Security Copilot Agent — Scenario

Entity Guard Investigator

Version: v1.2 (Defender-only)
Owner: adaQuest
Last Updated: 2025-09-17

1) Purpose of the Agent (Customer benefit)

The Entity Guard Investigator determines whether identities and related entities referenced by a Microsoft Defender incident are compromised. Starting from a Defender IncidentId, the agent extracts entities (users, email addresses & message IDs, devices, IPs, URLs/domains, file hashes), enriches them with Microsoft Entra identity signals, Defender device posture, Threat Intelligence (DTI) reputation, and optionally Intune compliance. It then produces a human-readable verdict per entity with concise evidence and prioritized actions.

Customer value

- **Accelerated triage:** Automates enrichment across Microsoft security products and focuses strictly on entities present in the incident.
- **Improved accuracy in email cases:** Message-level reasoning (EmailEvents, URL/Attachment artifacts, SPF/DKIM/DMARC results, ZAP, sender history) prevents false "benign" verdicts based only on domain reputation.
- **Consistent, explainable scoring:** Transparent per-entity risk model with thresholds (OK / Suspicious / Compromised).
- **Noise reduction:** Incident-scoped entity processing with a MaxEntities cap avoids overload.
- **Actionable outcomes:** Provides clear, prioritized next steps (e.g., isolate device, reset identity, block IoCs, review CA/MFA).
- **Easy consumption:** Final output is concise markdown only (no raw JSON exposed).

2) Functional Design

High-level steps the agent performs to achieve its purpose:

Inputs (minimal):

- IncidentId (Defender): required
- Optional: LookbackDays (default 30), MinTIConfidence (default 60), UseIntune (default true), MaxEntities (default 30)

High-level flow:

1. **Entity extraction (Defender)**
    - Retrieve incident: GetDefenderIncidentReport.
    - Extract AlertIds (GPT utility).
    - Expand entities from AlertEvidence (local KQL): users, emails, message IDs, devices, IPs, URLs, domains, file hashes.
    - Normalize & deduplicate (GPT) → authoritative scope.
2. **Email context (when present)**
    - From AlertIds, collect EmailEvents, EmailUrlInfo, EmailAttachmentInfo (local KQL).
    - Aggregate sender reputation within the lookback window (local KQL).
3. **Identity signals (Entra)**
    - GetEntraSignInLogsV1, GetEntraRiskyUsers, GetEntraAuditLogs, GetEntraServicePrincipalDetails.
4. **Device posture (Defender / Intune)**
    - GetDefenderDeviceSummary, optionally GetDefenderIdentitySummary.
    - Local KQL fallbacks: GetDeviceContext, GetUserDeviceCoOccurrence.
    - If UseIntune=true: GetIntuneDevices, GetPoliciesPerDevice, DescribeIntunePolicy.
5. **Threat Intelligence (DTI)**
    - Build indicator sets from entities & email artifacts (IPs, URLs, domains, sha256).
    - GetIndicatorsByIndicators, GetReputationsForIndicators, GetSummaryForIndicators.
6. **Consolidation (GPT)**
    - ConsolidateInvestigationSignals merges signals by entity (type:id).
    - Email-specific signals: classified Phish/Spam/Malware, SPF/DKIM/DMARC auth failures, ZAP/post-delivery actions, malicious URLs/attachments, sender phish history.
    - Clear distinction between missing sources (tool unavailable) and low signal (empty results).
7. **Risk scoring (GPT)**
    - BuildRiskScoringJson applies weighted scoring per entity.

- Verdicts: 0–29 OK, 30–69 Suspicious, 70–100 Compromised (requires ≥2 strong signals).

8. **Final output (human-readable)**

    - SummarizeEntityInvestigation produces markdown only.
    - Includes: Executive verdict per entity, evidence highlights, recommended actions, coverage notes.
    - Never returns raw JSON to end users.

