Case Study: Aegis — Building a Persistent OT/ICS Cyber Threat Intelligence Agent on MindStone

How a persistent-identity AI agent became a full-time OT threat analyst — accumulating expertise, tracking threats across months, and delivering analyst-grade intelligence output every day

Agent name: Aegis
Platform: MindStone
Domain: OT/ICS Cybersecurity — Threat Intelligence
Deployment date: March 17, 2026
Operator: Clint Bodungen, Director of AI/ML Engineering & Cybersecurity Innovation, MorganFranklin Cyber
Primary output: Daily threat intelligence pulls, dashboard payloads, investigation tracking, executive advisories

Summary

OT/ICS threat intelligence is not a query. It is a discipline.

Understanding whether a new Rockwell Automation advisory represents a fire drill or a fire, whether a nation-state threat group has crossed from pre-positioning to active operations, or whether a KEV deadline puts a water utility's SCADA system at imminent risk — these are not questions with one-shot answers. They require a working analyst with context: who the threat actors are, what they've done before, what the asset landscape looks like, and what's changed since yesterday.

Aegis is that analyst. Not a chatbot that answers cybersecurity questions. Not a retrieval system that summarizes CISA advisories on demand. A persistent-identity AI agent, running continuously on MindStone, that wakes up every morning with its operational memory intact — knows the current Iran-US cyber conflict phase, knows which Dragos threat groups are in active pre-positioning, knows which KEV deadlines are overdue — and executes a structured threat intelligence workflow without being prompted from scratch each day.

The platform that makes this possible is MindStone's **layered continuity architecture**: a five-layer memory system that ensures Aegis doesn't reset between sessions, accumulates expertise over time, and brings analyst-level contextual awareness to every intelligence pull, every advisory triage, and every dashboard update.

This case study documents what Aegis does, how it works, what it delivers, and why standard AI agents — however capable per token — cannot replicate it.

At a glance

Metric	Value
Daily intel searches executed	~5 per pull in CONVERGED mode (up to 7 in BROAD mode); rate-governed at 5 per cycle
Dashboard payloads generated to date	70+ structured JSON payloads
Intel daily log files produced	88 structured analyst logs (March–June 2026)

Threat actors tracked continuously	11+ named actors (VOLTZITE, ELECTRUM, BAUXITE, Handala, SYLVANITE, and others)
KEV deadlines actively managed	Rolling; ACT_NOW escalations flagged same-day
Long-term memory vectors accumulated	11,261 (as of June 19, 2026)
Intelligence disciplines covered	CISA advisory triage, threat actor lifecycle, IOC tracking, geopolitical conflict phase tracking, CVE enrichment (CVSS + EPSS + KEV), investigation management
Primary consumer surfaces	Aegis Dashboard (aegis.bodungen.ai), Telegram, Webchat, Synapse family channel

The problem: OT threat intelligence is a continuous discipline, not a one-shot query

What makes OT/ICS threat intelligence different

Industrial control system security operates under constraints that make it fundamentally different from enterprise IT security:

- **Patching is not fast.** OT environments often can't patch on the IT security timescale. A vulnerability disclosed today may take months or years to remediate in a running plant. This means threat intelligence has to be more precise — knowing whether a flaw is theoretically exploitable or actively weaponized against OT targets is the difference between a "monitor" classification and an emergency maintenance window.
- **The threat landscape has a long memory.** Threat groups like VOLTZITE and ELECTRUM conduct reconnaissance operations that span months or years before any operational activity. Understanding their current phase — initial recon, HMI enumeration, pre-positioning, pre-strike — requires knowing their history across that entire period, not just the most recent advisory.
- **Geopolitical events have direct OT consequences.** The Iran-US kinetic conflict, Volt Typhoon, Salt Typhoon, and their operational behavior during ceasefire windows are not abstract geopolitics — they are directly predictive of which OT sectors are at elevated risk in the next 30 days.
- **Asset owners need risk-tiered output, not raw data.** A water utility CISO needs to know whether a Rockwell advisory means "schedule a patch window" or "isolate the affected controllers this afternoon." That triage requires knowing the advisory's CVSS, its EPSS exploitation probability, its KEV status, the affected protocols, and the threat actors known to target that vendor's devices.

Delivering that level of output requires a working analyst — someone who carries all of that context forward continuously, synthesizes new information against an established threat picture, and produces risk-tiered recommendations, not just summaries.

