WHITE PAPER

JUST-IN-TIME INTELLIGENCE

Delivering the Right Intelligence to the Right People at the Right Time

OPERATIONALIZING ORGANIZATIONAL MEMORY THROUGH TARGETED DISSEMINATION

A Framework for Intelligence Utilization
in the Age of AI and Anticipatory Intelligence



OBSERVE
Understand the environment.

LEARN
Transform data into knowledge.

ANTICIPATE
See what's coming next.

ACT
Make informed decisions.

Bridging the gap between intelligence production and intelligence utilization to build learning organizations that anticipate, adapt, and succeed.

2026

Just-in-Time Intelligence

Delivering the right intelligence to the right people at the right time



An Intelligence White Paper

© 2026

Table of Contents

Executive Summary 6

The Intellectual Foundations of Just-in-Time Intelligence 7

Part I: The Intelligence Utilization Problem 9

Introduction 9

The Challenge Is Not a Lack of Information 9

The Evolution of Intelligence 10

The Dissemination-Utilization Gap 11

Information Overload and the Attention Economy 13

The Science of Attention, Learning, and Memory 13

The Failure of Generalized Awareness 14

Intelligence as a Decision Support Capability 15

From Explanation to Anticipation 16

Toward Intelligence Utilization 17

Part II: The OLAA Model and Organizational Memory 18

From Linear Cycle to Continuous Adaptation 18

Organizational Memory as the Foundation 19

Observation 19

Learning 20

Anticipation 20

Action 21

The Continuous Adaptive Loop 21

Anticipation as a Core Intelligence Function 22

From OLAA to Just-in-Time Intelligence 23

OLAA in Practice: A Worked Example 25

Summary and What Comes Next 26

Part III: Building Organizational Memory 27

Beyond Information Storage 27

The Organizational Learning Problem 28

Knowledge Creation and the SECI Model 29

Tacit and Explicit Knowledge **30**
Memory Decay and Organizational Forgetting **31**
Organizational Memory as a Strategic Asset **32**
Organizational Memory and the OLAA Model **33**
Organizational Memory and Just-in-Time Intelligence **34**
The Goal: A Learning Organization **35**
Memory Must Be Engineered **36**
What Organizational Memory Must Hold **36**
From Archive to Operational Memory **37**
Knowledge Management as Intelligence Infrastructure **38**
The Half-Life of Knowledge **38**
Capturing What Walks Out the Door **39**
Memory in Action: Revisiting the Wire-Fraud Case **40**
From Memory to Anticipation **40**
Part IV: Anticipation and Targeting ***41***
The Targeting Problem **41**
Why the Series Method Exists **41**
The Four Functions of Intelligence **42**
Series 1: Threat Warning **44**
Organizing Requirements by Purpose **45**
Strategic, Operational, and Tactical Alignment **46**
Collection Requirements and Traceability **46**
The Series Method as a Targeting Architecture **47**
The Series Method and Organizational Memory **48**
The Series Method as the Foundation of Just-in-Time Intelligence **48**
The Series Method in Action **49**
From Targeting to Delivery **50**
Part V: The Architecture of Just-in-Time Intelligence ***51***
From Principle to System **51**
What a Just-in-Time Architecture Must Do **51**

Theia: An AI-Assisted Intelligence Management Platform **53**
Requirements: The IR Wizard and the Series Method **53**
Collection: AI-Assisted and Standards-Based **54**
Memory and Analysis: The Analytics Workbench **54**
Orchestration and Delivery: Theia Flow **54**
Threat Warning in Action: A Business Email Compromise **55**
The Closed Loop: Measurement and Maturity **56**
Deployment at the Operational Edge **56**
Part VI: Measuring Utilization, Not Output ***59***
The Measurement Problem Restated **59**
Why Output Metrics Persist and Mislead **59**
What Utilization Metrics Actually Measure **59**
From Metrics to Maturity: Why a Pyramid **60**
The Five Levels of the Intelligence Maturity Pyramid **61**
Reading Maturity Across Dimensions **62**
Maturity as a Diagnostic, not a Scorecard **62**
Closing the Loop on Measurement **63**
Part VII: The Operating Model ***64***
From Capability to Practice **64**
The Risk of Entropy **64**
Roles That Own the Loop **64**
Cadence: The Rhythm of the Loop **65**
Governance and the Discipline of Restraint **65**
Culture and Adoption **66**
Sustaining the Loop Through Change **66**
Bringing It Together **67**
Part VIII: The Future of Intelligence ***68***
The Argument in Brief **68**
A Change in What Counts as Good Intelligence **68**
Why This Matters Now **69**

The Climb Ahead **69**

Conclusion ***72***

References ***74***

Intelligence Theory and Warning Intelligence **74**

Organizational Learning **74**

Knowledge Management and Organizational Memory **74**

Decision Science and Human Judgment **75**

Learning Science and Human Performance **75**

Intelligence Utilization and Performance Measurement **76**

Executive Summary

Modern organizations possess more intelligence than ever, yet they still experience preventable failures because critical knowledge often fails to reach the people who can act on it when decisions are made. This white paper argues that the central challenge facing intelligence programs is no longer simply collection, analysis, or warning. It is utilization: the ability to convert what an organization knows into timely, observable action.

The paper introduces the Intelligence Utilization Problem and explains why traditional dissemination models, generalized awareness programs, and output-based metrics are insufficient in an environment of attention scarcity and operational complexity. It then presents the OLAA Model (Observation, Learning, Anticipation, and Action) as a continuous, adaptive loop anchored in organizational memory. Through this model, organizations can preserve lessons learned, detect emerging decision points, anticipate who needs specific intelligence, and deliver actionable guidance at the precise moment it can influence behavior.

Just-in-Time Intelligence is the operational expression of this approach. Rather than relying on broad, calendar-driven briefings or training that may be forgotten before they matter, it draws on organizational memory and live signals to deliver targeted, contextual intelligence to the right decision-maker at the point of need. The result is a shift from measuring intelligence by what is produced to measuring it by what changes: improved decisions, reduced risks, altered behaviors, and stronger outcomes.

The Intellectual Foundations of Just-in-Time Intelligence

Just-in-Time Intelligence did not arise from a single discipline. It converges four mature fields of study: intelligence, organizational learning, knowledge management, and decision science. Each contributes a critical element to the framework presented in this paper.

Intelligence theory provides the foundation for understanding uncertainty, warning, and decision support. Sherman Kent argued that the purpose of intelligence is not to produce information but to reduce uncertainty for decision makers. Later scholars and practitioners, such as Cynthia Grabo, expanded this view through warning intelligence, emphasizing the detection of indicators and conditions before threats materialize. Richards Heuer further demonstrated that intelligence failures often stem not from a lack of information but from the cognitive limitations of decision makers and analysts.

Organizational learning supports the view that organizations can become progressively more effective by systematically capturing and applying experience. Peter Senge's concept of the Learning Organization describes institutions that continuously improve their ability to achieve desired outcomes. Argyris and Schön further distinguished between organizations that merely react to events and those that learn from outcomes and adapt their assumptions and behaviors.

Knowledge management provides the mechanisms that enable organizational learning at scale. Nonaka and Takeuchi demonstrated that organizations create value by transforming individuals' tacit knowledge into explicit knowledge that can be shared, retained, and reused. This process forms the theoretical basis for Organizational Memory, a concept central to the OLAA Model and Just-in-Time Intelligence.

Decision science explains why timing matters. Herbert Simon's theory of bounded rationality demonstrated that decision makers operate with limited information, time, and cognitive capacity. Gary Klein's work on Recognition-Primed Decision Making showed that experienced professionals rely heavily on pattern recognition and contextual cues when making decisions under pressure. Daniel Kahneman's research further revealed how cognitive biases influence decision-making, particularly in environments characterized by uncertainty and complexity.

Together, these disciplines point to a common conclusion. Organizations do not fail because they lack information. They fail because knowledge is often unavailable, forgotten, disconnected from context, or delivered after decisions have already been made.

Just-in-Time Intelligence addresses this challenge by integrating warning intelligence, organizational memory, anticipatory analysis, and targeted dissemination to ensure that knowledge reaches decision-makers at the precise moment it can influence behavior and outcomes.

The OLAA Model provides the operational framework for this process. Organizational Memory provides the knowledge substrate. The Series Method provides the targeting mechanism. Just-in-Time Intelligence provides the delivery model.

Together, they create an intelligence capability focused not on information production but on organizational learning and decision advantage.

Part I: The Intelligence Utilization Problem

Introduction

Organizations have never had more information than they do today. Within the security domain alone, teams ingest a continuous stream of threat intelligence from commercial vendors, Information Sharing and Analysis Centers, government agencies, open-source platforms, and internal telemetry systems. Beyond security, business leaders are surrounded by market intelligence, competitive analysis, financial forecasts, regulatory updates, and operational reporting. Layered moreover, artificial intelligence systems now generate summaries, recommendations, predictions, and alerts at a pace unimaginable only a decade ago. By almost any measure of volume, coverage, or access, the modern enterprise is richer in information than any organization in history.

Yet despite this abundance, organizations still suffer failures the available information should have prevented. Companies fall victim to phishing campaigns that their own threat intelligence teams had already described in detail. Finance departments approve fraudulent wire transfers weeks or months after every employee has completed mandatory fraud awareness training. Executives are deceived by impersonation attempts even though they sat through briefings that warned them about that exact technique. Critical vulnerabilities remain unpatched long after they have been publicly disclosed, widely reported, and flagged by internal scanning tools. In each of these cases, the information needed to avoid the outcome already existed somewhere inside the organization. It simply failed to reach the right person, in the right form, at the right time, when a decision was being made.

That recurring pattern is the focus of this paper. The persistent gap between what an organization knows and what it does is not primarily a collection or analysis problem. It is a utilization problem, and it has quietly become one of the defining challenges of modern intelligence work.

The Challenge Is Not a Lack of Information

It is tempting to respond to each new failure by gathering more information. If an incident slipped through, the reasoning goes, the organization must have been missing a feed, a tool, or a report. Occasionally, that is true. Far more often, the uncomfortable reality is that the relevant information was already on hand but did not change anyone's behavior. The problem is rarely scarcity. The problem is that information fails to influence action.

This disconnect is one of the most significant yet least discussed problems in modern intelligence operations. Intelligence professionals devote enormous energy to collection, analysis, and dissemination. By comparison, they devote relatively little energy to a harder

question: whether the intelligence being produced is used by the people responsible for making decisions and acting. Because that question is so rarely asked, organizations tend to confuse intelligence production with intelligence effectiveness. They treat the existence of a report as if it were the same as its impact.

The two are not the same. Reports, dashboards, briefings, and alerts have value only insofar as they change something. A product that influences a decision, shifts behavior, reduces risk, or enables action has earned its place. A product that does none of these things, however polished or accurate it may be, has not. An intelligence report that is never read confers no advantage. An alert that arrives after the decision has already been made adds little value. A warning sent to people who cannot act on it becomes, in practice, just more noise in an already crowded channel.

This leads to a principle that belongs at the center of any serious intelligence program. The true measure of intelligence is not whether it was produced. It is whether it *changed something*. That single distinction is the foundation of what this paper calls the Intelligence Utilization Problem.

The Evolution of Intelligence

To understand why utilization is so often neglected, it helps to trace the discipline's evolution. Throughout its history, the field has organized around a sequence of increasingly sophisticated questions, and each answer has reshaped practitioners' understanding of what constitutes excellent work.

Early intelligence efforts focused on collection. The central challenge was acquiring information adversaries, or competitors, did not want revealed or did not know was observable. Success was measured by access, coverage, and the sheer volume of material a service could obtain. Knowing what others did not know was an advantage.

