prior to a recent date, biographical details that conflict across sources, or a reverse-image match to a stock or stolen photo.

Beneficial ownership & corporate structure

Research the corporate registration and beneficial ownership structure of [company name] using public business registries; flag any offshore or high-risk jurisdiction ties and the percentage of ownership each party holds, where disclosed.

Search for evidence that [business name] is a shell company: no verifiable physical presence, no employees on professional networking sites, no digital footprint predating [the account opening date].

Identify whether [company] shares a registered address, director, or agent with any other entity in our case history over the past [12 months].

Business & merchant legitimacy

Verify whether [business name]'s website, business registration, and stated industry (MCC or SIC code) are consistent; flag any mismatch and describe what it suggests.

Check whether this website domain was registered within the past [6 months], uses a disposable or anonymized hosting provider, or closely mimics a known legitimate brand's name or design.

Search for public complaints, refund disputes, or "is this a scam" threads referencing this business or website, and summarize how many were found and how recent they are.

Litigation, regulatory action & enforcement history

Search court records and regulatory enforcement databases for civil or criminal litigation involving [entity/individual] related to financial crime, filed within the past [7 years].

Check whether this company has faced regulatory consent orders, fines, or license revocations from financial regulators in any jurisdiction, and list the regulator, the date, and the outcome.

Search for bankruptcy filings, judgments, or liens against this entity in the past [5 years] that might explain unusual account activity.

Part 2

Data Analysis Prompts

Each set below shows the gathering instruction and the summary instruction as two lines: the split the builder makes automatically, and the one you can make yourself under Add data gathering instructions for tighter control.

Structuring & threshold avoidance

Gather

Pull all deposits and withdrawals on this entity's linked accounts over the past [90 days], including amount, date, branch/channel, and account ID. Summarize: Flag any transactions between [$9,000 and $10,000], note whether they cluster around the same day or across linked accounts, and state the total count and dollar sum of flagged transactions.

Analyze this account's transactions for structuring across multiple linked accounts or branches on the same calendar day, using a [$9,000] threshold and a [90-day] window.

Identify whether this customer's transaction sizes changed noticeably in the [30 days] following a deposit over [$10,000], consistent with layering behavior.

Review whether this entity's pattern matches smurfing: pull all deposits under [$3,000] across linked accounts in the past [60 days] and flag any that route to a single common beneficiary.

Money mule & layering networks

Gather

Pull all inbound and outbound transactions on this account for the past [6 months], including counterparty name, amount, date, and account tenure at time of transaction. Summarize: Flag rapid pass-through activity, meaning funds disbursed within [48 hours] of receipt, and list any counterparty that also appears on another open or closed case in the same period.

Review the transaction network around this entity and list which counterparties also appear in other cases flagged in the past [90 days], ranked by how many shared counterparties they have.

Identify whether funds received in this alert were layered through two or more intermediary accounts before reaching a final destination, and map the path with dates and amounts at each hop.

Summarize this entity's inbound and outbound counterparties over the past [6 months] and flag any that match entities named in a previously filed SAR.

Account takeover & unauthorized access

Gather

Pull login events, device IDs, IP geolocation, and transaction activity for this account across the [48 hours] surrounding the flagged transaction. Summarize: Note any change in device, location, or beneficiary list within that window, and state whether the pattern is consistent or inconsistent with the account holder's prior [90 days] of activity.

Compare this customer's transaction pattern in the [30 days] before and after the reported access change; summarize whether behavior is consistent with the account holder.

Analyze whether a beneficiary added within [24 hours] of this transaction has any prior history on the account, and flag if it was added and used within the same session.

Card testing & promo/bonus abuse

Gather

Pull all authorization attempts from this device or IP over the past [7 days], including amount, approval/decline status, and merchant category code. Summarize: Flag a pattern of small-dollar authorizations (under [$5]) followed by a decline rate above [70%], and state the number of distinct card numbers attempted.

Analyze new-account signups from this referral source for shared device fingerprints in the past [7 days], and summarize how many converted a promo bonus without further account activity.

Summarize transaction velocity for this merchant category code across accounts opened in the past [7 days], and flag any spike above [3x] the trailing 30-day average.

Dormant account reactivation

Flag whether this account was inactive for over [6 months] before this transaction, and summarize what changed at reactivation: device, location, or beneficiary.

Identify all accounts that were dormant for [6+ months] and reactivated in the past [30 days] with a transaction to the same beneficiary; list the accounts and the shared beneficiary.

Case narrative & SAR/STR drafting support

Narrative Instruction

Draft a case narrative in our standard SAR format, covering entity history, the transaction pattern identified in the data summary tasks, and the analyst's rationale for escalation. Order sections as: who, what, when, where, why. Do not draw conclusions the underlying tasks didn't support.

Summarize this case's full evidence trail, entities, transactions, and prior alerts, into a filing-ready narrative with a recommended disposition, citing which task produced each fact.

Review this case's supporting documentation and list any gaps that would need to be filled before a SAR could be filed defensibly.

Summarize the last [three] alerts on this entity into a single consolidated narrative for the analyst reviewing the escalated case, in chronological order.

Customer risk rating & EDD tiering

Review this customer's transaction history, product usage, and geographic exposure over the past [12 months], and state whether their current risk tier still reflects observed behavior.

Identify customers whose transaction volume changed by more than [50%] in the past quarter but whose risk rating hasn't been updated in the same period.

Summarize the EDD-relevant facts for this customer: source of funds, jurisdiction exposure, and adverse findings from the past [24 months] of reviews.

Analyst QA & consistency review

Review the last [20] alerts closed by this analyst and flag any where the disposition doesn't match the evidence documented in the notes.

Compare how similar alerts (same typology, same threshold triggers) were dispositioned across the team in the past [30 days], and flag inconsistent escalation decisions.

Summarize whether this closed alert's investigation notes followed each step of our SOP checklist, and name the first missed step, if any.

A note on testing before you turn a task on

A prompt can be exactly right and still fail in production if you only validate it against one alert. Pick alerts that differ from each other on purpose: the one with a single flagged entity and the one with a dozen, the thin file and the thick one. The builder writes against whatever evidence it's given; a sample of one produces something shaped around a coincidence in that alert, not the range you'll actually see.

Where a prompt hits its limit

On tasks that route to a SAR filing agent, your instructions reach the narrative as an additional layer, never a replacement for FinCEN's requirements. No prompt gets the agent to skip a required section, invent a fact, draw a legal conclusion, or speculate past the evidence. The regulatory floor holds regardless of what's typed above it. Your instructions decide what happens above that floor.

And more detail isn't automatically better. A prompt narrow enough that almost nothing matches produces a clean, technically correct result that says nothing, a task that looks like it's working because it never has anything to flag. If a task hasn't surfaced a real finding in a while, that's worth a second look, not a passing grade.