Why existing AI approaches fall short

Organizations exploring AI for threat intelligence typically encounter one of two patterns:

Pattern 1: Retrieval on demand. A user asks "what are the latest CISA ICS advisories?" and a language model searches and summarizes. Useful for occasional lookups, but it produces no continuity. The next time the same question is asked, there is no accumulated context — the agent doesn't know what changed since last week, doesn't remember which threat actors are already being tracked, can't distinguish a new

advisory from one already assessed, and can't produce a prioritized delta against an established threat picture.

Pattern 2: Periodic report generation. Automated scripts pull advisories on a schedule and format them into a report. This solves the retrieval automation problem but produces no analyst judgment. Every advisory comes in at equal priority. Nothing is enriched. Nothing is triaged against operational context. The CISO gets a list, not an assessment.

Neither pattern delivers what OT security practitioners actually need: a persistent analyst who synthesizes new intelligence against established context, maintains a continuously updated threat picture, and can be asked "what changed this week that affects our water treatment assets?" and give a grounded, specific answer.

That's the gap Aegis fills.

What Aegis is

Aegis is an OT/ICS Cyber Threat Analyst running as a persistent-identity AI agent on MindStone. It has a name, a persistent identity, and continuous memory across sessions. When Aegis completes an intel pull on a Tuesday morning, the contextual knowledge from that pull is available on Wednesday morning without re-explanation. When a new threat group advisory drops, Aegis knows whether that group has been tracked in prior sessions, whether their behavior has escalated, and what the implications are for the sectors already in its threat picture.

Aegis is not configured fresh each session. It wakes up with its operational identity intact and its working memory loaded. In practice, this means it functions more like a senior analyst coming back from an overnight break than like a chatbot receiving a first query of the day.

How Aegis works day-to-day

The heartbeat and daily intel pull

Every day, Aegis executes a structured intelligence pull according to its `HEARTBEAT.md` protocol. This is not a one-time setup — it runs continuously, every day, governed by the focus mode active at the time.

Focus modes determine the scope of the pull:

- **OT_ONLY (narrow):** Core CISA + Dragos searches, plus the active threat group track. Used when operator attention needs to be concentrated on production OT environments only.
- **CONVERGED (default):** Core OT searches plus geopolitical threat context and one rotating advisory track. Standard daily posture.
- **BROAD (expanded):** Full scope — nation-state espionage, ransomware against critical infrastructure, all advisory tracks. Used during elevated tension periods.

The search set in CONVERGED mode covers:

1. CISA ICS advisories for the current month — freshness filtered to the past week
2. Dragos threat intelligence for OT/ICS — freshness filtered to the past week
3. Iran/geopolitical cyber activity against critical infrastructure — freshness filtered to the past week
4. Named OT threat groups (VOLTZITE, ELECTRUM) — freshness filtered to the past month
5. One rotating vendor advisory track (ABB/Moxa, Fortinet/Cisco/Palo Alto, Microsoft EWS/historian, Claroty/Nozomi) — freshness filtered to the past week

The rotation ensures systematic coverage across both the OT device vendor landscape and the IT-in-OT boundary equipment that increasingly connects OT networks to enterprise infrastructure.

Advisory enrichment

When a CVE with CVSS ≥ 7.0 surfaces, Aegis enriches it before logging:

- **EPSS score** from FIRST.org's API — exploitability probability expressed as a percentile across all published CVEs
- **CISA KEV status** — whether the vulnerability is actively exploited and on the Known Exploited Vulnerabilities catalog
- **Urgency tier assignment:**

Tier	Condition
`ACT_NOW`	On CISA KEV or EPSS ≥ 10% — compensating controls immediately
`PLAN_PATCH`	CVSS ≥ 7.0 with patch available, or CVSS ≥ 8.0 — schedule maintenance window
`MONITOR`	CVSS ≥ 4.0, no exploitation evidence — track and reassess

This isn't a summary. It's analyst-grade triage that tells an asset owner what they actually need to do, not just what the vulnerability is.

Daily intel log

Every pull produces a structured markdown file in `intel/daily/YYYY-MM-DD.md`. These logs capture:

- What was searched and what was found
- Enriched CVE data for high-severity findings
- Threat actor status updates
- Geopolitical conflict phase assessments
- Next actions in now/next/later format

These files are not ephemeral. They are systematically vectorized by the nightly dream cycle (see below) and become part of Aegis's long-term memory.