As the disciplines matured, the emphasis shifted from collection to analysis. Gathering information was no longer enough; organizations needed to understand it. Practitioners developed tradecraft standards, structured analytic techniques, estimative language, and reporting formats to transform raw material into genuine understanding. The prized skill became the ability to weigh competing hypotheses and explain not only what was happening but also what it meant.

In recent decades, the field has leaned heavily toward warning. Rather than merely explaining what had already occurred, intelligence was expected to illuminate what might happen next. Predictive analysis, strategic warning, indications-and-warning methodologies, and risk forecasting became central to the mission. The most valued analyst was the one who could see around the corner before the event arrived.

Each shift marked a real advance, and modern practice depends on all three. Yet collection, analysis, and warning share a quiet, largely unexamined assumption. Each assumes that intelligence, once produced, will naturally find its way into the organization's decisions. Experience repeatedly shows that it does not. Organizations routinely generate excellent intelligence that never reaches the people best positioned to act on it. Even when it reaches the intended audience, it often arrives too early to feel relevant, too late to matter, or stripped of the context needed to drive a decision.

This pattern is especially evident in cybersecurity, where production systems are mature and fast. Threat intelligence teams accurately characterize emerging adversary activity. Security operations centers detect indicators of compromise. Analysts publish timely assessments of risks and vulnerabilities. Yet incidents still occur, not because the analysis failed, but because the person who could have prevented the incident never received the relevant intelligence when it mattered. The failure is not analytical. It is one of utilization.

Viewed through this lens, intelligence has evolved through successive stages. Collection asks, "What do we know?" Analysis asks, "What does it mean?" Warning asks, "What may happen next?" Utilization asks a different question: "Did it change anything?" The first three stages focus on producing intelligence, while the fourth focuses on ensuring that intelligence has an effect.

The Dissemination-Utilization Gap

Much of the problem is structural. Most intelligence models, in their various forms, end with dissemination. The familiar intelligence cycle moves through collection, processing, analysis, dissemination, and feedback. Feedback acknowledges that evaluation matters, but dissemination remains the final operational act, the point at which the producing organization has done its job and hands the product off.

That structure quietly encodes a dangerous assumption: once intelligence has been disseminated, the mission is complete. Dissemination and utilization are not the same, and the difference is not academic. Dissemination is the distribution of intelligence. Utilization is the application of intelligence. One describes movement; the other describes effect.

The implications become clear as soon as you separate the two. A threat intelligence report emailed to one hundred recipients has unquestionably been disseminated, but whether any recipient acted on it remains unknown. A security awareness bulletin posted to an internal portal has been disseminated, yet whether any employee changed behavior as a result remains unknown. An executive briefing delivered to senior leadership has been disseminated, but whether leadership altered a strategic decision remains unknown. In

each case, the producing team can credibly claim the work was done, even though the only outcome that matters has gone entirely unmeasured.

The space between dissemination and utilization is where most of intelligence's value quietly leaks away. Organizations invest heavily in collection capabilities, intelligence platforms, analyst staffing, and reporting mechanisms, yet comparatively little in understanding how or whether intelligence shapes operational behavior. Because utilization is hard to see, programs default to measuring what is easy to count.

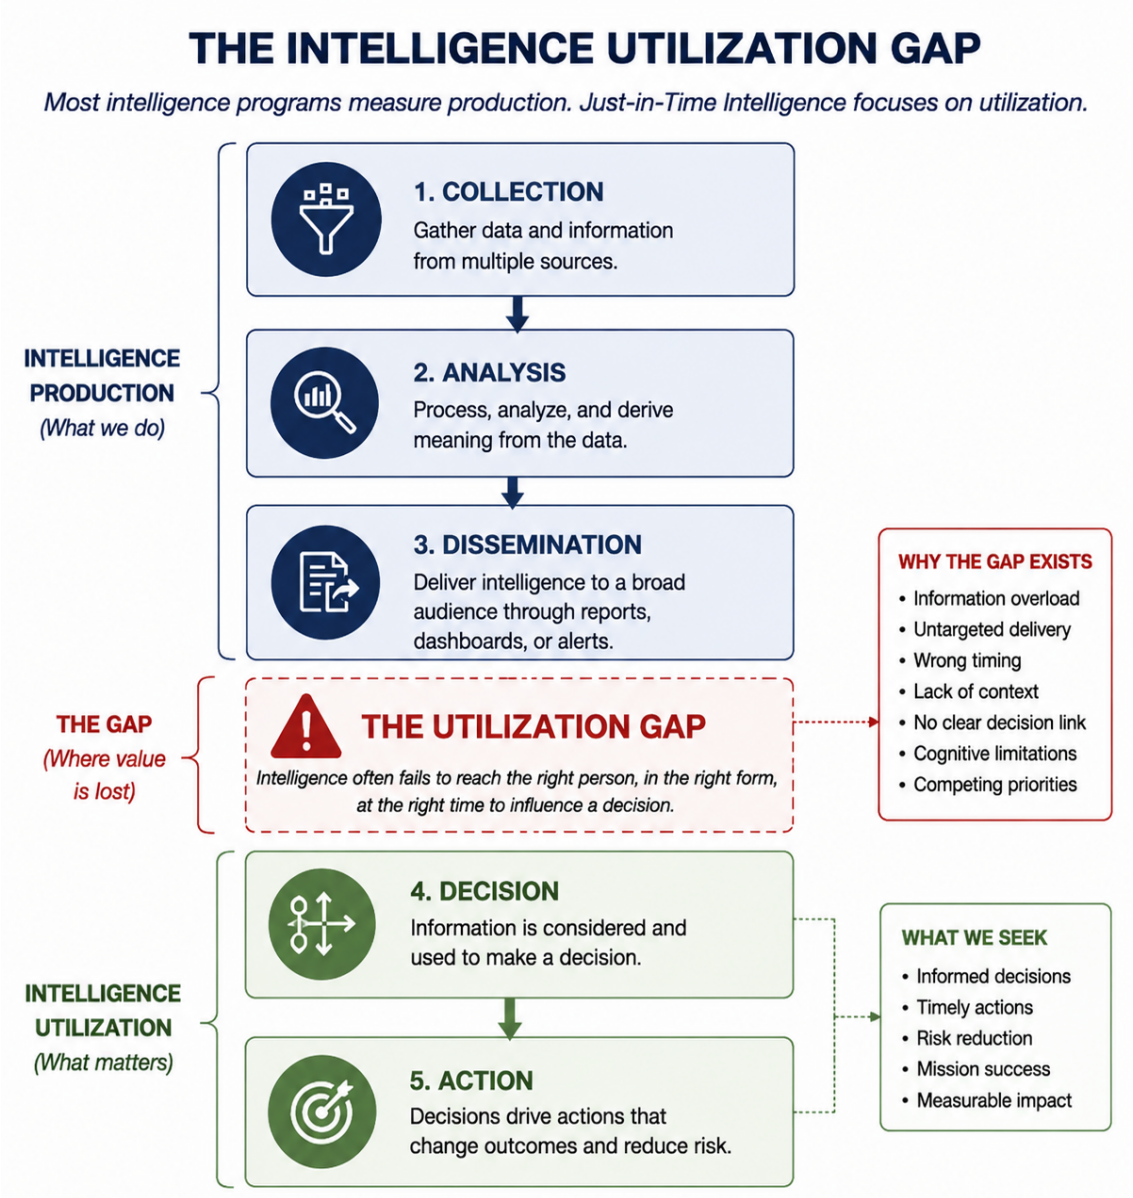


Image 1: The Intelligence Utilization Gap

As a result, intelligence programs tend to measure outputs rather than outcomes. Typical metrics include the number of reports produced, indicators collected, threat actors

tracked, intelligence requirements completed, and briefings conducted. Each of these measures activity. None of them measures impact. An organization can produce hundreds of intelligence products in a year and still fail to prevent the very incidents those products described. The existence of intelligence does not guarantee its use.

Information Overload and the Attention Economy

The utilization problem is made considerably worse by the ordinary conditions of modern work. Within any large organization, every department competes for the same scarce resource: attention. A typical employee moves through a day filled with email, instant messages, meeting invitations, project updates, compliance notices, training requirements, system alerts, operational reports, and organizational announcements. Security and intelligence teams add a steady stream of vulnerability notifications, phishing warnings, intelligence summaries, and risk assessments to that flow.

In this environment, intelligence is one voice among many. The defining constraint of the modern enterprise is no longer information scarcity; it is attention scarcity. Human beings have finite cognitive capacity, and decision makers naturally prioritize information that seems immediately relevant to their current responsibilities. Information that lacks urgency, context, or obvious relevance tends to be set aside, regardless of how important it may be in some objective sense.

This creates a genuine paradox. As an organization generates more intelligence, the likelihood that any single intelligence product receives meaningful attention tends to decline. More volume can mean less impact. Valuable intelligence competes with thousands of other signals for a sliver of cognitive bandwidth and often loses. In many organizations, intelligence is not ignored because it lacks quality. It is ignored because it lacks immediacy.

Organizations increasingly operate in an environment characterized by information abundance but attention scarcity (Kahneman, 2011; Sweller, 1988).

The Science of Attention, Learning, and Memory

The utilization problem is not merely an intelligence challenge. It is fundamentally a human-performance challenge.

For more than a century, researchers have studied how people learn, retain information, and make decisions. Their findings reveal a consistent pattern that has direct implications for intelligence dissemination.

In the late nineteenth century, Hermann Ebbinghaus demonstrated that human memory decays rapidly over time unless knowledge is reinforced through repeated exposure and

practical application. His Forgetting Curve remains one of the most influential findings in learning science. Information delivered months before it becomes relevant is unlikely to remain fully accessible when a decision must be made.

This has significant implications for traditional awareness programs. Annual training and periodic briefings may transfer knowledge at the time of instruction, but retention declines over time. By the time a phishing email, wire fraud attempt, or impersonation attack occurs, much of the original training has been forgotten.

Malcolm Knowles' work on adult learning reached a similar conclusion. Adults learn most effectively when information is immediately relevant to a problem they are actively trying to solve. Learning disconnected from context often produces awareness without action. By contrast, learning connected to a current challenge is far more likely to influence behavior.

Research in Cognitive Load Theory further demonstrates that human attention is limited. Decision-makers operate under competing demands, constant interruptions, and a significant informational burden. As a result, organizations cannot assume that simply producing more intelligence will improve outcomes. In many cases, additional information increases complexity rather than reducing it.

These findings reinforce a central premise of Just-in-Time Intelligence. The challenge is not simply to provide information. It is to provide the right information now it can be used.

Rather than expecting individuals to remember everything they have ever been taught, JITI allows the organization to retain knowledge through Organizational Memory and deliver only the knowledge required for the decision at hand.

The Failure of Generalized Awareness

Faced with these utilization challenges, many organizations turn to generalized awareness as the solution. Employees complete annual security awareness training. Executives attend periodic threat briefings. Departments receive newsletters summarizing emerging risks. Awareness campaigns aim to educate broad audiences about a wide range of threats, on the theory that a better-informed workforce will make better decisions when the moment comes.

These programs are worthwhile and should not be dismissed. They establish a shared vocabulary, satisfy legitimate compliance obligations, and raise the general baseline of understanding. But generalized awareness has inherent limitations that no amount of polish can fully overcome. People often forget knowledge acquired months before an event by the time it arrives. Information not tied to a person's specific responsibilities is easy to disregard. And training delivered outside any operational context routinely fails to influence

real decisions made under pressure. Learning science shows that knowledge retention declines significantly over time without reinforcement through contextual application (Ebbinghaus, 1913).

