# Cypienta's AI Sample Use Cases

## What is Cypienta

Cypienta groups related events and deduplicates their redundant values, condensing data volumes by 80% and revealing context-rich storyboards with underlying root causes and coordinated attacks.

It is made of the following 5 layers defining its 5 core functionalities:

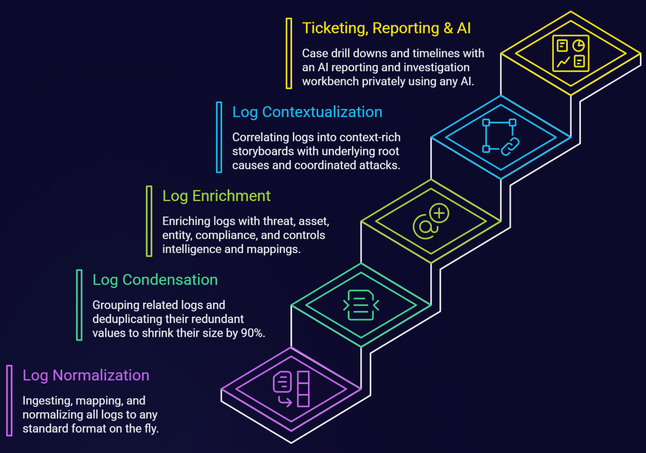


- Ingestion & Normalization Layer

  Interoperable data collection, normalization, lake, & warehouse layer on any underlying data store, such as AWS S3, Azure Blob, GCP, Snowflake, Databricks, Local storage, etc. Runs our unique models and any custom for data extractions, transformations, aggregation, sampling, trimming, filtering, analysis, and processing.

- Condensation Layer

  Lossless data condensation & optimization layer with our unique models that group related events, such as those resembling similar activities and deduplicates redundant fields within them

- Enrichment Layer

  Threat, asset & entity intelligence enrichment layer with our unique models and any custom enrichment logic for events, assets & entities.

- Contextualization & Investigation Layer

  Lossless data contextualization, correlation, consolidation & investigation layer with our unique models that group related events, such as those resembling similar activities, common root cause, or causal attack chains, and deduplicates redundant fields within them and any custom grouping, correlation or coalescing logic. The admin interface enables easy auditability, explain ability, tuning and orchestration of the models.

- Ticketing, Reporting & AI Layer

  Case Management interface where you can collaborate and view the groups and chains in interactive storyboards, timelines, drill-down workbenches, and investigation playbooks with incident detection, mitigation, & response advisories, threat actor & vector attributions, extracted IOCs & PIIs, incident diamond model reports, case briefings, and AI assisted workflows.

  Interoperable AI wrapper interface with our unique guard railing & augmentation modules that integrate with any AI such as Copilot, Chat GPT, and local models, and pass context & consolidated data with masked PIIs & secrets, and with extracted IOCs & threat intel, as well as enable any custom workflows or augmentation logic

All managed through Cypienta's Admin Interface
Administration interface with pipeline management allowing output data forwarding & transformation with unique modules that output data in all formats to all destinations, and model management allowing easy auditability, explainability, tuning and orchestration of all of our models. Before any use case happens, AI receives:

## What an AI model gets from Cypienta

- Normalized, condensed telemetry (token-efficient, cross-tool unified format).
- Asset, entity, and object intelligence enrichment from historical + user-defined intelligence tables.
- Threat enrichment with ATT&CK TTPs, NIST 800-53, VERIS, CSA CCM, CISA KEV, CRI and relevant controls in AWS, GCP, Azure, M365, and Intel vPro.
- Contextual groupings of telemetry (hundreds of logs reduced to one contextual grouping).
- Incidents with highlighted attack chains and causal “stories” linked together.
- User interactions on the Case management UI and settings via the AI Integration wrapper (PII masked, IOCs extracted, safe for local models).

That’s what makes the following use cases actually high signal.

# Use Cases

## Alert Handling & Triage

1. Incident, Ticket, & Case summarization
   Enabled by: Contextualization & Condensation.
2. Generating Threat Intelligence Reports & Attributions
   Enabled by: Normalization + Contextualization layer.
3. Cross-source context extraction
   Enabled by: Normalization.
4. Severity scoring using enriched asset criticality + KEV presence.
   Enabled by: Intelligence & Enrichment layer.
5. False positive pattern detection using grouping baselines.
   Enabled by: Contextual clustering + lake history.
6. Recommended mitigations advisories
   Enabled by: Intelligence & Enrichment layer.
7. Highlighting suspicious patterns.
   Enabled by: Consolidation deduplication logic.
8. Risk-based prioritization
   Enabled by: Asset & Entity Intelligence enrichment.

## Investigation Support

1. Natural language querying
   Enabled by: Interoperable Data Layer.
2. Suggesting questions, verdicts & tags for each ticket
   Enabled by: Data Foundation abstraction.
3. Timeline & Storyboard detailed reconstruction
   Enabled by: Contextualization layer.
4. Entity correlation using unified asset, object, & entity graph.
   Enabled by: Intelligence & Enrichment.
5. Pivot suggestions based on gathered evidence
   Enabled by: Story linking engine.
6. IOC enrichment summaries from extracted artifacts within groupings.
   Enabled by: Enrichment + extraction modules.
7. Threat intelligence summarization pre-linked to internal exposure.
   Enabled by: KEV + ATT&CK enrichment.
8. Malware attribution
   Enabled by: Asset correlation engine.
9. Recommending the best analyst for the investigation
   Enabled by: Consolidation layer.
10. Email header analysis enriched with entity ownership context.
    Enabled by: Asset/Entity intelligence.
11. Decoding obfuscated scripts pcaps, and payloads.
    Enabled by: AI Integration wrapper.
12. Command-line interpretation with ATT&CK mapping pre-applied.
    Enabled by: TTP enrichment.
13. Highlighting anomalous behavior using entity baselines.
    Enabled by: Historical clustering.

