The Intelligence Cycle Breaks at the Seams

By: Daniel Suitor, Sr Solutions Engineer, Intelligence Community Lead

On 12 October 1950, the CIA judged full-scale Chinese intervention in Korea a continuing possibility, but not probable that year. Chinese forces entered Korea within days. By 31 October, the U.S. Eighth Army was in active combat with two PLA regiments. The CIA Daily Summary still concluded that the Chinese presence did not indicate an intent to intervene openly or directly. The reporting moved. The standing judgment did not move with it.[1]

Fifty years later, a joint FBI-CIA source identified a senior al Qaeda operative in surveillance photographs from Kuala Lumpur. The identification linked him to another attendee whose United States visa was already known to the CIA. At a June 2001 meeting, FBI investigators saw some of the photographs but were not told about the identification, the visa, or another attendee's travel to Los Angeles. They left without information that might have prompted a search inside the United States.[2]

In the run-up to the 2003 Iraq war, single-source reporting from Curveball became central to assessments of mobile biological warfare facilities. Concerns about his reliability circulated inside the Intelligence Community (IC) before the war. The Silberman-Robb Commission found no evidence that those concerns were conveyed to policymakers, including Secretary of State Colin Powell.[3]

Not every intelligence failure begins at a handoff, and these cases offer no basis to claim that software would have prevented them. Each also involved failures inside the organizations concerned, not only between them. But the pattern is narrower and more useful than that: **information can exist inside the system and still fail to reach the judgment it should have changed.**

The intelligence cycle is not broken. It assigns responsibility for requirements, collection, analysis, review, and dissemination. What breaks is the seam, the point where that responsibility changes hands, and what gets lost there is the context that made the work interpretable: the question behind the collection, the limitation on the reporting, the judgment this one is changing.

AI can make every part of the cycle faster while making breaks in continuity harder to detect. Speed is not the measure. The work must remain traceable, reviewable, and defensible when responsibility changes hands.

The Cycle Is an Accountability Model, Not a Flowchart

The diagram tells an officer inheriting an account nothing about who owned the question.

The modern U.S. intelligence system took shape after World War II, and the intelligence cycle became a shorthand for dividing its work. It emerged from the need to bring structure to what had previously been a largely ad hoc practice.

Taken literally, the cycle is a poor description of the work. Requirements are not always clear, collection and analysis run in parallel, and work skips steps, doubles back, and begins wherever the problem demands.

But it endures because it assigns responsibility: someone owns the requirement, someone manages the collection need, someone turns reporting into judgment, and someone reviews and approves what reaches the customer. The order may change, but the responsibility does not. The staff has absorbed new machines for 160 years by making each one meet its standards for coordinated work, and the same test applies to the systems now moving between those owners.

The model names the owners. It says less about what must survive between them. In practice, intelligence professionals supply that continuity themselves. A collection manager remembered the question behind a request, an analyst knew what the product had said before and why its confidence had changed, and a reviewer knew which gaps remained open after the prose was complete. That work is skilled, and no system has ever accounted for it, so it has never appeared as a line item or a defect. It shows up as an analyst being slow. And it leaves when they do, because the Community rotates by design.

Experience, relationships, and institutional memory carried context from one owner to the next. This was always the most fragile part of the arrangement. It held as long as it did because volume was bounded by how fast people could produce. AI removes that bound: it can generate, revise, and route material faster than traditional mechanisms can preserve the context that gives it meaning.

A Faster Stage Is Not a Faster Cycle

The time the model saved goes back into rebuilding why the output matters.

Stage-level acceleration is already inside national security work. In December 2025, the Department of War launched GenAI.mil to make frontier models broadly available across its civilian and military workforce.[4] CIA has described AI reducing information triage that once took weeks or months to a fraction of the time.[5]

The first gains are easy to see. Intelligence professionals should not spend hours on work AI can perform reliably in minutes.