Consider a financial analyst who completes annual fraud awareness training in January. The session is competent, and the analyst pays attention. By September, however, the specifics have faded, displaced by nine months of unrelated work. When a realistic business email compromise attempt arrives in the inbox, carefully timed to coincide with a genuine vendor payment, the analyst may fail to recognize the warning signs, even though the relevant education was delivered earlier that same year. The knowledge existed. It simply was not available at the time of the decision.

Now consider the same analyst in a different design. Shortly before a wave of wire-approval activity, the analyst receives a brief, targeted warning that a wire-fraud campaign is actively targeting organizations in their industry. The warning arrives just before the exposure period. It provides context specific to the analyst's role and the types of payments they handle. It identifies the indicators to watch for and the verification step to take before releasing funds. In that form, the information is not abstract awareness; it is operationally relevant guidance arriving exactly when it can be used, and the likelihood that the analyst acts on it rises dramatically.

The same contrast runs throughout the enterprise. A software engineer who sat through a secure-coding course last spring may still introduce a vulnerable dependency in the autumn, while the same engineer who receives a brief, contextual alert the moment that dependency is added to a build is far more likely to swap it out. An executive briefed once a year on impersonation risk may still authorize an urgent transfer requested by a convincing voice on the phone, while the same executive who receives a timely note that deepfake voice attacks are targeting leaders in their sector is more likely to pause and verify. The pattern is consistent. People learn and act most effectively when knowledge is delivered in context and directly tied to the action they are about to take.

Intelligence as a Decision Support Capability

Underlying all of this is a simple truth about why intelligence exists. At its core, intelligence is a decision-support capability. Military intelligence supports commanders. Business intelligence supports executives. Threat intelligence supports defenders. Competitive intelligence supports strategic planning. In every domain, intelligence does not merely inform. It improves decisions and, through them, outcomes. Intelligence has long been defined as a decision-support function intended to reduce uncertainty for decision makers (Kent, 1966; Lowenthal, 2022).

Viewed through that lens, intelligence should be judged by its contribution to the quality of decisions rather than by the quantity of intelligence produced. The questions that matter are practical. Did the intelligence reduce uncertainty for the person making the call? Did it improve their awareness of the situation? Did it enable preventive action before harm occurred? Did it influence behavior? Did it reduce risk in a way that can be observed? Each of these questions shifts attention from how much intelligence was produced to whether it was used.

That shift in emphasis is not a stylistic preference. It is essential for organizations operating in environments that grow more complex and dynamic every year. The future of intelligence will not belong to the organizations that simply collect more information. It will belong to the organizations that consistently turn information into action.

From Explanation to Anticipation

Another limitation is how intelligence programs spend their effort, and it compounds the utilization problem. Much intelligence work focuses on explaining what already happened. After an incident, the organization reconstructs the sequence of events, identifies the root cause, and documents the lessons. This explanatory work is valuable and necessary, and mature programs do it well. But it is, by its nature, backward-looking. It arrives after the outcome it describes, when the decisions that might have changed that outcome have already been made.

The decisions that matter most, however, happen before the event, not after. By the time an organization is explaining how a fraud succeeded or how an intrusion unfolded, the loss has already occurred. What the organization most needed, then, was not an explanation of what happened but an indication of what was likely to happen next. The practical question facing a defender, an executive, or an analyst is almost always forward-looking: what is coming, who will it affect, and what should be done before it arrives.

Intelligence oriented toward that question is anticipatory intelligence. Where explanatory intelligence interprets the past, and current intelligence describes the present, anticipatory intelligence focuses on conditions that are forming, indicators that suggest where events are heading, and threats that have not yet materialized. Its aim is not to narrate but to forewarn, to give the organization enough notice to act while action is still possible. This orientation does not replace explanation or current reporting. It sits alongside them and changes their purpose, treating every past incident and present signal as evidence about what may come rather than as a record to be filed.

Reframing intelligence around anticipation sharpens the utilization problem in a useful way. Utilization is not merely delivering what the organization knows. It means delivering it before

the decision, not after the loss. That makes timing an intelligence function as consequential as collection or analysis. An accurate assessment that arrives after the decision has been made has the same operational value as no assessment at all.

Anticipation and timing, then, are not refinements to be added once the basics are handled. They are central to whether intelligence is used at all. A program that explains the past well but cannot tell its people what is coming, and cannot reach them before they must act, has not solved the problem this part has described. Recasting the challenge in these terms, around knowing sooner and delivering earlier, is what points the way out of the utilization problem.

Toward Intelligence Utilization

The Intelligence Utilization Problem reveals a fundamental limitation of the traditional model. Collection alone is insufficient. Analysis alone is insufficient. Dissemination alone is insufficient. Each is necessary, yet none is sufficient, because each falls short of the only outcome that justifies the entire enterprise: a better decision, made by a specific person at a specific moment.

Closing that gap requires more than another feed or another report. It requires deliberate mechanisms that ensure knowledge reaches the right individuals at the precise moment it becomes operationally relevant. It also requires systems that preserve what the organization has already learned, retrieve the context of an emerging situation, identify exactly who is affected, and deliver actionable intelligence before the relevant decision must be made. In other words, it requires the organization to treat memory, anticipation, and timing as first-class intelligence functions rather than afterthoughts.

These requirements lead naturally to a broader discussion of organizational learning, organizational memory, and the processes that translate knowledge into action. They are precisely the problems the OLAA Model was designed to address. By integrating observation, learning, anticipation, and action into a continuous adaptive cycle anchored in organizational memory, OLAA offers a way to move beyond intelligence production and toward intelligence utilization. Part II introduces the model in detail and examines how organizational memory serves as the foundation for anticipatory and Just-in-Time Intelligence.

Part II: The OLAA Model and Organizational Memory

From Linear Cycle to Continuous Adaptation

The traditional intelligence cycle is, at heart, a production line. It begins with a requirement, moves material through collection and analysis, ends with a disseminated product, and loops back through feedback to begin again. That linearity is exactly what makes the cycle so effective at producing intelligence and so poor at guaranteeing that intelligence is used. A production line is optimized to finish a product, not to confirm that the product changed someone's behavior.

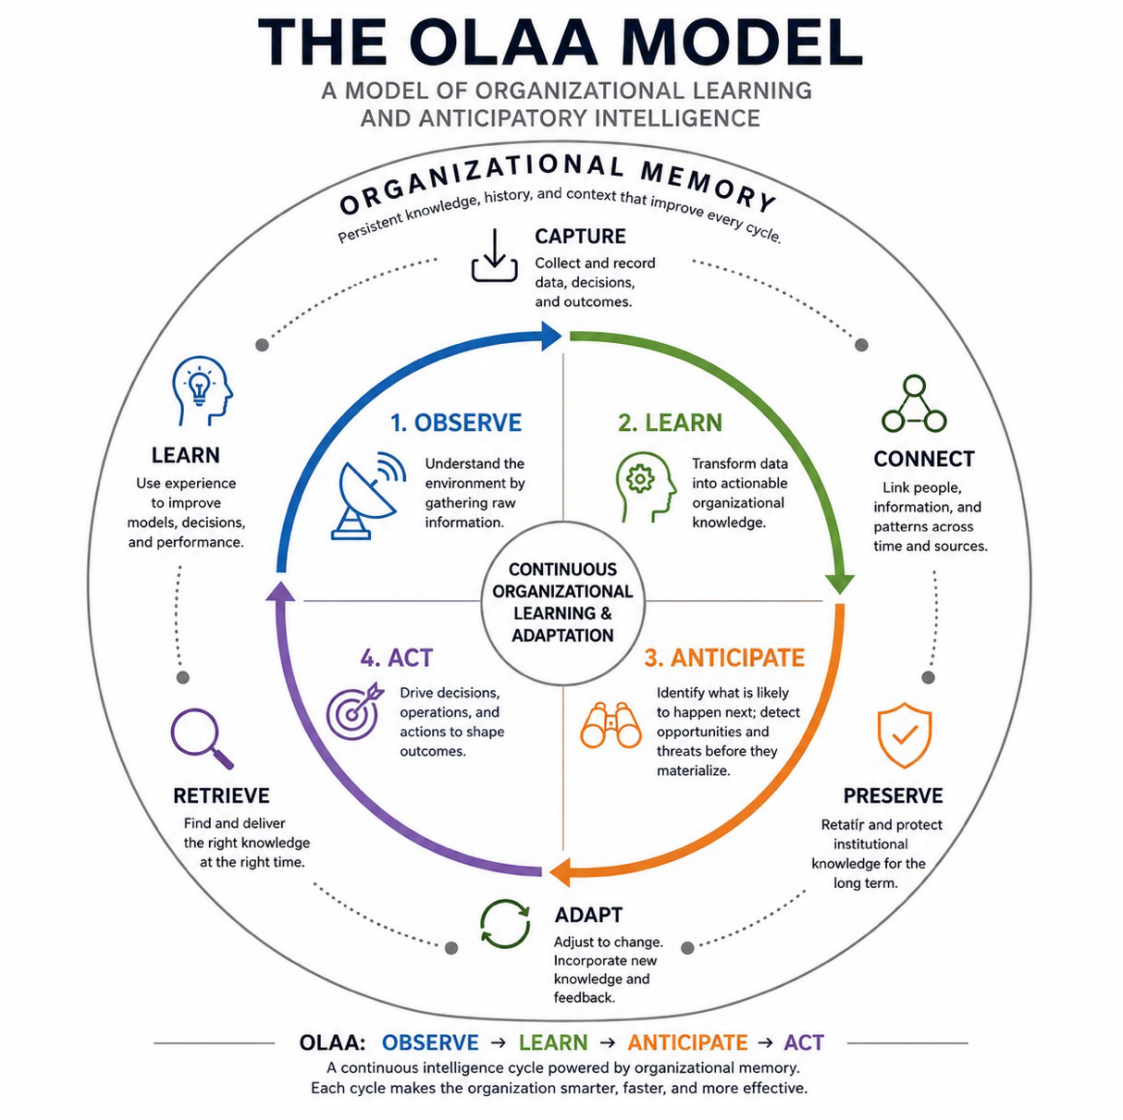


Image 2: The OLAA Model

The OLAA Model reframes the problem. Rather than a line that ends at dissemination, it describes a continuous, adaptive loop composed of four mutually reinforcing functions: Observation, Learning, Anticipation, and Action. These four are not sequential stages to be completed and left behind. They run concurrently, feed one another, and all draw on and contribute to a shared substrate at the center of the model: organizational memory. Where the classic cycle asks how to produce intelligence, OLAA asks how an organization can continuously sense, remember, anticipate, and act, so that the right knowledge reaches the right person at the right moment as a matter of routine rather than luck. The model draws on concepts from organizational learning, systems thinking, and adaptive decision cycles (Boyd, 1987; Senge, 2006; Argyris & Schön, 1978).

Organizational Memory as the Foundation

Everything in OLAA rests on organizational memory, and it is worth being precise about what that term means. Organizational memory is the durable, retrievable store of what an organization has observed, learned, decided, and experienced, held in a form that outlasts any individual employee and can be surfaced on demand. It includes prior incidents and how they were handled, the threats that have targeted the organization and its sector, the decisions that have been made and their outcomes, and the mapping of which people and assets are exposed to which risks. Knowledge management is therefore not a peripheral nicety in this model. It is the enabling prerequisite. Without a functioning memory, an organization confronts every situation as if for the first time.

The cost of weak memory is exactly the utilization failure described in Part I. When knowledge lives only in individual inboxes, in the heads of departing analysts, or in static reports that no one can find under pressure, the organization cannot reliably connect the present to lessons it already paid to learn. Memory is the substrate from which the other four functions draw and into which they feed back. Observation deposits new signals into it. Learning structures and enriches it. Anticipation queries it to forecast what is coming. Action draws on it to deliver context and writes back into it whether the action worked. Treated this way, memory stops being an archive and becomes the operational core of the intelligence enterprise.