Guardrails & behavior:

- If incident is missing/inaccessible → stop gracefully with user-friendly message.
- Only incident entities are processed (owners/assignees ignored unless also entities).
- Entity processing capped by MaxEntities.
- This build is Defender-only (no Sentinel dependencies).

3) Triggers for Agent Activation

Manual (v1): Analysts run the agent on demand by providing a Defender IncidentId. Optional inputs: LookbackDays (default 30), MinTIConfidence (default 60), UseIntune (default true), MaxEntities (default 30).

Optional/Future: Scheduled sweeps (e.g., every 30–60 minutes) using Defender incident listing; event-driven activation on incident creation/update via product webhooks.

4) Plugins or Data Signals

Global plugins/data sources (required):

- Microsoft Entra — GetEntraSignInLogsV1, GetEntraRiskyUsers, GetEntraUserDetailsV1, GetEntraAuditLogs, CAPolicyAnomalyDetectionSkill, GetEntraServicePrincipalDetails
- Microsoft 365 Defender — GetDefenderIncidentReport, GetDefenderIncidents, GetDefenderDeviceSummary, GetDefenderIdentitySummary, GetFileAnalysis (optional)
- Microsoft Threat Intelligence (DTI) — GetIndicatorsByIndicators, GetReputationsForIndicators, GetSummaryForIndicators, FindThreatIntelligence
- Microsoft Intune — GetIntuneDevices, GetPoliciesPerDevice, DescribeIntunePolicy, GetIntuneDeviceGroupMemberships, GetIntuneDeviceDiscoveredOrManagedApplication, GetAppOrPolicyDeviceGroupTargeting, GetPolicySettingUsage

Local KQL skills (Defender-only):

- FetchEntitiesFromDefenderAlerts — expand entities via AlertEvidence using AlertIds
- GetEmailContextFromDefenderAlerts — EmailEvents/EmailUrlInfo/EmailAttachmentInfo by NetworkMessageId

- SummarizeEmailSenderRisk — sender risk aggregation within lookback window
- GetDeviceContext, GetUserDeviceCoOccurrence — device/user context fallbacks

5) Customer Onboarding/Deployment

Prerequisites

Reader-level permissions for:

- **Microsoft Defender XDR** (to access incidents, AlertEvidence, Email tables, device context)
- **Microsoft Entra ID** (to access sign-in logs, risky users, audit logs)
- **Microsoft Threat Intelligence (DTI)** (for indicator reputation and enrichment)
- **Microsoft Intune** (optional — only if you want device compliance data).
- Advanced Hunting access in **Defender** (must be able to query AlertEvidence, EmailEvents, EmailUrlInfo, EmailAttachmentInfo).

Enable the agent

2) In the **Security Copilot Agents gallery**, locate **Entity Guard Investigator** and enable it.
3) When prompted, **consent** to the required plugins (Defender, Entra, DTI, Intune optional).

How to run

- Provide a **Defender IncidentId** when starting the agent.

- Optionally adjust investigation parameters:
    - LookbackDays (default 30)
    - MinTIConfidence (default 60)
    - UseIntune (true/false, default true)
    - MaxEntities (default 30).

What to expect

- The agent analyzes all entities in the incident: users, email addresses, message IDs, service principals, devices, IPs, URLs, domains, and file hashes.
- Each entity is enriched with context (Entra sign-ins, risky users, device posture, threat intel, compliance).
- Output is a clear markdown report with:
  • Executive verdict per entity (OK, Suspicious, Compromised)
  • Evidence highlights and reasoning
  • Prioritized recommended actions
  • Data Coverage Notes (missing sources vs. low signal).
- No raw JSON is shown to analysts.

Operational guidance

- Only entities present in the incident are analyzed; owners/assignees are ignored unless they also appear as entities.
- Email analysis is message-centric, so malicious evidence tied to a message will outweigh benign domain reputation.
- Use the output as an accelerator for triage — confirm key findings and apply recommended containment steps.