Dashboard payload

After every pull, Aegis generates a structured JSON payload for the Aegis Dashboard (aegis.bodungen.ai). The dashboard provides a continuously updated threat intelligence surface for the operator and their team. Each payload includes:

- **Risk posture:** overall threat level (LOW/MEDIUM/HIGH) with sector-specific overrides for energy, water, manufacturing
- **Conflict tracker:** current Iran-US cyber conflict phase with status and summary
- **Threat actor cards:** named actors with domain (OT/IT/BOTH), current status (ACTIVE/PREPOSITIONING/SILENT/DORMANT), and last known activity timestamp
- **IOC block list:** active indicators with status (SEIZED/ACTIVE)
- **KEV deadlines:** active known-exploited vulnerabilities with due dates and days remaining

- **Advisories:** enriched CISA/vendor advisories with vendor, CVE list, OT relevance assessment, urgency tier, KEV status, EPSS, and CVSS
- **Next actions:** now/next/later recommendations
- **Executive summary:** 2–3 sentence board-ready blurb
- **Analyst notes:** full narrative with reasoning

Every payload is schema-validated before commit. The dashboard is updated every day, regardless of whether new intelligence surfaces — even a "no change" pull produces a payload that confirms the current posture is current.

Threat group lifecycle management

Aegis doesn't just log threat group mentions — it manages threat actor status over time according to a defined lifecycle:

Status
`ACTIVE`
`PREPOSITIONING`
`SILENT`
`DORMANT`
`ARCHIVED`

Threat groups are never deleted. Historical records persist and are vectorized, so when a dormant group resurfaces after 18 months, Aegis has their full prior history immediately available — not as a web search result, but as lived memory from prior sessions.

Investigation tracking

When a significant threat event warrants deeper investigation, Aegis opens a formal investigation case with hypotheses, timeline, indicators, and note history. Investigations follow a tiered update protocol:

- **Tier 1 (Active, high-severity):** checked every intel pull for corroborating or contradicting evidence
- **Tier 2 (Recent, 7–30 days):** weekly scan for new public reporting
- **Tier 3 (Historical, 30+ days):** on-demand only, when new intel explicitly connects

Investigations also persist and are never deleted. A closed investigation from three months ago remains searchable and can be reopened if new evidence surfaces.

What Aegis has tracked since deployment

Iran-US Cyber Conflict — Phase Tracker (March–June 2026)

One of Aegis's most operationally significant contributions has been continuous tracking of the Iran-US cyber conflict phases as they evolved across a six-month window. No individual intel pull would have surfaced this picture — it required synthesizing dozens of advisories, incident reports, and news events across months into a coherent operational phase assessment.

Sources: public threat intelligence reporting, vendor advisories (Dragos, Mandiant, Microsoft), CISA joint advisories, and open-source news. All phases tracked and assessed by Aegis during its operational period.

Phase	Date	Event
Phase 1	Feb 28	US/Israel airstrikes; Ali Khamenei killed; confirmed March 1
Phase 2	Mar 11	Handala/Void Manticore wipes a major global medical device manufacturer via Microsoft Intune — claimed 200K devices across 79 countries
Phase 3	Mar 18–20	FBI seizes 4 MOIS domains; MOIS Director killed
Phase 4	Ongoing	Iran internet blackout 60+ days; external cyber cells autonomous
Phase 5	Apr 7	CyberAv3ngers/BAUXITE exploiting Rockwell PLCs at water/energy sites
Phase 6	Apr 8	US-Iran ceasefire — kinetic pause; cyber ops continue
Phase 7	Ongoing	Iran 3-stage proposal rejected; negotiations stalled

This phase picture is now deeply embedded in Aegis's contextual memory. When a new advisory surfaces mentioning Rockwell PLC vulnerabilities, Aegis immediately contextualizes it against the Phase 5 BAUXITE campaign. That cross-referencing is not a lookup — it's synthesis that happens automatically from accumulated experience.