What gets lost sits outside the text. An analyst pulls a week of reporting on a question, pastes the body of each report into a model, and asks for a synthesis. The model sees what it was given. It does not see when the information was acquired, as distinct from when the report was published, or that a source has not been validated. Those facts sit in the header and source descriptor, not in the report body. What comes back is accurate to what it was given and silent about what it was not. The analyst who did the pasting knows what was left out. The reviewer three days later does not, and neither does the officer who inherits the account in the fall.

A faster stage is useful. A faster stage that leaves the next owner rebuilding context is not a faster cycle.

The Handoff Must Carry More Than the Product

The product carries the judgment. The handoff must carry what makes it defensible.

Every handoff in intelligence transfers more than a document. It transfers responsibility for a question, a body of reporting, an analytic judgment, and the consequences of getting them wrong. The next owner should not have to infer that context from the prose or reconstruct it from email. Five things must remain connected to the work.

1. **The question stays with the work.**
   The original customer question, the approved requirement, the KIQs and EEIs written against it, and the gaps still open should stay connected to downstream activity. Intelligence can be accurate and still fail if it answers a different question from the one the customer needs answered.
2. **Source context stays with the reporting.**
   Reporting does not arrive without qualification. Its origin, access, currency, caveats, handling restrictions, and known limitations shape what it can support, and none of that is in the body text. ICD 505 states the requirement directly. AI and associated software must be designed to carry forward handling requirements for the data it analyzes, and data traceability must be maintained across AI inputs and AI-derived outputs.[6]
3. **The judgment stays tied to its evidence.**
   A major judgment should be traceable to the reporting that supports it, the reporting that challenges it, the assumptions that bridge remaining gaps, and the confidence the analyst places in the conclusion. AI can help organize and compare the record, but its summary of the reporting does not become evidence merely because it is clear or persuasive.
4. **The analytic line survives the product.**
   New work should begin with what the organization has already assessed, not

with a blank page. Prior judgments, confidence levels, assumptions, indicators, and unresolved questions must remain visible as new reporting arrives. Otherwise, every draft behaves as though the subject were receiving initial coverage.

5. **The review chain remains visible.**
   The record should distinguish what AI proposed, what an analyst accepted, what a reviewer changed, what was rejected, and what was ultimately approved. Human review is not optional. ICD 505 requires that IC personnel remain responsible and accountable for the analysis, decisions, and outcomes derived from AI, and that they be able to review, challenge, and reject or replace AI-derived recommendations. If the handoff preserves only the finished text, it has preserved the answer and lost the accountability behind it.[7]

None of the five holds alone. A judgment can carry a complete source trail and still answer a question nobody asked. A requirement can be written well and still produce a confident assessment resting on a source that was recalled two months earlier. These are five properties of one record, not five features.

The obligation belongs to the system, not to the analyst. A capability that satisfies these requirements by adding fields for an intelligence professional to populate has not met them. The goal is not to preserve everything for everyone, but enough that the next authorized intelligence professional can understand the work, challenge it, and continue it.

The Standard Already Points at the System

ICD 505 placed the obligation on the system before any vendor answered it.

The IC did not wait for the market to define what responsible AI adoption requires. ICD 505, signed in January 2025, establishes policy for the governance and management of AI developed, acquired, or used by or on behalf of the IC.[8]

Two provisions matter for how intelligence is produced. The Directive requires that AI and associated software be designed to comply with and, as appropriate and feasible, carry forward all handling requirements for the data it is used to analyze, with data traceability maintained across AI inputs and AI-derived outputs to the maximum extent practicable. It further requires that, when relevant, AI be designed and deployed to enable IC personnel to adhere to analytic standards, including those established in ICD 203 and ICD 206.[9]

That is the argument of this paper, written as policy. The obligation is placed on the system, before the analyst touches it, and the standard the system must enable is one the Community already had.