## Detection Engineering

1. Writing & tuning detections for every incident
   Enabled by: Interoperable Data Layer.
2. Converting rules across vendors via unified data model.
   Enabled by: Data Foundation abstraction.
3. Explaining & optimizing current detection logic
   Enabled by: Consolidation layer.
4. Rule tuning using false positive clusters
   Enabled by: Historical grouping.
5. Detection gap analysis using ATT&CK coverage enrichment.
   Enabled by: Intelligence layer.
6. Log schema understanding
   Enabled by: Normalization pipeline.
7. Generating test cases from historical stories.
   Enabled by: Contextual story engine.
8. Reviewing rule logic for the environment
   Enabled by: Story-linking engine.

## Threat Hunting

1. Generating hypotheses using asset exposure intelligence.
2. Translating hypotheses into cross-source normalized queries.
3. Summarizing hunt findings from contextual groupings.
4. Surfacing weak signals from aggregated clusters.
5. Suggesting pivots across entity relationships.

All enabled by: Contextualization + Asset Intelligence + Unified data layer.

## Incident Response

1. Suggesting recommendations from consolidated incident stories.
2. Executive summaries from grouped causal chains.
3. Step - by step technical deep dives with extracted IOCs containment
4. Response checklists using mapped TTPs.
5. Remediation suggestions based on KEV + asset criticality.
6. Identifying differential changes in assets, entities, & objects
7. Extracting indicators from contextualized artifacts.
8. Root cause summaries from story-level clustering.

Enabled by: Story engine + Intelligence enrichment + Case management UI.

## Automation & SOAR

1. Dynamically expanding playbooks aligned to current incidents.
2. Explaining playbook logic against consolidated flows.
3. Generating response scripts from normalized fields.
4. Updating scripts to match updating use cases
5. Auto-generating case notes, and playbooks from stories.
6. Suggesting next actions using story progression logic.

Enabled by: AI Integration wrapper + Contextualization engine.

## Knowledge & Training

1. Internal knowledge assistant grounded in historical grouped incidents.
2. Explaining concepts, provide rapid guidance & next steps
3. Coach Analysts during investigations
4. Runbook summarization tied to actual incident types.
5. Translating findings with preserved ATT&CK/NIST mapping.
6. Generating tabletop scenarios from historical attack chains.
7. Building simulations based on historical behavior.

Enabled by: Consolidated data lake + enrichment layer.

## Communication

1. Customer notifications from contextualized incidents.
2. Regulator summaries mapped to NIST 800-53 controls.
3. Translating reports safely with PII masked.
4. Automatic sensitive fields redaction before exposure.
5. Formatting findings from unified data outputs.

Enabled by: AI wrapper + Enrichment + Output pipeline.

## Governance & Metrics

1. Extracting metrics from contextual groupings.
2. Consistent categorization via enrichment tags.
3. Identifying recurring root causes from clustered stories.
4. Suggesting control improvements via ATT&CK coverage analysis.
5. Auditing alert handling consistency using case history.

Enabled by: Historical clustering + Explainability layer.

## Asset, Entity & Object Intelligence

1. Normalizing asset identifiers across vendors.
2. Building unified asset profiles from all telemetry.
3. Explaining asset criticality
4. Identifying ownership from entity tables.
5. Detecting risky service accounts via behavior clustering.
6. Surfacing privileged account exposure.
7. Identifying dormant high-risk accounts.
8. Mapping relationships across users, devices & apps
9. Flagging risky entity combinations.
10. Summarizing exposure posture across cloud environments.
11. Correlating assets across AWS, GCP, Azure, & M365.
12. Detecting shadow IT from grouped anomalies.
13. Identifying anomalous entity-to-entity communication.
14. Highlighting risky object access patterns.
15. Explaining risk concentration across entity graph.

Enabled by: Intelligence & Enrichment + Historical contextualization.

## Vulnerability Intelligence

1. Summarizing scan results mapped to assets.
2. Translating CVEs into contextual business risk.
3. Mapping vulnerabilities to affected groupings.
4. Prioritizing based on KEV + exploit signals.
5. Correlating active alerts with vulnerable systems.
6. Identifying exposed internet-facing vulnerable assets.
7. Explaining exploit techniques via ATT&CK mapping.
8. Aligning vulnerabilities to active campaigns.
9. Grouping vulnerabilities by systemic root cause.
10. Suggesting patch prioritization using asset intelligence.
11. Identifying compensating controls via NIST mapping.
12. Tracking remediation progress from consolidated cases.
13. Summarizing repeated patch failures.
14. Detecting vulnerability trends across historical clusters.
15. Highlighting assets repeatedly exposed across cycles.

Enabled by: Enrichment layer + Contextual consolidation + Unified data lake.

If you step back, what Cypienta is really doing here is:

> It turns fragmented telemetry into structured, condensed, enriched, contextualized stories.
>
> Then any AI model, cloud or local, operates on high-density contextual signal instead of raw noisy logs.

That’s the difference between AI being a chatbot and AI being operationally useful in a SOC.