Orchestration and subagents in an agent harness for security work

Delegation brief (one per subagent call)

Objective: one sentence, one subject.
Output format: the fields the parent will read, as a schema.
Tools and sources: the named tools this subagent may call; nothing else is in its session.
Boundaries: what it must not do (writes, tickets, queries outside the named sources).
Effort cap: maximum tool calls and maximum tokens for this call.
Escalation: what it returns when evidence is missing or contradictory instead of guessing.

Example
Objective: Find every sign-in to host app-17 from an address outside the corporate ranges between 14 Nov 2026 00:00 and 19 Nov 2026 00:00 UTC.
Output format: JSON list of {time, account, source_address, authentication_method, evidence_id}.
Tools and sources: siem_search on the auth index only.
Boundaries: read only; account lookups only for names that appear in results.
Effort cap: 10 tool calls, 60,000 tokens.
Escalation: if the auth index has no data for part of the window, return that gap as a finding.

Agentic AI in Cybersecurity, Orchestration and subagents in an agent harness for security work, edition 19 Nov 2026, https://agenticaiincybersecurity.com/agent-harness/orchestration-and-subagents/. License CC BY 4.0.