Security Copilot Agent — Scenario

Sentinel Data Leak

1) Purpose of the Agent (Customer benefit)

The *Sentinel Data Leak* agent performs a comprehensive investigation of a Microsoft Sentinel incident, correlating Defender XDR and Purview data to determine whether the incident is a true data exfiltration event or a false positive. It returns a single Markdown report with evidence, correlation insights, analyst reasoning, and executive recommendations.

Customer value

- Automates cross-signal correlation centered on a Sentinel incident.
- Validates exfiltration signals against corroborating device/user/file context and Purview data security alerts.
- One executive-ready report with verdict, evidence highlights, and prioritized recommendations.
- Uses the existing Sentinel + Purview signals.

2) Functional Design

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

Parameter	Required	Description
Incident	■	Sentinel incident identifier (GUID, integer Incident Number, or incident URL).

1. **Incident context**
    a. Retrieve Sentinel incident details and entities using GetIncident, GetIncidentEntities, and (when useful) GetSentinelIncidents.
2. **Enrichment & correlation**
    a. Enrich with device and file context: EnrichIncidentWithDeviceContext, EnrichIncidentWithFileContext.
    b. Hunt related activity for involved users, IPs, or hostnames: FindUserIpOrHostnameAccessRecords, FindUserScriptDownloads.
3. **Purview (data security) analysis**
    a. Zoom and summarize data-risk/user-risk and DLP/IRM signals:

        i. ZoomIntoPurviewDataRisk, GetDataRiskSummary, ZoomIntoPurviewUserRisk, GetUserRiskSummary,
        ii. SummarizePurviewAlert, TriagePurviewAlerts, DataSecurityAnalytics.
4. **Reasoning & verdict**
    a. Correlate all findings to assess whether sensitive data was accessed, shared, or exfiltrated and whether user/device behavior supports a **true exfiltration** vs **false positive** conclusion.
5. **Output**
    a. Produce a **Markdown** report: executive summary (verdict), incident context, correlated evidence tables, analyst reasoning, and prioritized recommendations; include coverage/limitations if any data is unavailable.

Behavior & guardrails

- **Manual execution only** (on-demand).
- **No additional parameters/thresholds** required beyond Incident.
- **Markdown-only output** (no raw JSON).
- **Graceful handling** of missing signals (report what was unavailable).

3) Triggers for Agent Activation

Trigger: Default

Type: Manual (no schedule/recurrence)

ProcessSkill: SentinelDataLeak.SentinelDataLeakInvestigator

Invocation: Analyst provides Incident (GUID / number / URL).

4) Plugins or Data Signals

Required Skillsets.

- Fusion
- Sentinel
- Purview
- MCP.Sentinel

Child skills used by the Process skill:

- GetIncidentEntities
- GetIncident
- list_sentinel_workspaces
- GetSentinelIncidents
- EnrichIncidentWithDeviceContext
- EnrichIncidentWithFileContext

- FindUserIpOrHostnameAccessRecords
- ZoomIntoPurviewDataRisk
- GetDataRiskSummary
- FindUserScriptDownloads
- GetUserRiskSummary
- SummarizePurviewAlert
- TriagePurviewAlerts
- DataSecurityAnalytics
- ZoomIntoPurviewUserRisk

5) Customer Onboarding/Deployment

Prerequisites

- Access to Microsoft Sentinel incidents and alerts.
- Access to Purview data security/alerts as required by the listed skills.
- Fusion/Sentinel/Purview skillsets enabled in Security Copilot.

Enable the agent

- Install/enable **Sentinel Data Leak** in Security Copilot; consent to the three required skillsets above.

How to run

- Run manually with:
  Run Sentinel Data Leak for Incident = "<GUID | IncidentNumber | IncidentURL>"

Expected output

- One **Markdown** report with: verdict (true exfiltration / false positive / inconclusive), incident context, evidence correlation (Defender/Purview), analyst reasoning, and prioritized recommendations; include any data-coverage notes.