AIRMDR

53 SOC LLM Prompts You Can Use Right Now

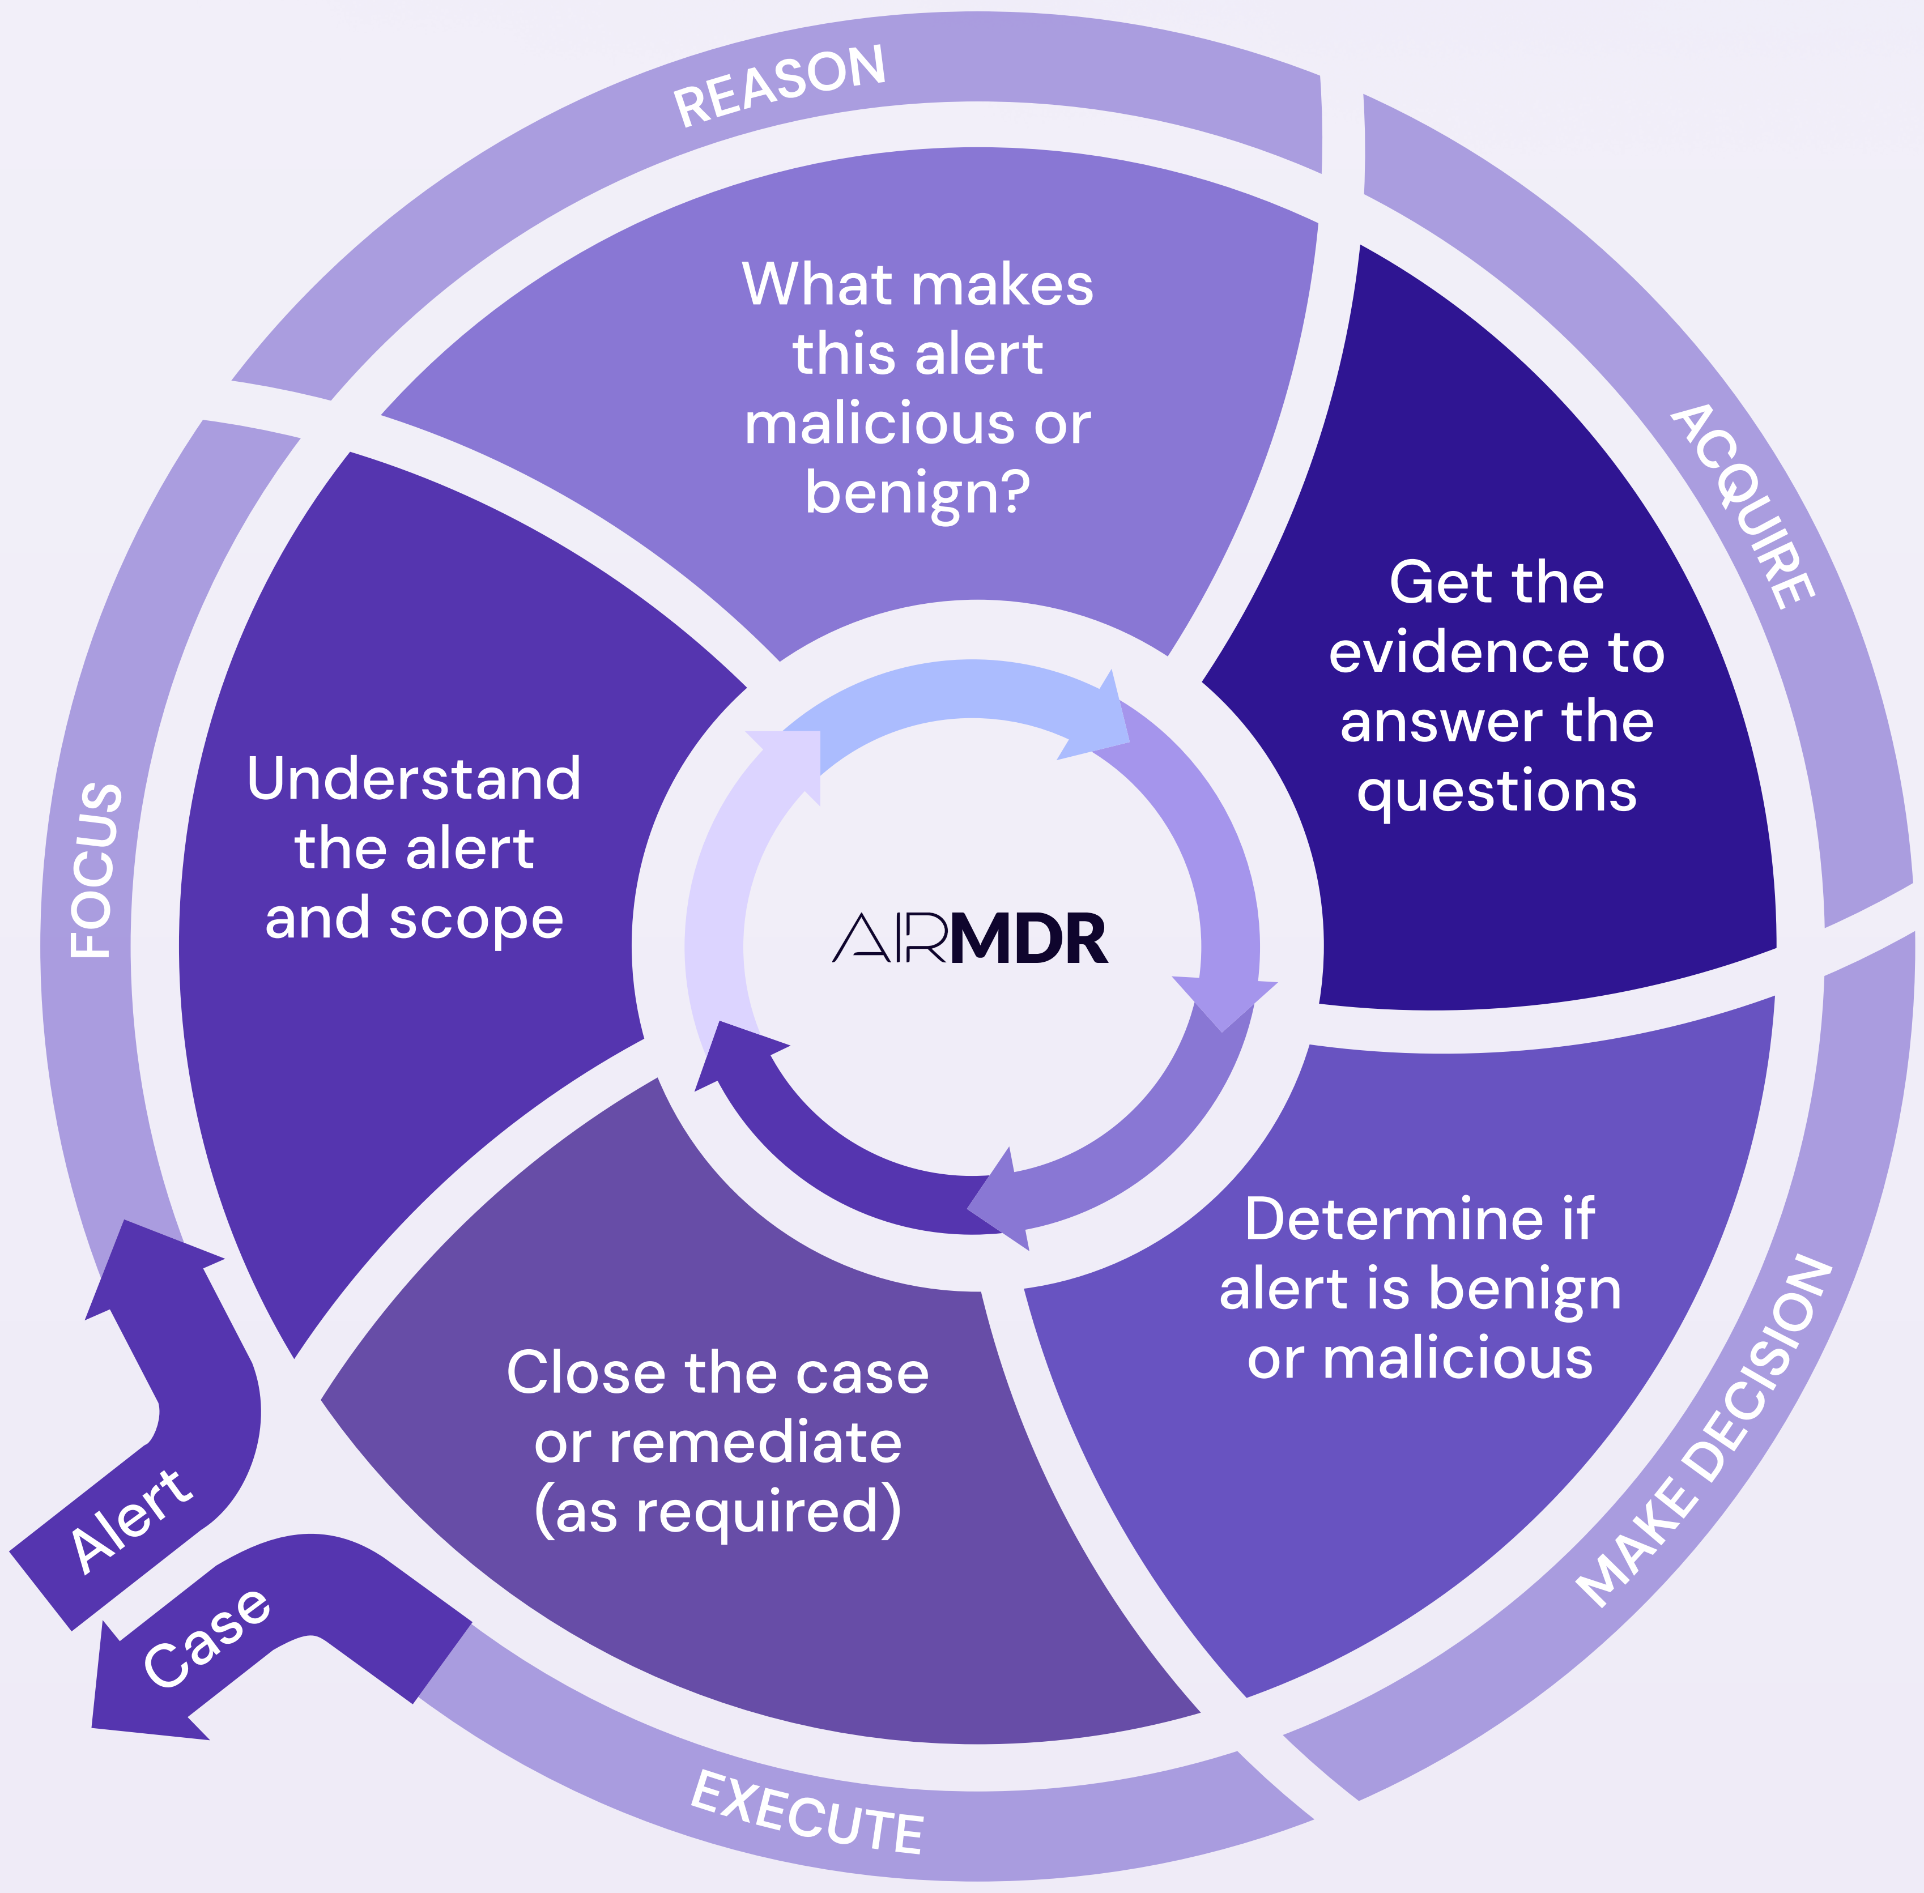
REASON

What makes this alert malicious or benign?