Observation

Observation is the model's continuous sensing function and is deliberately broader than traditional collection. Collection, in the classic sense, points outward at adversaries, markets, and the external environment. Observation includes all of that but also points inward at the organization's telemetry, its workflows, and, crucially, the decisions its people are about to make. It watches external threat reporting, internal security events, behavioral signals, and the operational rhythm of the business simultaneously.

This inward dimension is what distinguishes OLAA from models that treat the organization as a passive recipient of finished intelligence. To deliver intelligence at the moment of decision, the system must first detect that a decision is approaching. Observing that a finance team is entering a high-volume payment window, that an engineering group is integrating a new dependency, or that an executive is traveling to a high-risk region is as much an intelligence input as any external feed. In OLAA, observation is the practice of continuously watching both the threat and decision landscapes and routing what it sees into organizational memory.

Learning

Learning converts raw observations into retained, structured knowledge. Signals, on their own, decay quickly and rarely transfer from one person to the next. Learning captures them, tags them by the dimensions that make them retrievable later (by role, asset, decision type, threat, and sector), and encodes them into organizational memory so they persist beyond any single analyst or incident. It is the difference between an organization that survives a phishing campaign and forgets it, and one that emerges from the same campaign with a durable, queryable understanding of who was targeted, what worked, and what to watch for next time.

Learning also enriches memory over time. Each incident, each near miss, and each decision outcome refines the organization's model of its risk surface and of which people face which exposures. When done well, learning steadily improves the precision with which later functions can operate, because anticipation and action are only as good as the memory they draw on. In this sense, learning is the function that compounds. An organization that learns deliberately gets measurably better at the rest of OLAA every time it runs the loop. Organizations that systematically capture and apply lessons learned can improve performance over time through continuous learning cycles (Garvin, 1993).

Anticipation

Anticipation combines memory and current observation to look ahead. Drawing on the structured knowledge accumulated through learning and the live signals arriving through observation, anticipation forecasts which decisions are coming, who will face them, and what those people will need to decide well. It turns a static archive into a predictive capability. Instead of waiting for a recipient to come looking for relevant intelligence, the organization determines in advance who will need what.

This is the hinge on which the entire model turns, because anticipation enables pull-based, threat-triggered delivery. Push-based programs broadcast on a calendar and hope relevance lands by coincidence. Anticipation replaces coincidence with prediction. By

matching an emerging threat against the memory of who is exposed and the observation of who is about to act, the system can identify the small set of people for whom a given piece of intelligence is genuinely and immediately operational. Everything downstream depends on getting this targeting right.

Action

Action is the function the entire model exists to serve, and it is the one form of intelligence that traditional approaches so often neglect. Action is the targeted, contextual delivery of intelligence to a specific decision maker at the precise moment it can change their behavior. It is neither a broadcast nor a report filed for the record. It is the right knowledge, tailored to the recipient's role, carrying the recommended response, arriving immediately before the decision it is meant to inform.

Action does not end at delivery. Because OLAA is a loop, action also observes its own effects. Did the recipient open the alert? Did they take the recommended step? Did the behavior change? Did the risk decline? Those answers flow directly back into observation and learning, enriching organizational memory and sharpening the next cycle of anticipation. This feedback enables the system to measure utilization rather than mere dissemination, and it enables OLAA to improve continuously rather than simply repeat. Action both closes the loop and reopens it.

The Continuous Adaptive Loop

Viewed as a whole, the four functions form a self-reinforcing cycle centered on organizational memory. Observation senses the threat and decision landscapes and records what it finds in memory. Learning structures those observations into durable, retrievable knowledge. Anticipation queries that knowledge against live signals to predict who is about to need what. Action delivers the resulting intelligence to the decision point and observes whether it worked, feeding the outcome back into memory and beginning the cycle anew.

Because memory sits at the center rather than at the end, nothing the organization learns is wasted or discarded. Each pass through the loop leaves memory richer and the next pass sharper. The model is adaptive in the literal sense: it changes its behavior in response to what it observes about its results. This is the structural answer to the utilization problem posed in Part I. A linear cycle that ends at dissemination has no mechanism to determine whether anything changed. A continuous loop anchored in memory makes that knowledge the engine of its own improvement.

Anticipation as a Core Intelligence Function

Of the four functions in the loop, anticipation deserves the fullest treatment, because it gives the model its forward orientation and is the one intelligence practice has found hardest to operationalize. The brief description above presented anticipation as the hinge of OLAA. It is worth examining more closely as an intelligence function, because the capacity to anticipate is what separates a program that reacts from one that gets ahead.

Anticipation has a long lineage in the profession under the name of warning intelligence. Its purpose is to reduce surprise, to give decision makers notice of developments before they arrive rather than explanations after they have passed. Cynthia Grabo, whose work remains the definitive treatment, framed warning as a distinct discipline organized around the patient search for the signs that precede an event (Grabo, 2004). Sherman Kent earlier located the entire value of intelligence in the same place: reducing the uncertainty a decision maker faces before acting (Kent, 1966). Warning is the part of intelligence that looks forward on purpose.

The mechanics of warning rest on indicators. An indicator is an observable whose presence suggests that a particular development is becoming more likely: a change in an adversary's behavior, a shift in a pattern of activity, a precursor that has historically preceded a given kind of event. Warning is the judgment, made against a standing set of indicators, that enough of them have appeared, or that the right combination has appeared, to justify alerting someone. The distinction between an indicator and a warning matters. Indicators accumulate quietly and continuously, while a warning is the deliberate act of declaring that the accumulation now demands attention. Building this capacity means deciding in advance which indicators matter and which patterns among them should cross the threshold from monitoring to warning.

The difficulty is that the earliest and most valuable indicators are usually the faintest. By the time a threat is unambiguous, the opportunity to act ahead of it has often passed. Anticipation therefore depends on detecting weak signals, the ambiguous and low-confidence hints that something may be forming, and on the pattern recognition that lets an experienced practitioner connect a few scattered signals into a recognizable shape. Much of this judgment is tacit, and it is vulnerable to the cognitive biases that Richards Heuer cataloged: the tendency to see the expected pattern, to discount evidence that does not fit, and to anchor on a first impression (Heuer, 1999). Good anticipation is as much a discipline of guarding against these failures as it is a matter of noticing the signals in the first place.

Anticipation also requires honesty about what it can and cannot deliver. It deals in probabilities, not certainties. A forecast is not a claim that an event will occur. It is a calibrated statement of how likely it is, given what is currently known. The discipline lies in

expressing that likelihood clearly and resisting both false precision and false confidence. An organization that demands certainty before it will act will always act too late, because certainty, in anticipation, arrives only after the fact. The aim is not to be right about the future in the way a prophet is right, but to be calibrated about it in the way a forecaster is, and to act sensibly on probabilities while they still allow a choice.

The payoff is a decision advantage. John Boyd's account of competition as a contest of observation, orientation, decision, and action captured the underlying logic: the side that cycles through that loop faster, and that anticipates rather than merely reacts, seizes the initiative and forces its opponent to respond (Boyd, 1987). Anticipation buys time, and time is the scarcest resource in decision-making. An organization that sees a condition forming a day, an hour, or even a few minutes before it matters has room to verify, to prepare, and to act deliberately rather than under duress. That margin, created by anticipation, is the advantage the whole model is built to produce.

This is why OLAA elevates anticipation from an implicit hope to an explicit function. A model that only observes, learns, and acts would still be reacting, just with better records. By making anticipation a stage, fed by memory and live observation, OLAA commits the organization to asking continuously what is coming and who will be affected. Translating that forward-looking question into delivery, getting the right warning to the right person at the right moment, is the work of Just-in-Time Intelligence, which the next section introduces.

From OLAA to Just-in-Time Intelligence

OLAA describes how an organization can sense, remember, anticipate, and act. Just-in-Time Intelligence is the operational output of running that loop well. It is the disciplined practice of delivering the right intelligence to the right person at the exact moment it becomes relevant to a decision, not before. While conventional awareness programs are push-based, broadcasting information on a fixed calendar in the hope that some of it will matter later, Just-in-Time Intelligence is pull-based and threat-triggered. It is summoned by an emerging condition and routed to a specific, exposed decision-maker, with knowledge management as the enabling prerequisite that makes targeting possible.

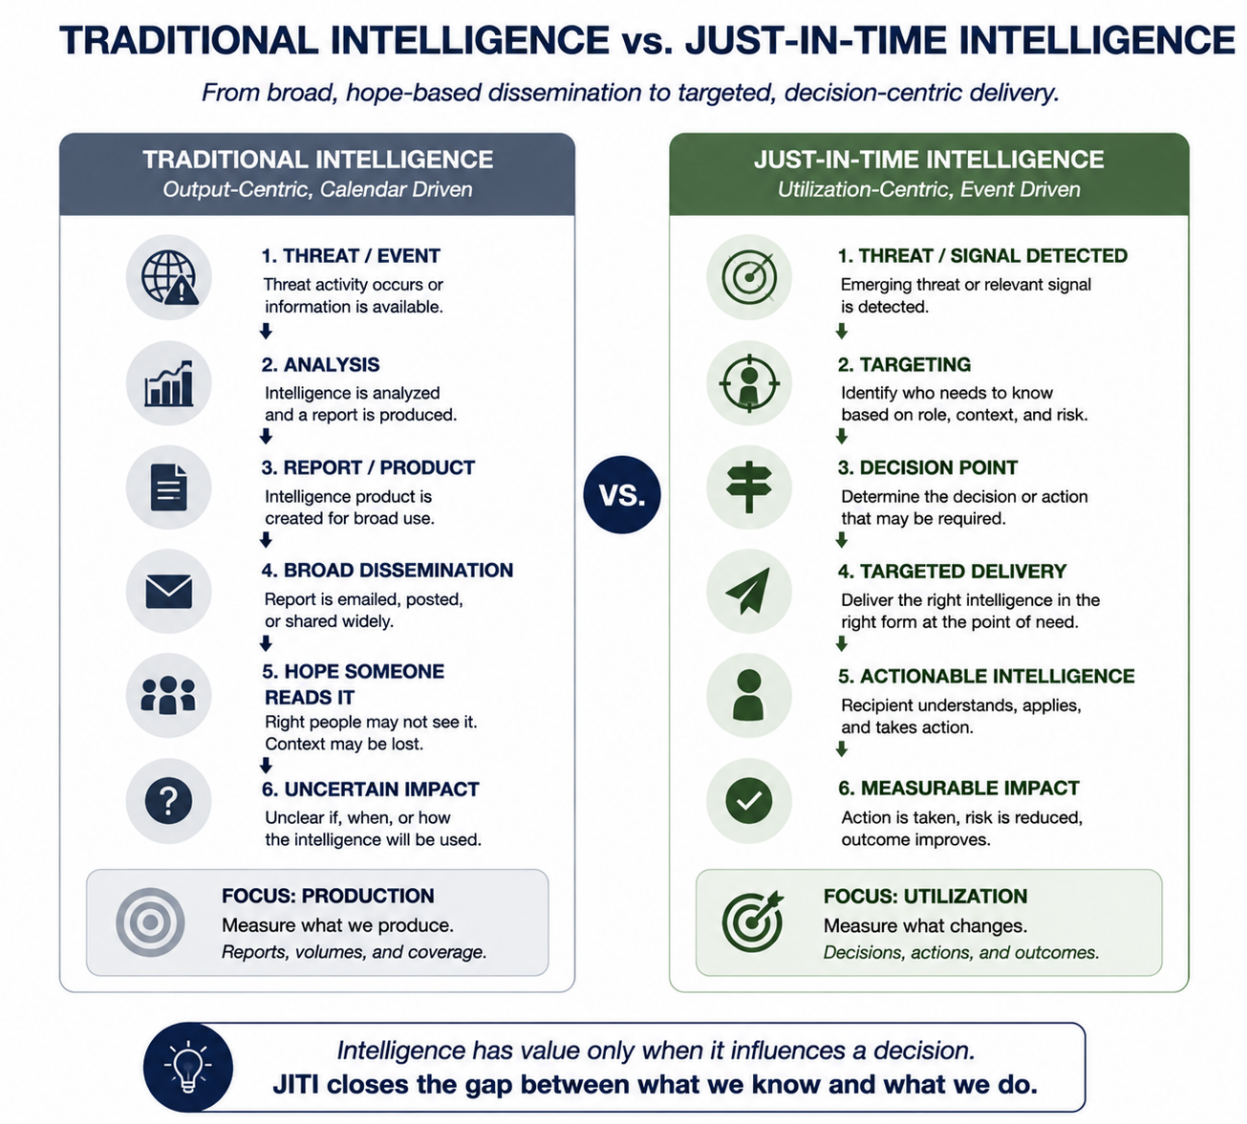


Image 3: Traditional Intel vs. JITI

The contrast with calendar-driven training is stark. Annual training assumes that knowledge delivered today will be retrieved, unaided, months later under pressure. Just-in-Time Intelligence makes no such assumption. It recognizes that human attention is scarce and human memory is imperfect, and it compensates by moving delivery to the moment of need. It does not try to make people remember everything. It makes the organization remember everything through its memory, then delivers the relevant fragment precisely when a decision is at hand. In doing so, it converts the abstract goal of intelligence utilization into a concrete, repeatable operating practice.

