Operation Ridgeline

Case Study

SYNTHETIC CASE STUDY — SUPPLY CHAIN THREAT ACTOR INVESTIGATION

The following is a fictionalized case study illustrating Summit Cognitive platform capabilities. All entities, organizations, and events are synthetic. Any resemblance to real persons, organizations, or incidents is coincidental.

SITUATION

A defense prime contractor's security team suspects that a series of anomalous network events across three supplier organizations over a 14-month period may represent coordinated reconnaissance activity by a nation-state threat actor. The events individually cleared routine investigation — each appeared unrelated to the others. The team has 340 entities across seven data sources: network logs, vendor contracts, procurement records, open-source corporate filings, personnel records, threat intelligence feeds, and unstructured analyst notes.

The problem: three analysts have spent six weeks building partial link charts in i2 Analyst's Notebook. The picture is incomplete. The timeline is inconsistent. No one can demonstrate whether the events are connected or coincidental, and the brief to leadership is due in 72 hours.

SUMMIT COGNITIVE ENGAGEMENT

Hour 1 — Data ingest & entity resolution

All seven data sources loaded via the STIX 2.1 importer and the CSV batch loader. Entity resolution identified 47 duplicate entities across sources — "Meridian Systems LLC," "Meridian Systems," and "Meridian Sys." collapsed to a single canonical node with three source provenance records. Graph populated: 293 confirmed entities, 812 relationships.

Hour 2 — Temporal reconstruction

Time slider set to 14 months prior. Network assembled event by event. At month 4, a previously unremarkable vendor contract renewal between Supplier A and an intermediary became visible as temporally proximate to a personnel change at Supplier B involving an individual with prior employment at a known front company. The connection was invisible in static link charts because it required correlating three separate data sources across a 12-day window.

Hour 3 — AI Copilot investigation

"What is the strongest evidentiary path connecting the Meridian Systems entity to the anomalous network events at Supplier C?" — answered in 23 seconds. The Copilot returned a 4-hop path: Meridian → shared infrastructure → shell entity → Supplier C network segment, with why-paths tracing each edge to its source document. Support score: 0.84 (high confidence). Each cited relationship verifiable in the source data.

Hour 5 — Hypothesis simulation

Narrative Simulation Engine ran 5 Monte Carlo rollouts on a hypothesis: "If Meridian is acting as a procurement intermediary for the threat actor, what acquisitions would we expect to see in the next 60 days?" Result: 80% confidence (4/5 rollouts) that a bid on a specific RFP category would appear. Marked as HYPOTHESIS in the graph, not confirmation. Leadership brief explicitly labeled the simulation layer as predictive inference.

Hour 7 — Case note export

Full investigation exported as a structured case note: the question asked, the answer, the confidence level, the why-paths with source citations, the hypothesis scenarios with probability distributions, and the temporal reconstruction sequence. The brief wrote itself.

OUTCOME

Time to actionable picture	7 hours (vs. 6 weeks of partial manual work)
Entities resolved	293 confirmed (47 duplicates collapsed automatically)
Key connection surfaced	4-hop path invisible in static link charts
Simulation accuracy	Hypothesis scenario validated at 80% confidence across 5 rollouts
Audit trail	Full tamper-evident record of every query, path, and hypothesis clearly labeled
Case note	Complete citation-ready brief generated automatically from investigation session

"The connection was in the data the whole time. Summit found it in seven hours. We had been looking at it for six weeks."

— Fictionalized composite of analyst feedback from platform testing