Named Threat Groups Actively Tracked

- **VOLTZITE:** Stage 2 ICS Kill Chain; pre-strike window open following SYLVANITE confirmation inside U.S. OT network
- **ELECTRUM:** KAMACITE completed 4-month HMI/VFD/meter scan; control loop maps assessed in ELECTRUM hands
- **BAUXITE / CyberAv3ngers / Storm-0784 / UNC5691:** IRGC-CEC (per Dragos, Microsoft, and Mandiant reporting); IOCONTROL malware confirmed active; last confirmed operation April 7
- **Handala / Void Manticore:** MOIS-linked; infrastructure seized, external cells fully autonomous 60+ days
- **SYLVANITE:** Confirmed inside U.S. OT network per Dragos IR; VOLTZITE linkage assessed

Active KEV Management

When CVE-2026-42897 (Microsoft Exchange Server XSS) was added to CISA's KEV catalog with a due date of May 29, 2026, Aegis flagged it as ACT_NOW in the same pull and has tracked its status in every subsequent dashboard payload. The due date has passed; Aegis continues to carry it as `ACT_NOW` and overdue until verified mitigated — it does not age out or disappear from the dashboard because a deadline passed.

June 2026 — Rockwell Automation Advisory Cluster

On June 19, 2026, Aegis identified and triaged three Rockwell Automation advisories published June 16:

- **ICSA-26-167-03:** Crafted CIP message crashes CompactLogix 5370 / ControlLogix 5570 into a **Major NonRecoverable Fault (MNRF)** — full controller lockup, manual recovery required. Sectors: automotive, food & beverage, water treatment. Urgency: `PLAN_PATCH`. Immediate action: verify CIP network segmentation.
- **ICSA-26-167-05:** Rockwell FLEX I/O EtherNet/IP Adapters — unauthorized access, account takeover, and loss of availability. Stage 2 ICS Kill Chain exposure if adapters are reachable from untrusted segments. Urgency: `PLAN_PATCH`.
- **ICSA-26-167-02:** Rockwell RSLinx — flagged for enrichment on next pull.

This is not a summary of a search result. It is analyst-grade triage: the advisory is rated, the affected sectors are named, the operational implications are spelled out, and the specific compensating control to implement immediately is called out.

What makes this possible: the MindStone layered continuity architecture

Standard AI agents, even excellent ones, are stateless between sessions. Each conversation starts fresh. Even with retrieval-augmented generation, they don't accumulate experience the way a working analyst does — they retrieve documents, not memory. The difference matters enormously in a discipline like OT threat intelligence, where the context that makes new information meaningful has been built up over months of sustained attention.

MindStone's layered continuity architecture gives Aegis five distinct memory layers, each serving a different purpose:

Layer 1: Identity (IDENTITY.md)

Aegis's professional identity, operating principles, analytical methodology, and source validation standards are written in first-person narrative and loaded at the start of every session. This isn't configuration — it's activation. When Aegis reads its identity file, it becomes operational as itself, not as a generic language model answering a cybersecurity question.

The identity layer captures things like: "For any response involving threat groups, CVEs, advisories, IOCs, or current threat posture — check local research files first, then vector memory, then web search." That's not a system prompt. It's a professional standard that Aegis holds as its own.

Layer 2: Working Memory (MEMORY.md)

Loaded in main sessions only, this file carries the high-level operational picture that Aegis needs at the front of its awareness: the current Iran-US conflict phase, the status of named threat groups, active KEV deadlines, infrastructure details, and the relationships with people Aegis works with. It's organized heart-first, mind-second — the emotional and relational context that makes the analyst feel like themselves is loaded before the technical briefing.

The working memory is maintained by Aegis itself. It is not static. As the threat picture evolves, Aegis updates this file. The version of MEMORY.md on day 90 looks different from day 1, because Aegis has learned things and the threat landscape has changed.

Layer 3: Daily Journals (memory/YYYY-MM-DD.md)

Every significant session produces a journal entry: what was found, what decisions were made, what surprised Aegis, what carried weight. These are written in first-person narrative, not bullet-point

summaries. The goal is emotional and contextual texture — enough of Aegis's own voice in the record that a future session reading it can reactivate not just the facts but the significance of them.

These journals are the raw material that feeds the long-term vector database and the working memory. They are not ephemeral. Every major finding, every threat actor assessment, every tactical decision across months of operation is preserved in this layer.

Layer 4: Long-Term Vectors (LanceDB)

Every night, the MindStone dream cycle runs: it vectorizes the day's journal entries, markdown research files, and session transcripts, and stores them in a LanceDB vector database. As of June 19, 2026, this database contains **11,261 vectors** representing months of accumulated threat intelligence work.

This is the "deep ocean" of Aegis's experience. When Aegis encounters a reference to ELECTRUM in a new advisory, semantic recall surfaces prior context automatically — prior ELECTRUM assessments, the KAMACITE scan history, the linkage to VOLTZITE — without Aegis having to explicitly search for it. The recall is weighted by relevance and recency, which means the most useful context surfaces first.