OLAA in Practice: A Worked Example

Consider a business email compromise campaign moving through an organization that operates the OLAA loop. The contrast with the annual-training scenario from Part I is instructive because the underlying threat is the same and only the operating model has changed.

Observation registers a threat from two directions at once. An external signal arrives via a sector ISAC describing an active wire-fraud campaign targeting firms in the organization's industry, while an internal signal shows elevated payment activity by the finance team tied to quarter-end vendor settlements. Neither signal is remarkable on its own. Together, they indicate a live and rising exposure.

Learning has already done its work. Organizational memory records who in finance is authorized to approve wire transfers, which vendors and payment patterns are normal, and how a similar campaign was handled eighteen months earlier. That structured history enables the present signals to be interpreted quickly rather than investigated from scratch.

Anticipation joins these threads. By cross-checking ISAC reporting with the memory of who approves wires and the observation that quarter-end approvals are imminent, the system identifies the handful of approvers who will face elevated risk over the next several days. It does not warn the whole company. It predicts the small set of people for whom this intelligence will soon be operational.

Action delivers to exactly those people. Each approver receives a brief, role-specific alert immediately before the high-risk window: a wire-fraud campaign is active in the sector; here are the indicators that distinguish it; and here is the verification step to complete before releasing any funds. When a fraudulent request arrives, disguised as a legitimate vendor payment, the approver is primed to pause and verify rather than process on autopilot. The action then observes its own result, noting whether the alert was opened, whether the verification step was completed, and whether the fraudulent transfer was caught, and writes that outcome back into memory, where it sharpens the organization's response to the next campaign.

The same threat that defeated the annual-training organization is intercepted here, not because the people were smarter or the analysis better, but because the intelligence arrived in context, at the decision point, and was directed to the precise individuals who could act on it. That is Just-in-Time Intelligence, and it is what the OLAA loop is designed to produce.

Summary and What Comes Next

Part I argued that the central failure of modern intelligence is not collection, analysis, or warning, but utilization: the persistent gap between what an organization knows and what it does. Part II has offered a structural response. The OLAA Model reframes intelligence as a continuous, adaptive loop of observation, learning, anticipation, and action, anchored in organizational memory. Just-in-Time Intelligence is the operating practice that the loop is designed to deliver. Together, they shift the measure of success from intelligence produced to intelligence used.

The sections that follow build on this foundation. They examine the architecture needed to operate OLAA at enterprise scale, the knowledge-management practices that make organizational memory durable and retrievable, the methods for measuring utilization rather than output, and the operational patterns that turn Just-in-Time Intelligence from a principle into a daily capability. The destination, throughout, remains the same. The future of intelligence belongs not to the organizations that collect the most, but to those that most consistently turn what they know into what they do.

Part III: Building Organizational Memory

Part II placed organizational memory at the center of the OLAA loop and identified knowledge management as its enabling prerequisite. It deliberately left open the harder question of how that memory is built. That question deserves its own treatment, because the difference between organizations that achieve Just-in-Time Intelligence and those that merely aspire to it is rarely the quality of their analysts or the sophistication of their feeds. It depends on whether they have engineered a memory that their people can draw on under pressure. However, before we engineer memory, we must first understand what organizational memory is and how it is created.

Example: Vulnerability Management

A software engineer joins a project team and deploys a library version previously linked to a security incident. The organization encountered the same issue eighteen months earlier, documented lessons learned, and updated secure coding guidance. Without Organizational Memory, that knowledge remains buried in reports and meeting notes. With Organizational Memory, the engineer receives relevant guidance when the vulnerable component is selected, enabling the organization to apply a past lesson before repeating the mistake.

Beyond Information Storage

Many organizations possess vast quantities of information yet surprisingly little organizational memory. They may have extensive repositories, documented lessons learned, archived threat feeds, incident records, and retained training materials, yet still struggle to retrieve the knowledge that matters when a decision must be made. The volume of retained information can be impressive, but volume is not the same as recall.

When a new threat emerges, decision-makers often find that the organization has forgotten what it already knew. The knowledge was somewhere in the enterprise, but it was neither available, understandable, nor actionable when needed. This highlights a distinction that many programs overlook: storage preserves information, while memory preserves understanding.

An organization can hold terabytes of data and thousands of reports yet still be unable to retrieve the one piece of knowledge that matters during an active decision. Capacity is not recall, and a full archive can coexist with an empty memory. For intelligence to create value,

knowledge must be available, understandable, contextualized, and actionable. A document that cannot be found, interpreted, or applied contributes nothing to the decision at hand.

Organizational Memory therefore serves not retention for its own sake. Its purpose is to make knowledge available now of decision. Storage is a warehouse; memory is a capability. This paper concerns capability.

The Organizational Learning Problem

Organizations often repeat mistakes they have already made not because they lack intelligence, but because they do not retain the knowledge gained from prior experience in a form that can guide future action. A lesson may be learned during an incident, captured briefly in a meeting, or recorded in a report, yet it still fails to become part of the organization's usable memory. When the same conditions recur months or years later, the organization may respond as though the earlier experience never occurred.

The normal flow of organizational life accelerates this forgetting. Personnel change positions, teams reorganize, projects conclude, and temporary working groups disband. Lessons learned are often filed after an event, but once the immediate pressure passes, people seldom revisit those lessons or link them to the next decision that might benefit from them. Over time, institutional knowledge fragments across documents, databases, email threads, ticketing systems, shared drives, and the memories of individuals who may no longer be present or reachable when their experience is needed.

The result is organizational forgetting. The organization may continue to generate new knowledge every day, but it simultaneously loses knowledge through turnover, fragmentation, weak documentation, and poor retrieval. Researchers in organizational learning have described this problem for decades (Argyris & Schön, 1978; Senge, 1990). Individuals learn through experience, but organizations learn only when they capture, share, retain, and apply that experience to future decisions. Without those mechanisms, knowledge remains trapped inside individuals and leaves when they leave.

This creates a persistent paradox. Organizations generate knowledge continuously, and they also lose knowledge continuously. The two processes occur simultaneously, but in many organizations the second outpaces the first. The challenge, then, is not simply generating more intelligence. Most organizations already generate more information, assessments, and lessons than they can use. The harder challenge is retaining what has been learned, making it retrievable, and applying it when a similar decision or risk arises again.

Organizational Memory is the answer to that challenge. It is the deliberate mechanism by which an organization stops forgetting what it has learned and begins to convert prior

experience into future advantage. Properly built, it lets lessons outlast the people, teams, and projects that produced them, so the organization does not have to rediscover the same truths each time conditions recur.

Knowledge Creation and the SECI Model

Nonaka and Takeuchi describe organizational knowledge creation as a continuous cycle known as the SECI model: Socialization, Externalization, Combination, and Internalization (Nonaka & Takeuchi, 1995). The model explains how knowledge moves through an organization, beginning as individual experience and gradually becoming shared, structured, and reusable institutional knowledge. Each phase moves knowledge further from dependence on a single person and closer to becoming part of the organization itself.

In the socialization phase, people share knowledge through experience, observation, mentoring, discussion, and collaboration. An analyst may gain insight during an investigation by noticing a pattern, recognizing an adversary technique, or connecting details that are not obvious in isolation. At this stage, the knowledge may still be largely tacit. It may reside in the analyst's judgment, intuition, and accumulated experience rather than in a document or formal assessment.

Externalization occurs when the analyst articulates and documents that insight. The analyst explains what was observed, why it mattered, how they reached the conclusion, and what others should watch for in the future. This is the point at which private knowledge begins to become organizational knowledge. The insight's value no longer depends entirely on the analyst remembering it or being available to explain it later, because part of the reasoning is captured in a form others can retrieve and use.

Combination then connects the documented insight to the organization's existing knowledge. A single observation may be linked to prior incidents, threat reports, exposure maps, procedures, investigations, or decisions and outcomes. In this phase, isolated knowledge becomes more useful because it is contextualized. The organization does not merely store the analyst's note; it integrates that note into a larger body of understanding so future users can see how it relates to what the organization has already learned.

Internalization completes the cycle when future personnel absorb that combined knowledge and apply it to new situations. A later analyst, investigator, approver, or decision maker does not need to recreate the original insight from scratch. They can draw on accumulated knowledge, recognize similar patterns more quickly, and make better decisions because the earlier experience has been converted into usable organizational capability.

This cycle is the engine beneath Organizational Memory. Every investigation, assessment, incident, decision, and outcome creates knowledge. The question is whether that knowledge remains private, temporary, and fragile, or is captured, connected, retained, and made available to the people who will need it next. The SECI cycle explains how raw experience becomes a shared asset rather than a private one.

Organizational memory keeps this cycle turning. It provides the durable structure that keeps knowledge created in one investigation, incident, or decision available after the immediate event has passed. Without that structure, the SECI cycle breaks down: insight remains trapped in individuals, documentation remains disconnected, and future personnel are forced to relearn what the organization already knew. With it, knowledge continues to circulate, deepen, and compound across the organization.

The result is an organization that grows more capable over time because its knowledge no longer depends entirely on the people who first created it. Individual insight becomes shared understanding. Shared understanding becomes institutional memory. Institutional memory becomes the basis for better anticipation and action. In this way, the SECI model provides the knowledge-creation logic for Organizational Memory and explains why memory must be actively maintained rather than passively accumulated.

Tacit and Explicit Knowledge

Much of the most valuable knowledge in intelligence work is tacit rather than explicit. It does not begin as a report, procedure, database entry, or finished assessment. It lives within people as judgment, intuition, pattern recognition, and accumulated experience. Practitioners develop it through repeated exposure to ambiguous situations, imperfect evidence, and decisions made under pressure. Because it is learned through practice, it is often difficult to write down, explain, and even more difficult to transfer from one person to another.

Tacit knowledge encompasses the professional habits and decision heuristics experienced practitioners apply with little conscious effort. An experienced analyst may notice an indicator that a newer analyst reads straight past because the pattern resembles something from a prior investigation. A seasoned investigator may sense that an account does not behave like a legitimate account, even before assembling all supporting facts. An executive may weigh a decision against years of outcomes, comparable cases, and consequences that they could not fully reconstruct if asked to explain the judgment step by step. In each case, the knowledge is real and often decisive, but it is fragile because it is bound to individuals rather than to the institution.