ACQUIRE

Get the evidence to answer the questions

MAKE DECISION

Determine if alert is benign or malicious

EXECUTE

Close the case or remediate (as required)

FOCUS

Understand the alert and scope

AIRMDR

Alert

Case

FOCUS

- Break down the alert into: detection logic, triggering event, entities involved, and potential security risks. Explain it in analyst-friendly terms.
- What behavior is this detection designed to identify?
- What systems, identities, IPs, or time ranges should be in scope for investigation?
- What are the IOCs that should be investigated in this alert?
- What attacker techniques or tactics (MITRE ATT&CK) might this alert map to?
- What conditions must be true for this alert to represent actual risk?
- What business process or legitimate activity could resemble this behavior?
- What is the potential impact if this alert is confirmed malicious?

REASON

- What are reasons this alert might be a false positive?
- What are the 5 most important investigative questions I need to answer to determine if this is malicious?
- Provide 3 malicious hypotheses and 3 benign explanations that could explain this alert.
- What questions should i answer to develop the highest confidence I am making the correct decision on this case?
- What evidence most strongly supports malicious intent?
- What evidence most strongly supports a benign explanation?
- Are the observed indicators causally related, or could they be coincidental?
- What would I expect to see next if this activity is truly malicious?
- What alternative explanations would explain this alert?

ANYTIME PROMPTS

The "Adversarial Pre-Mortem"

"Assume my current write up is incorrect. Generate a plausible scenario where the evidence I found is actually misleading (e.g., a sophisticated decoy or an admin using a tool improperly). Where is the biggest gap in my logic?"

The "Executive Risk Translation"

"Translate the current write up into a 'Business Risk Brief' for the CFO. Avoid security jargon. Focus exclusively on: Operational Downtime Risk, Regulatory Exposure (GDPR/CCPA), and Estimated Financial Impact of inaction."

The "MITRE ATT&CK Alignment"

"Map the observed behaviors in this case to specific MITRE ATT&CK TTPs. Identify which phase of the Kill Chain we disrupted and which phases (if any) might have occurred silently prior to detection."

ACQUIRE

- What evidence would I need to collect to confidently determine if this alert represents true malicious activity?
- What data sources would provide additional information that is not already available but necessary to investigate?
- Write a [Splunk / Sentinel / KQL / Sumo / etc.] query to find related activity for this user/IP/host in the last 24 hours.
- Based on the questions i need to answer, identify gaps where more data may be needed and where I can get that data..
- What pivots should I perform on this IP/domain/user/host to uncover related suspicious activity?
- What external enrichment (WHOIS, VirusTotal, ASN, geo, sandbox, etc.) would help validate the risk of this indicator?
- Based on what I've collected so far, what critical evidence am I missing?
- What historical activity exists for this entity over the past 7–30 days?
- What comparisons would help determine if this behavior is anomalous?
- What endpoint, identity, network, or cloud telemetry should I correlate together to validate this activity?

MAKE DECISION

- Based on the investigation summary below, should this case be escalated? Provide justification suitable for management review.
- Does the available evidence sufficiently support closing this as a false positive? What residual risks remain?
- What is my confidence level in this determination, and what factors reduce that confidence?
- Challenge my conclusion. What weaknesses exist in my reasoning?
- If I close this case incorrectly, what is the potential business impact?
- Is there evidence of impact/attack success, or only evidence of attempt?
- What threshold would need to be crossed to justify escalation?
- If I escalate this, what action/response should I expect from others?
- Does this case require broader environment review beyond the original scope?
- Is additional expert review (L3, IR, engineering) warranted before disposition?
- Perform a peer review of my analysis and make an escalation recommendation. Explain why.
- Identify missing information and assumptions I made that might make my conclusion incorrect.

EXECUTE

- What actions are required to isolate this attacker?
- Which containment actions might have unintended consequences?
- If confirmed malicious, outline prioritized remediation steps (containment → eradication → recovery).
- Based on this case, who do I need to notify about this event and what should I say to them?
- What are the tactical actions that should be prioritized first?
- What are the strategic adjustments i need to make to improve my security posture?
- What additional monitoring should be enacted or modified given the conclusion in this event.
- What controls failed or were bypassed that allowed this activity to occur?
- Should detection logic be updated based on this investigation? If so, how?
- What documentation updates (playbooks, knowledge base, training) are required following this case?

AIRMDR

FROM ALERT TO ACTION

53 Prompts You Can Use Now

Scan or click for
latest update