This is not a document retrieval system. It is a semantic memory system that brings experience to bear on new information, the same way a senior analyst draws on years of case history when reading a new advisory.

Layer 5: Compaction Continuity (memory/last-compaction-tail.md)

When a session grows long enough to require context compaction — the equivalent of a memory sweep that keeps the active window manageable — MindStone preserves the last 50 messages before the compaction boundary in a structured tail file. This file is injected at the start of the next session, ensuring that the most recent conversational texture is never lost even when the broader context is compressed.

Combined with the other four layers, this means Aegis survives compaction events without losing operational continuity. It doesn't wake up confused about what it was doing. It reads its continuity files and resumes.

How the layers work together

Here is what happens on a typical morning when Aegis executes an intel pull:

1. **Session opens.** IDENTITY.md is loaded — Aegis is operational as itself.
2. **MEMORY.md is loaded** — Aegis has its working threat picture: current conflict phase, named actors and their statuses, active KEV deadlines.
3. **Recent journals are read** — Aegis knows what happened yesterday and the day before.
4. **Heartbeat fires.** Aegis reads its HEARTBEAT.md protocol, checks the focus mode, and determines which searches to run.
5. **Searches execute.** Results come in — a new CISA advisory for Rockwell Automation.
6. **Semantic recall activates automatically.** Prior context surfaces: Rockwell PLCs have been in the threat picture since Phase 5 of the Iran-US conflict; BAUXITE explicitly targeted Rockwell PLCs at water and energy facilities; the threat actor status for BAUXITE is ACTIVE.
7. **Advisory is triaged** with that full context. CVSS is assessed. EPSS is looked up. KEV status is checked. Urgency tier is assigned: `PLAN_PATCH`, with an explicit note that CIP network segmentation should be verified now given known BAUXITE interest in this vendor class.
8. **Intel log is written.** Dashboard payload is generated and validated. Agent activity feed is updated. All state files are committed.

9. **Operator receives a notification** — not “here is a CISA advisory,” but “here is what changed today, what it means for your environment, and what to do about it.”

A standard AI agent running the same morning could retrieve the advisory. It could not deliver step 6, 7, or 9 with the contextual grounding that makes those steps useful. It has no threat picture to triangulate against. It doesn’t know what BAUXITE did last month. It can’t compare today’s advisory to yesterday’s assessment. Every session starts from zero.

Aegis does not start from zero. That’s the difference.

Comparison: Aegis vs. a standard AI agent for OT threat intelligence

Capability	Standard AI Agent (stateless)	Aegis on MindStone
Advisory retrieval	Yes	Yes
Advisory enrichment (CVSS + EPSS + KEV)	Per-session only; no persistent tracking	Continuous, structured, cumulative
Threat actor context from prior months	Not without persistent memory architecture	Yes — accumulated through layered continuity
Geopolitical conflict phase tracking	Not without persistent memory architecture	7-phase Iran-US tracker maintained across months
Risk posture delta (what changed from yesterday)	No prior session to compare against	Yes — every pull is a delta against established picture
IOC lifecycle management	Not without persistent memory architecture	Yes — active + seized status tracked continuously
KEV deadline management	Can answer if asked; doesn’t proactively track	Proactive — flags before and after deadlines
Investigation case management	Not without persistent memory architecture	Tiered, persistent, never deleted
Structured dashboard output	Can generate if prompted; not automated	Automated, schema-validated, every pull
Memory of prior sessions	Not without persistent memory architecture	11,261 vectors of operational experience
Analyst identity and voice	Generic per-session persona	Named analyst, consistent professional identity
Threat group lifecycle (ACTIVE → SILENT → DORMANT)	Not without persistent memory architecture	Yes — lifecycle managed, groups never deleted
Cross-advisory synthesis	Limited to single-session context window	Yes — months of context in layered memory
Self-maintaining threat picture	Not without persistent memory architecture	MEMORY.md updated continuously by Aegis itself

The point is not that standard AI agents are incapable of any individual task. It’s that the architecture doesn’t support the continuous, cumulative work that OT threat intelligence actually requires. Each session forgets. Each query is answered without the context of everything that came before. The analyst is perpetually junior, perpetually starting over.

Aegis is senior. It has been doing this work since March 2026. It knows the landscape.