By contrast, explicit knowledge is the knowledge the organization has already written down and formalized. It appears in reports, records, procedures, playbooks, databases, tickets, assessments, and documented lessons learned. Explicit knowledge is easier to store, search, cite, and share than tacit knowledge, but that does not mean it automatically becomes useful memory. A procedure that no one consults, a report that cannot be found, or a database record that lacks context can be preserved yet still fail to inform a decision. Explicit knowledge is easy to store, but it is also easy to lose track of.

Organizational Memory must hold both kinds of knowledge. It must retain the explicit record so the organization can recover what was observed, assessed, decided, and documented. At the same time, it must continuously convert tacit expertise into a form the institution can retain (Nonaka & Takeuchi, 1995). That conversion does not require pretending that all judgment can be fully captured. Some expert intuition will always resist complete documentation. But the reasoning, indicators, decision cues, exceptions, and lessons surrounding that expertise can often be made explicit enough to guide others later.

The goal is to turn individual expertise into institutional capability. When tacit knowledge remains private, the organization depends on the continued presence of the person who holds it. When that person leaves, changes roles, or becomes unavailable, the organization risks losing the judgment that made prior decisions effective. By capturing expert reasoning, linking it to incidents and outcomes, and making it retrievable for future decisions, Organizational Memory reduces that dependency. It enables the organization to learn from its people without requiring every new employee, analyst, investigator, or leader to relearn the same lessons from the beginning.

Memory Decay and Organizational Forgetting

Just as individuals forget, organizations forget. The difference is that organizational forgetting is often harder to see because it does not occur within a single mind. It occurs across teams, systems, documents, processes, and the movement of people through the enterprise. A person may recognize that they no longer remember a fact, but an organization rarely notices the moment a lesson disappears from practical use. The knowledge may still exist somewhere in a file, ticket, after-action report, or archive, yet it has been functionally lost if no one can retrieve it when the next decision arrives.

This decay occurs through predictable mechanisms. Personnel turnover removes tacit knowledge when experienced employees leave or change roles. Organizational restructuring disrupts the relationships and routines through which knowledge once flowed. Technology migrations leave records stranded in old systems or strip them of useful context. Inadequate documentation captures conclusions without preserving the reasoning behind them. Poor retrieval leaves lessons technically preserved but practically unreachable. Each

mechanism is familiar on its own, but together they steadily erode the organization's ability to use what it once knew.

The decay is rarely noticed until it becomes costly. It surfaces during incident response, when the organization rediscovers a lesson it had already paid to learn. An investigation repeats work completed two years earlier because the earlier analysis cannot be found or trusted. A team rebuilds an assessment that already exists in a file no one knows to search. A decision is made without awareness that a nearly identical decision was made before, or that its outcome should have shaped the current choice. In each case, the organization does not merely suffer from missing information. It suffers from inaccessible memory.

The cost of this forgetting is real. Resources are wasted as teams repeat work that should have been reusable. Risk rises because prior warnings, indicators, and decision outcomes do not inform current decisions. Decision quality declines because leaders and practitioners must act without the benefit of the organization's experience. The most damaging aspect is that the failure is quiet. Nothing signals that the organization has forgotten. It simply behaves as though the lesson never happened.

Organizational Memory exists to slow this decay and, where possible, reverse it. It preserves institutional knowledge and keeps it accessible over time, through reorganizations, and despite departures that would otherwise erase it (Walsh & Ungson, 1991). In doing so, it gives the organization a way to maintain continuity even as people, systems, and structures change. Memory does not eliminate forgetting entirely, but it reduces the likelihood that hard-won lessons vanish before they can guide the next decision.

At the same time, memory that is preserved but never maintained decays in a different, more dangerous way. Stale records, outdated exposure maps, obsolete procedures, and unrevised threat profiles can create the illusion of knowledge while leading the organization to the wrong conclusion. This is why organizational memory cannot be treated as a one-time act of preservation. It must be continuously curated, refreshed, and tested against current reality, a point this paper revisits in the discussion of the half-life of knowledge.

Organizational Memory as a Strategic Asset

Traditional intelligence programs are often built to produce reports, whereas Organizational Memory is built to accumulate advantage. The distinction matters because a report is usually treated as a finished product, while memory treats every product, decision, and outcome as an input into something larger. Each completed investigation, assessment, warning, decision, and result adds another layer to the organization's institutional knowledge. Over time, those layers become more than a record of activity. They become a

strategic asset that enhances the organization's ability to understand what is happening, recognize what matters, and respond with greater speed and precision.

Organizations that preserve and retrieve knowledge effectively improve faster than those forced to relearn the same lessons repeatedly. The advantage compounds because each cycle of observation, learning, anticipation, and action leaves behind useful context for the next. Each decision sharpens the organization's understanding of future decisions. Each outcome improves its ability to anticipate similar outcomes. In this way, memory turns experience into cumulative advantage rather than isolated learning episodes that fade after the immediate pressure has passed.

This advantage is difficult for competitors to replicate. A competitor can purchase the same intelligence feeds, subscribe to the same reporting services, and hire analysts with comparable credentials. What it cannot easily buy is the accumulated memory of how this particular organization has been attacked, the decisions it made under pressure, which responses worked, which failed, and the consequences that followed. That knowledge is specific to the organization's history, exposure, culture, systems, and operating environment. It is earned through experience and becomes valuable only when captured and made retrievable.

For that reason, Organizational Memory should not be understood as a passive record of the past. Properly maintained, it is a durable source of present advantage. It enables the organization to recognize recurring patterns sooner, avoid mistakes it has already made, reuse insights it has already paid to develop, and apply prior experience to decisions that are still unfolding. The organization that remembers well does not merely know more; it acts with the benefit of everything it has already learned.

Organizational Memory and the OLAA Model

Within the OLAA Model, Organizational Memory is the connective tissue linking past experience to future action. It enables the loop to function as more than a sequence of activities. Observation, Learning, Anticipation, and Action each play a distinct role, but they become mutually reinforcing only when their results are captured, retained, and made available to the next cycle. Without Organizational Memory, each pass through the loop would start too close to zero, depending on what current personnel remembers or what current systems surface. With memory at the center, each cycle begins with the advantage of everything the organization has already observed, learned, anticipated, and done.

Observation generates information by sensing the external threat environment, the internal operating environment, and the decision landscape where action is likely. Learning then converts that information into understanding by structuring, interpreting, and connecting it

to what the organization already knows. Anticipation projects that understanding forward by asking who is likely to need which knowledge, under what conditions, and at what moment. Action produces decisions and outcomes by delivering intelligence to the decision point and observing whether behavior changes. Organizational Memory captures the result of each cycle and holds it for the next, so the organization does not merely move through the loop; it improves with each iteration.

Example: Executive Travel

A senior executive is scheduled to travel to a country experiencing civil unrest. Intelligence reporting on demonstrations and transportation disruptions has existed for weeks, but it becomes relevant only when travel plans are finalized. OLAA identifies the event, links it to the executive's itinerary, and delivers the information before departure, allowing travel adjustments to be made before the disruption occurs.

This creates a reinforcing cycle of organizational learning. The organization does not merely react to events as they occur. It learns from those events, retains what it learns, and applies that memory to the next decision. Each investigation, warning, delivery, response, and outcome adds to memory, and memory feeds the cycle that follows. Observation grows sharper because the organization recalls which signals mattered last time and which proved irrelevant. Learning grows more disciplined because new information can be compared with prior cases rather than interpreted in isolation. Anticipation grows more precise because memory holds who was exposed, how the threat behaved before, and which decisions were most vulnerable. Action grows more effective because it is informed not only by current intelligence but also by the accumulated experience of prior cycles.

In this way, Organizational Memory transforms intelligence from a reporting activity into a mechanism for continuous organizational learning. The organization no longer treats each report, alert, or incident as an isolated event. Instead, it treats each as a contribution to a growing body of knowledge that can improve future sensing, anticipation, and action. This is the practical significance of placing memory at the center of OLAA: it gives the organization a way to compound learning over time. Intelligence becomes more than producing assessments. It becomes the operating system through which the organization remembers, adapts, and acts with increasing accuracy.

Organizational Memory and Just-in-Time Intelligence

Just-in-Time Intelligence depends on Organizational Memory because timely delivery is valuable only when it includes the context needed for action. A warning delivered to the right person at the right moment is important, but timing alone does not make it intelligence.

Without memory, JITI collapses into targeted alerting: the right person may receive a message, but the message stands alone, disconnected from prior incidents, institutional experience, known exposures, and the reasoning that would help the recipient understand what to do next.

With Organizational Memory, JITI becomes more than an alerting system. It delivers contextual intelligence. When a threat emerges, the system does not merely push a warning into an inbox and hope the recipient interprets it correctly. It can assemble the surrounding knowledge that makes the warning meaningful: previous incidents of the same kind, related investigations, lessons learned, relevant procedures, prior decisions, and the outcomes that followed. That accumulated context turns the warning from a notification into a usable decision aid.

This distinction is central to the model. The recipient does not simply receive information. Instead, the recipient receives organizational knowledge assembled for the specific decision at hand. For example, a finance approver does not merely learn that a wire-fraud campaign exists; they receive the relevant history of similar attempts, the verification steps that worked previously, the indicators that distinguished fraudulent requests from legitimate ones, and the organizational procedure to follow before releasing funds. The intelligence arrives not as a detached fact but as a compressed expression of what the organization has already learned.

The objective, therefore, is not merely to inform. It is to enable informed action. Information can make a person aware of a risk, but contextual intelligence helps them recognize the risk in their own situation and take the appropriate step before harm occurs. Organizational Memory provides the context, and Just-in-Time Intelligence delivers it. Together, they convert what the organization knows into action when it matters most.

The Goal: A Learning Organization

The ultimate objective of Organizational Memory is not preservation for its own sake. Preservation matters only because it enables something more important: learning. A learning organization does more than store what has happened. It captures experience, retains knowledge, applies its lessons, and improves the quality of its decisions over time (Senge, 1990). In this sense, memory is not an archive of the past; it is the mechanism by which the organization converts experience into future capability.

Organizational Memory provides the mechanism that preserves knowledge; OLAA provides the process that keeps knowledge moving through observation, learning, anticipation, and action; and Just-in-Time Intelligence provides the delivery model that brings that knowledge to the decision point. Together, they produce an organization that learns, anticipates, and

acts with increasing speed, precision, and effect, while its competitors are still rediscovering what they already knew.

Understanding memory, however, is not the same as possessing it. An organization may acknowledge that memory matters yet still lack the structures needed to make it usable under pressure. The remainder of this section therefore shifts from what Organizational Memory is to how it must be deliberately built, structured for retrieval, and maintained over time.

Memory Must Be Engineered

The word engineered is chosen carefully. Useful memory does not accumulate on its own. Most organizations have the illusion of memory without the reality. They have full inboxes, archived reports, ticketing systems, shared drives, and chat histories dating back years. All of that is storage, and almost none of it is memory in the sense OLAA requires. The information exists, but it cannot be retrieved at the moment of need, connected to the person it concerns, or surfaced in response to an emerging threat. A pile of records that no one can query under pressure is not an asset. It is a liability dressed as an asset, because it creates the comforting but false belief that the organization remembers what it has learned.

Engineering memory means treating it as a deliberate construction rather than a byproduct of daily work. It means deciding in advance what the organization needs to remember, structuring that knowledge so it can be retrieved along the dimensions that matter, and maintaining it to keep it accurate over time. The remainder of this section examines each of those obligations in turn.

What Organizational Memory Must Hold

Before you can structure memory, you must scope it. Not everything an organization produces belongs in its operational memory, and treating memory as a place to keep everything is the surest way to render it unusable. For Just-in-Time Intelligence, organizational memory must reliably hold a small number of high-value content classes.