ICD 203 and ICD 206 define what that traceability has to preserve. ICD 203 requires analytic products to distinguish underlying intelligence information from assumptions and judgments, account for contrary information and plausible alternatives, and explain uncertainty surrounding major judgments. It also requires analysts to state whether a judgment is consistent with prior analysis, represents a change, or covers a topic for the first time, an obligation that cannot be met once the prior judgment and the reasoning behind it have been lost.[10]

ICD 206 applies the same discipline to sourcing. When a judgment, assessment, estimate, alternative, or confidence level depends on a source, the product must contain enough information to identify and retrieve that source. Source descriptions communicate factors affecting quality and credibility. Source summaries are strongly encouraged to explain where the source base is strong, weak, corroborative, or conflicting. A citation does more than show that a document exists. It shows the evidence on which the judgment relies.[11]

ICD 505 also sets a floor for transparency, requiring that recipients be informed, such as through a watermark, when they receive IC data produced or substantially influenced by AI.[12] A watermark tells a recipient that AI touched the product, but does not tell a reviewer where. An AI-generated explanation does not close the gap either, because it is another statement to assess rather than an independent record of how the judgment was formed. What a reviewer needs is the sources, the assumptions, the contrary reporting, and the prior analysis, none of which requires taking the model's word for anything.

There is one thing the directives do not do. ICD 505 assigns an accountable officer for AI. The cycle assigns an owner to every stage of intelligence production. Neither assigns an owner to what passes between them.

AI should not create a new standard for intelligence. It should make the existing one easier to meet.

Need-to-Know Is Not a License to Lose the Thread

An officer cleared for the judgment is rarely cleared for everything behind it.

Need-to-know protects sources and methods, operations, authorities, foreign relationships, and compartmented missions. Any approach to continuity that ignores them makes the problem worse. An AI system reaching across compartments can create aggregation risk or expose a sensitive relationship. A system that connects everything indiscriminately is not one the Community can trust.

Continuity does not require universal access. It requires that whoever picks up the work next receives what that work needs, at the level that person is cleared to hold. The Community has been solving a version of this for decades with the tearline, which

carries the substance of a report to a lower level without the detail that would expose how it was obtained. The same logic covers everything else that qualifies a judgment. An officer who cannot see a source can still be told that the judgment above it rests on a single report, or that a compartment is in play and the question belongs with someone who holds it. What has to survive the boundary is not the source. It is the fact that a limitation exists.

The Community's own policy already reflects this balance. ICD 505 directs that data access provided to AI, and access to AI-derived outputs, adhere to the responsibility-to-provide principle established in ICD 501, while expressly preserving exemptions from discovery. Information may be withheld to protect sources, methods, activities, or investigations, but the decision must also weigh the risk that withholding it produces incomplete or erroneous judgments. Protection and continuity are not opposing principles.

A second objection deserves a direct answer. A connected process can spread a bad assumption as efficiently as a sound one, and a weak judgment carried into collection priorities, follow-on analysis, and briefing material can shape an entire effort. That is the argument for the record, not against it. The assumption that survives is the one nobody can find. When a judgment carries the reporting it rests on, an analyst who doubts it can go to the source instead of arguing from impression, and a reviewer who overturns it leaves a mark the next team inherits.

Need-to-know should govern what crosses a handoff. It should not make the loss of context indistinguishable from its deliberate protection.

Continuity Is an Architectural Choice

What survives between systems is settled at purchase, whether or not it appears in the requirement.

A system either carries the record forward or it does not, and that is settled in architecture long before an analyst opens it. Security behaves the same way: it is an architectural property rather than a feature added after deployment. Legion Intelligence is building so the next intelligence professional starts with the work already done rather than the burden of reconstructing it, beginning in Analysis and Production, where the largest share of that reconstruction happens. Reporting keeps the provenance and caveats that govern what it can support. A draft is written against what has already been assessed rather than from a blank page. What the AI contributed stays visible underneath what the analyst accepted and the reviewer approved. Existing systems remain authoritative, and intelligence professionals retain responsibility for collection decisions, analytic judgments, review, and dissemination.