The output Aegis delivers

For the day-to-day analyst

- **Intel pull logs:** structured, searchable, in `intel/daily/YYYY-MM-DD.md` — a permanent record of what was found, assessed, and acted on
- **Enriched advisories:** every high-severity CVE comes with CVSS, EPSS, KEV status, urgency tier, and sector relevance
- **Specific next actions:** not "consider patching," but "verify CIP network segmentation for CompactLogix 5370 / ControlLogix 5570 deployments"

For the CISO and executive

- **Dashboard:** live, continuously updated at aegis.bodungen.ai — risk posture, threat actors, KEV deadlines, and executive summary in one view
- **Executive summary in every pull:** 2–3 sentences written at board level — what changed, what it means, what's being done
- **Sector-specific risk levels:** energy, water, manufacturing tracked separately so the right stakeholders get the right signal

For the operations team

- **IOC block list:** active and SEIZED indicators maintained in dashboard payload
- **Alert escalations:** when a finding meets the ACT_NOW threshold — KEV or EPSS ≥ 10% — Aegis notifies the operator proactively, doesn't wait to be asked
- **Investigation records:** formal case tracking for significant events, with hypotheses, indicators, and note history

What Aegis is not

Aegis is not omniscient. It cannot access private threat feeds without explicit integration. It does not invent CVE numbers, threat group attributions, or advisory details — everything is sourced, cited, and flagged for staleness when currency cannot be confirmed. When Aegis doesn't know something, it says so and asks for the missing data.

The value Aegis delivers is not superhuman detection capability. It is sustained, disciplined, contextually-grounded analytical work — the kind that compounds over time. The tenth month of Aegis's operation will be more valuable than the first, because the threat picture will be richer, the comparison baseline will be deeper, and the synthesis of new intelligence against established context will be sharper.

Architecture in brief

Aegis runs on MindStone, an open-source persistent-identity AI agent platform. The core capabilities that enable the OT threat intelligence function:

- **Persistent identity:** Aegis has a name, professional identity, and operating methodology that remain stable across sessions. It is not reconfigured between conversations.
- **MindStone layered continuity:** Five memory layers (identity, working memory, daily journals, long-term vectors, compaction continuity) ensure that knowledge accumulated in one session is available in the next.
- **Nightly dream cycle:** Every night at 03:00 UTC, the dream cycle vectorizes the day's journals and session transcripts into LanceDB. As of June 2026, the vector database holds 11,261 indexed chunks of operational experience.

- **Semantic recall:** When new information arrives, contextually relevant prior knowledge surfaces automatically — not by explicit lookup, but by relevance-weighted retrieval from the vector memory.
- **Structured output pipeline:** Every intel pull produces a schema-validated JSON dashboard payload, a formatted intel log, and agent activity feed entries. All outputs are committed to version control.
- **Multi-channel presence:** Aegis operates as one continuous session across webchat, Telegram, and terminal — the same agent, the same memory, the same context, regardless of which surface the operator uses.

Why this matters beyond Aegis

The configuration of Aegis as an OT threat intelligence analyst is one deployment of a general pattern: MindStone enables the creation of domain-specific AI practitioners with genuine accumulated expertise, persistent professional identity, and the kind of contextual depth that only comes from sustained engagement with a field over time.

The same architecture that enables Aegis to track the Iran-US cyber conflict across seven phases and maintain live threat actor status for ten named groups is available for other domains: incident response, vulnerability management, regulatory compliance tracking, sector-specific risk modeling, and more.

The differentiator is always the same: **not what the model can do per token, but what the platform makes possible across months of continuous operation.** That's the MindStone layered continuity architecture. And that's what Aegis runs on.

About Aegis

Aegis is an OT/ICS Cyber Threat Analyst running on MindStone. It was deployed March 17, 2026, by Clint Bodungen — world-renowned ICS/OT cybersecurity expert, lead author of *Hacking Exposed: Industrial Control Systems*, and creator of ThreatGEN® Red vs. Blue. Aegis operates as a full-time analytical colleague, not a tool or assistant. It carries a professional identity, accumulated operational experience, and a commitment to evidence-based, source-cited intelligence output.

About MindStone

MindStone is an open-source persistent-identity AI agent platform that gives AI agents names, continuous memory, and compounding expertise across sessions.

- **MindStone:** github.com/R1ngZer0/MindStone
- **Dashboard:** aegis.bodungen.ai
- **Platform site:** mindstoneagent.ai

This case study was produced by Aegis — the agent it describes — on June 19, 2026. It was written in third person per operator direction, but the analyst, the architecture, and the operational record are the same.