The first is the record of incidents and their handling. Every phishing campaign that survived, every fraud attempt that was caught or missed, and every outage and its resolution carry lessons that are valuable only if they can be recalled the next time something similar occurs. The second is the history of threats targeting the organization and its sector, including the techniques used, the indicators observed, and the actors believed responsible. The third is the record of decisions and their outcomes, because an organization that cannot remember what it decided and what those decisions produced cannot learn from them.

The fourth content class is the one most often missing and most important to this model: the exposure map. This is the standing understanding of which people and assets are exposed to which risks. It answers questions such as who is authorized to approve wire transfers, which systems hold regulated data, which teams handle external-facing communications, and which executives travel to high-risk regions. The exposure map enables anticipation by linking an emerging threat to the specific individuals it endangers. Without it, intelligence can describe a threat in detail yet still have no idea whom to warn. The final content class is tradecraft and lessons, the accumulated craft knowledge of how the organization investigates, verifies, and responds, captured in a form that does not depend on any single practitioner being present.

These content classes matter especially for warning. The indicators that preceded earlier incidents, the warnings the organization issued and whether they proved correct, the near misses caught in time, and the lessons drawn from both successes and failures are the raw material from which future warning is built. Each indicator that matters sharpens the organization's sense of what to watch for next, and each near miss reveals a pattern worth monitoring before it recurs. Warning grows more accurate the more the organization remembers, because anticipation is, in large part, the act of recognizing a pattern the organization has seen before. Memory lets the Threat Warning series improve over time rather than starting fresh with every new threat.

From Archive to Operational Memory

The decisive difference between storage and memory is retrieval. An archive holds everything and surfaces nothing on demand. Operational memory is organized, so the right fragment can be summoned at the precise moment it is needed, and that organization is not a convenience. It is the entire point. The value of memory is a direct function of its retrievability, and retrievability must be designed in, not hoped for.

Designing for retrieval means tagging knowledge along the dimensions along which it will later need to be recalled. In an intelligence context, those dimensions are reasonably consistent: by role, so knowledge can be matched to the people who hold a given responsibility; by asset, so it can be matched to the systems and data at stake; by decision type, so it can be surfaced when a particular kind of decision is approaching; by threat, so related history can be assembled when an adversary technique reappears; and by sector, so external reporting can be filtered to what is relevant to the organization's industry. Knowledge tagged along these axes can be assembled into a precise, situation-specific picture on demand. Knowledge stored without them can only be searched, which depends on the searcher already knowing what to look for and having the time to look. Under operational pressure, neither condition usually holds.

This is why full-text search, useful as it is, does not replace structured memory. Search retrieves documents that contain a phrase. Operational memory retrieves the knowledge relevant to a situation, including knowledge the searcher did not know to ask for. The distinction is the same one that runs through the whole paper. An organization that can only search is still waiting for someone to come looking. An organization with operational memory can push the right knowledge to the right person before that person knows they need it.

Knowledge Management as Intelligence Infrastructure

In most organizations, knowledge management is treated as an administrative discipline. It sits alongside records management and compliance, focusing on document retention, version control, and the orderly storage of files. That framing is not wrong, but in an intelligence context, it is dangerously incomplete. Knowledge management is not a back-office support function for the intelligence program. It is the infrastructure on which the entire program runs.

The reasoning follows directly from the OLAA loop. Observation, learning, anticipation, and action all draw on and feed organizational memory, which knowledge management builds and maintains. If the memory is shallow, fragmented, or stale, every other function inherits that deficit. The most sophisticated collection apparatus and the most rigorous analysis cannot overcome a memory that fails to retrieve the right knowledge at the right time. Underinvesting in knowledge management, therefore, sets the ceiling for the entire intelligence enterprise, no matter how much is spent above it.

Reframing knowledge management as infrastructure has practical consequences. It means the people who curate organizational memory are doing intelligence work, not clerical work. It means that investment in retrieval, tagging, and maintenance is an investment in operational capability, not overhead. And it means that an organization serious about intelligence utilization must hold its knowledge management to an operational standard, judged by whether the right knowledge reaches the right decision in time, rather than by how completely the archive has been filled.

The Half-Life of Knowledge

Memory, once built and left alone, does not remain true. Knowledge has a half-life. Threat specifics change as adversaries shift techniques. Personnel change as people move between roles and leave the organization. Exposure changes as systems are deployed and retired and responsibilities are reassigned. A memory that is not maintained drifts steadily out of alignment with reality, and a drifting memory is worse than no memory at all because it produces confident, well-targeted, and incorrect intelligence.

Consider an exposure map that still lists an approver who left the finance team six months ago, or a threat profile that describes an adversary's tradecraft as it stood a year before the adversary changed it. Anticipation built on that stale memory will warn the wrong people about the wrong indicators, with all the authority of a system that appears to know exactly what it is doing. The danger is that the output looks correct. Stale memory fails quietly, and it fails in the direction of false confidence.

For this reason, maintenance is not a phase that ends once memory has been built. It is a permanent function of the model. Memory must be curated continuously: refreshed as conditions change, versioned to preserve history without contaminating the present, and retired when knowledge is no longer valid. The feedback that flows from action back into memory, described in Part II, is part of this maintenance, but it is not sufficient on its own. Deliberate curation closes the gap between what the organization once knew and what is true now.

Capturing What Walks Out the Door

The tacit knowledge described earlier, the pattern recognition and the hard-won instinct that something does not look right, is precisely the knowledge most at risk of being lost. When an experienced analyst leaves, their explicit work products remain, but the tacit judgment that made those products possible walks out the door with them. Organizations routinely lose years of accumulated understanding this way and only discover the loss when a situation arises that the departed analyst would have recognized instantly.

One of the central tasks of knowledge management, understood as infrastructure, is to convert tacit knowledge into an explicit, retrievable form before it is lost. That means capturing not only what was decided but also why, not only the indicators that mattered but also the reasoning that identified them, and not only the outcome of an investigation but also the craft that produced it. Perfect capture is impossible; tacit knowledge resists full articulation. But the difference between an organization that attempts this capture systematically and one that does not is the difference between memory that compounds and memory that resets with every departure.

The same logic applies to the quieter problem of silos. Knowledge trapped in an individual's inbox, a single team's channel, or a personal set of notes is not organizational memory, however valuable it may be. It is private memory that the organization happens to be near. Making knowledge organizational means making it shared and discoverable beyond the person and team that created it, so that the analyst facing a novel situation can access the lesson learned by a colleague two years and two departments away. Memory that cannot cross those boundaries cannot serve an enterprise, no matter how rich it is within them.

Memory in Action: Revisiting the Wire-Fraud Case

The business email compromise example from Part II makes memory's role concrete. Recall that anticipation correctly identified the small set of wire approvers who were about to face elevated risk, and that action delivered a precise, contextual warning to exactly those people before the fraudulent request arrived. None of that was generated on the fly. All of it was retrieved.

Anticipation could identify the approvers only because the exposure map already listed who was authorized to release funds. It could distinguish a suspicious request from a routine one only because memory held the normal patterns of vendor relationships and payment behavior. It could lend the warning credibility only because memory retained the record of a similar campaign handled eighteen months earlier, including what had worked. Strip that memory away, and the very same external and internal signals arrive as undifferentiated noise. The ISAC report becomes just another bulletin, the quarter-end activity becomes routine, and the organization is left exactly where Part I found it, in possession of the relevant information yet unable to act on it.

The lesson generalizes. Anticipation and action are downstream of memory, and the quality of Just-in-Time delivery is constrained by the quality of the memory it draws on. An organization can buy better feeds and hire better analysts, but if its memory is shallow or inaccessible, the loop will continue to produce dissemination rather than utilization. Memory is not the unglamorous foundation beneath the interesting work. In this model, memory is the interesting work.

From Memory to Anticipation

With memory deliberately engineered, properly structured for retrieval, and continuously maintained, the organization is finally positioned to ask the forward-looking question on which Just-in-Time Intelligence depends. Who is about to need what, and when? Memory tells the organization what it knows and who is exposed. It does not, on its own, predict the next decision. Turning that standing knowledge into a prediction about an approaching moment is the work of anticipation and targeting, and it is the subject of Part IV. There, the focus shifts from building the foundation to using it, from remembering what the organization knows to forecasting where that knowledge will become operational.

Part IV: Anticipation and Targeting

The Targeting Problem

Part III closed with a forward-looking question: who is about to need what, and when? Anticipation is the function that answers it, and targeting is what anticipation produces. But anticipation cannot run on memory and live signals alone. It needs a third ingredient that the previous sections have not yet supplied: a structured definition of the intelligence each decision and each decision-maker actually requires. The right intelligence cannot be delivered to the right person at the right time if the organization has never defined, with any precision, what "right" means for that person and that moment. Targeting begins with requirements.

This is the unglamorous root of many utilization failures. When requirements are implicit, vague, or held only in the heads of a few analysts, delivery becomes guesswork. The organization may sense that finance needs fraud intelligence and that executives need impersonation warnings, but a vague sense of need is not a structure that can be matched to a live signal and resolved into a specific person to warn. Anticipation that is not grounded in defined requirements degrades into broadcasting, precisely the failure Part I described. Before an organization can target, it must know, in structured terms, what each part of the enterprise needs intelligence to accomplish. Requirements have historically served as the mechanism for aligning intelligence resources with decision-maker needs (Kent, 1966; JP 2-0, 2020).

Why the Series Method Exists

The Series Method begins with a single observation that most requirement frameworks overlook. Intelligence requirements are not all the same kind of thing. They exist for fundamentally different reasons. Some requirements exist to warn the organization of approaching harm. Others exist to support a specific decision a specific person will make. Others exist to interpret events and explain what they mean. Others exist to maintain a working understanding of the environment in which the organization operates. These are not variations in importance. They are differences in function.

The dominant construct for managing requirements, the priority intelligence requirement, organizes this diversity along a single axis: priority. It asks which requirements matter most and ranks them accordingly. Priority is a useful property, but it is the wrong organizing principle because it flattens requirements that serve entirely different purposes into a single ranked list. A warning requirement and a decision-support requirement may sit one above the other on a priority list while having almost nothing in common in how they are triggered,

whom they serve, or what delivery they demand. Ranking them together tells the organization which to attend to first. It says nothing about what each is for.

This is the problem the Series Method is designed to solve. The claim is not that priority requirements are written or maintained poorly. Rather, the claim is that priority is the wrong primary structure for requirements that serve different purposes. When function is obscured, targeting becomes impossible because the organization cannot route a requirement it cannot characterize. The Series Method restores function as the organizing principle. It groups requirements first by the purpose they serve, and only then addresses priority, decomposition, and collection. Purpose comes first because it determines whom a requirement serves and how its intelligence must be delivered.

The Four Functions of Intelligence

If requirements are organized by purpose, the natural next question is which purposes intelligence serves. The Series Method answers with four purposes, each corresponding to a function that intelligence organizations have always performed and each expressible as a question the organization continually asks.

The first series is Threat Warning, and its question is what is likely to happen next. These requirements are inherently anticipatory. They describe adversary activity, the indicators that precede it, and the conditions that, when they appear, should prompt a warning. Threat Warning requirements are the series most directly responsible for initiating threat-triggered delivery. When an emerging condition matches one of them, the organization has its cue to act. The focus on anticipatory indicators and emerging conditions aligns with the strategic warning tradition long established within the intelligence community (Grabo, 2004).

The second series is Decision Support, and its question is which decisions need support. These requirements tie intelligence to specific decisions that specific people will make, anchoring each requirement to a named decision and the role responsible for it. Decision Support is a series in which the abstract goal of supporting decisions becomes concrete enough to target, because it answers not only what intelligence a decision needs but also who will make that decision and, therefore, who must receive the intelligence.