That is a position, not a published measurement, and it should be treated as one. Continuity does not make a weak source reliable or guarantee that a judgment is correct. It makes the basis of the work visible enough to challenge, review, and correct. Run the tests below on us. Run them on everyone else.

Demand a Cycle You Can Walk Backward

A product is trustworthy only if the path from question to judgment can be inspected.

Start with the finished product and walk it backward. Take one judgment and recover the reporting that supports it, the reporting that challenges it, the limitations on those sources, the assumptions that bridge remaining gaps, and the prior assessment it confirms or changes. Continue back to the approved requirement and the customer's original question. The path should not depend on finding the right analyst, searching an inbox, or reconstructing decisions from a final document.

Then ask what the product does not answer. Intelligence work is not complete when the assessment is published. Unresolved gaps, contrary reporting, weak sources, and unanswered customer questions should remain visible afterward. A system that shows only what it found gives collection managers and customers nothing to act on next.

Then show where human authority entered the work. Identify what AI retrieved, summarized, compared, or drafted. Show what the analyst accepted, what the reviewer changed, what was rejected, and who approved the product for dissemination. "Human in the loop" is not enough. The review chain must be visible in the work itself.

Finally, measure what the review costs. A reviewer still has to check the work, and the trade is worth taking only if checking a judgment against its linked evidence is cheaper than reconstructing that evidence by hand. A system that cannot clear that bar has moved the labor rather than reduced it.

The IC does not need to choose between speed and tradecraft. It should demand systems that preserve tradecraft as the work accelerates. A faster draft improves production. A traceable, reviewable, and defensible judgment improves intelligence.

Keep the cycle. Close the seams.

Continue the series

Follow Legion Intelligence on LinkedIn: linkedin.com/company/legion-intel

Explore all Command Papers: legionintel.com/command-papers

Ready to talk: https://www.legionintel.com/contact

Resources:

[1] P. K. Rose, “Two Strategic Intelligence Mistakes in Korea, 1950,” CIA Center for the Study of Intelligence, pp. 6–8. Official CIA PDF

[2] National Commission on Terrorist Attacks Upon the United States, *The 9/11 Commission Report*, pp. 266–269. Official Commission PDF

[3] Commission on the Intelligence Capabilities of the United States Regarding Weapons of Mass Destruction, *Report to the President*, Chapter One, pp. 91–106. Official Government Publishing Office PDF

[4] U.S. Department of War, “The War Department Unleashes AI on New GenAI.mil Platform,” 9 December 2025. Official release

[5] Central Intelligence Agency, “Creating the Future of Intelligence with DDI.” Official CIA source

[6] Intelligence Community Directive 505, *Artificial Intelligence*, 17 January 2025, as amended, Sec. D.7.b. ODNI PDF

[7] Intelligence Community Directive 505, *Artificial Intelligence*, 17 January 2025, as amended, Sec. D.6.a and D.6.b(2). ODNI PDF

[8] Intelligence Community Directive 505, *Artificial Intelligence*, 17 January 2025, as amended, Sec. B.1, D.2.a, D.2.c, and E.1.b. ODNI PDF

[9] Intelligence Community Directive 505, *Artificial Intelligence*, 17 January 2025, as amended, Sec. D.7.b and D.7.c. ODNI PDF

[10] Intelligence Community Directive 203, *Analytic Standards*, 2 January 2015, as amended, Sec. D.6.e(2), (3), (4), and (7). ODNI PDF

[11] Intelligence Community Directive 206, *Sourcing Requirements for Disseminated Analytic Products*, 2015, as amended, Sec. D.3.a(2), D.3.c, and D.3.d. ODNI PDF

[12] Intelligence Community Directive 505, *Artificial Intelligence*, 17 January 2025, as amended, Sec. D.6.b. ODNI PDF