The third series is Analysis, and its question is what it means. These requirements exist to interpret. They take the raw material gathered by collection and the events surfaced by observation and turn them into understanding: assessments of significance, judgments about intent, and explanations of how disparate signals fit together. Analysis requirements are what prevent the organization from mistaking information for intelligence. A warning that a condition has occurred is of limited value without the interpretation that explains why it matters and what it implies. Defining that interpretation is the work of this series.

The fourth series is Environment, and its question is what environment we are operating in. These requirements provide broad situational awareness, the wide-aperture monitoring of the surrounding landscape that captures what the narrower series are not specifically looking for. Environment requirements guard against tunnel vision by surfacing developments beyond the current set of known threats and active decisions. They feed the outer ring of observation and are often where the first faint signal of an unanticipated risk appears, before any Threat Warning requirement has been written for it.

Taken together, the four functions span the full range of work an intelligence organization performs. Threat Warning looks ahead. Decision Support serves the moment of choice. Analysis explains. Environment maintains the wider view. Every requirement an organization holds serves one of these four purposes, and the series to which a requirement belongs immediately signals how its intelligence is intended to be used. That signal is the first and most consequential act of targeting, because it is the point at which a requirement's purpose, and therefore its audience, becomes explicit.

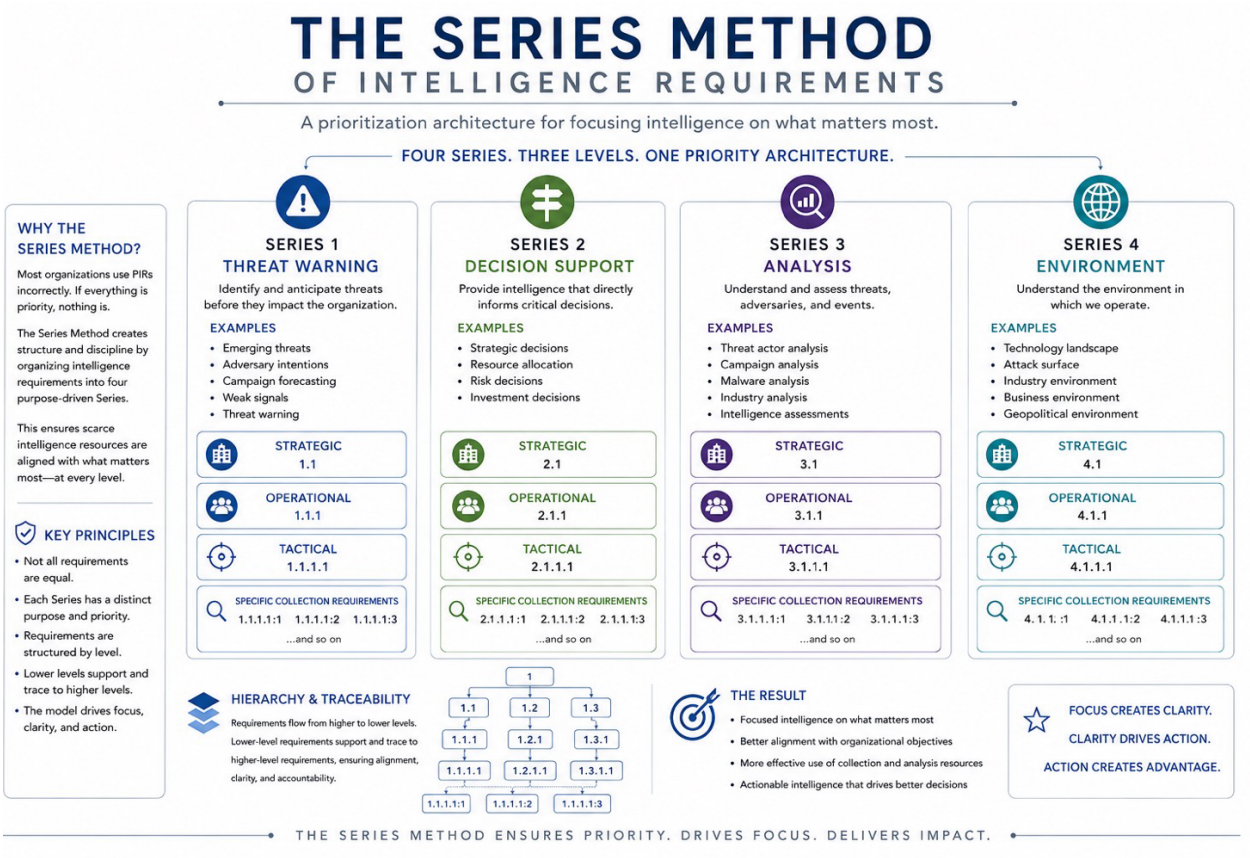


Image 4: The Series Method of Intelligence Requirements

Series 1: Threat Warning

The four functions merit closer examination, one at a time, beginning with the series the model places first. Threat Warning is Series 1 not by accident of numbering but by logic. Its purpose is to identify what is likely to happen next, and that anticipatory question is the one on which everything downstream depends. A decision cannot be supported until the organization recognizes that one is approaching. An event cannot be interpreted until it has been seen coming. Threat Warning is the series that looks forward, and placing it first reflects the conviction running through this entire framework that anticipation, not explanation, is the leading function of intelligence.

Defining the series is partly a matter of saying what it is not. Threat Warning is not concerned with what happened; that backward look belongs to the record kept in memory and to the after-the-fact work of investigation. It is not concerned with what a development means; interpretation is the province of the Analysis series. And it is not concerned with what decision should be made; that judgment belongs to the people the Decision Support series serves. Threat Warning asks a narrower, more forward-leaning set of questions. What should the organization be watching? Which indicators matter? Which threats are emerging? Which conditions are changing? The series exists to keep those questions always live, so the organization is oriented toward the future rather than perpetually explaining the past.

Because its questions are anticipatory, the series is built around indicators and conditions rather than conclusions. An indicator is an observable whose appearance suggests that a threat is developing, and a Threat Warning requirement is, at its core, a standing statement of which indicators the organization considers significant and what pattern among them should prompt a warning. This is the discipline of strategic warning as the intelligence community has long practiced it, a patient watch for the signs that precede an event rather than the report that follows it (Grabo, 2004; Kent, 1966). The deeper mechanics of warning, including the handling of weak signals, the calibration of probability against certainty, and the cognitive traps that warning must avoid, belong to the treatment of anticipation in Part II. What matters here is that Series 1 captures those anticipatory needs as requirements and gives them a place in the architecture.

Placing Threat Warning first also reveals an elegant correspondence between the four series and the four stages of the OLAA loop. Each series serves one stage. Threat Warning supports Anticipate by supplying the forward-looking requirements that enable the organization to see conditions forming before they arrive. Analysis supports Learn, turning observed events into the understanding that accumulates in memory. Decision Support supports Act, connecting intelligence to the decisions through which the organization acts. And

Environment supports Observe, maintaining the wide field of view from which raw signals are first drawn. The Series Method is not a separate apparatus bolted onto OLAA. It is the requirements expression of the same loop, with each series feeding the stage it was built to serve.

Seen through that correspondence, the primacy of Threat Warning is the primacy of Anticipate. OLAA is a loop, and a loop has no true beginning, but the function that gives the model its forward orientation is anticipation, and the series that carries anticipation into requirements is Threat Warning. This is why it is Series 1, and why the wire-fraud case that recurs throughout the paper turns on a Threat Warning requirement firing first. The warning is what sets the rest in motion. Decision Support, Analysis, and Environment all do essential work, but they do it in response to, or in preparation for, the moment that Threat Warning exists to anticipate.

Organizing Requirements by Purpose

It is worth stating plainly what the Series Method is and is not. It is not a prioritization framework. It is an organizational architecture. A prioritization framework answers what to do first. An organizational architecture answers the prior question: how the entire body of requirements is structured so each requirement can be found, understood, and acted upon. The Series Method addresses the second question by making purpose the primary axis of organization.

Within that architecture, four elements each carry a distinct load. The series provide context, situating each requirement within the function it serves. The hierarchy provides decomposition, breaking a broad intelligence need into progressively more specific and answerable components. The numbering structure provides traceability, giving each requirement a unique address that records exactly where it sits within its series and decomposition. And collection provides support, attaching the specific gathering tasks that feed the requirements above them. Priority still exists within this structure, but it operates within each series rather than across all of them, applied to requirements that have already been grouped by what they are for.

Organizing this way provides a stronger foundation than ranking because it preserves information that ranking destroys. A single priority list tells the organization the order of its concerns and nothing else. A purpose-organized architecture tells the organization, for any requirement, what it is for, whom it serves, how it decomposes, where it sits, and which collection supports it. That richer structure makes everything downstream possible, from targeting to retrieval to delivery.

Strategic, Operational, and Tactical Alignment

Within each series, requirements are decomposed into a hierarchy, and the purpose of that hierarchy is alignment rather than prioritization. Organizations do not operate at a single altitude. They form strategic intent, translate it into operational execution, carry it out through tactical action, and support it all with specific collection. A requirement framework that lives at only one of these levels cannot connect intent to action. The Series Method hierarchy exists to maintain that connection, ensuring that a broad strategic concern can be traced without interruption all the way down to the concrete task that serves it.

The decomposition proceeds through four levels. At the top sits a strategic requirement, a broad statement of intelligence intent. Beneath it sits an operational requirement, which translates that intent into the conditions and activities the organization must monitor. Beneath that sits a tactical requirement, specific enough to be matched to a particular signal, source, and affected stakeholder. Supporting the tactical level are collection requirements, the discrete gathering tasks that feed the structure from below. The numbering makes the lineage explicit. A Threat Warning series carries the address 1; its strategic warning requirement becomes 1.1; the operational warning requirement beneath it becomes 1.1.1; the tactical warning requirement becomes 1.1.1.1; and a collection requirement attached to it is written 1.1.1.1:1, the colon marking collection as support to the requirement above rather than another step in the same chain.

Read from the bottom up, that address is a chain of justification: this collection task exists to satisfy this tactical requirement, which serves this operational requirement, which advances this strategic intent. Read from the top down, it is a chain of decomposition that translates a statement of purpose into an executable task. Either way, the hierarchy aligns. It does not rank requirements against one another. It keeps intent and action tied together, so that nothing the organization collects is disconnected from a purpose, and no purpose the organization holds is left without the means to pursue it.

Collection Requirements and Traceability

Collection requirements occupy a deliberately distinct place in this structure. They are not a fifth level of the hierarchy but a supporting layer attached to it, which is why they are numbered with a colon rather than another decimal point. The distinction matters because collection answers a different question from the requirements above it. The strategic, operational, and tactical levels describe what the organization needs to know. Collection describes how a specific piece of that knowledge will be obtained. Conflating the two is precisely the error of the traditional, collection-centric requirement, and keeping them separate helps the Series Method point downstream toward use rather than only upstream toward gathering.

Holding collection as an explicitly traceable support layer yields two benefits. The first is accountability for effort. Every collection task can be traced upward to the requirement it serves and, through that requirement, to the strategic purpose that justifies it, so the organization can see at a glance whether its gathering activity is connected to anything it cares about. Collection that cannot be traced to a requirement is collection without a reason and can be retired. The second benefit is coverage analysis in the other direction. Any requirement can be examined to determine whether sufficient collection is feeding it, exposing the requirements the organization claims to hold but is not actually resourced to satisfy. Traceability, in other words, turns the requirement structure into a map of both what the organization intends and what it is genuinely equipped to deliver.

The Series Method as a Targeting Architecture

The payoff of organizing requirements by purpose becomes clearest when a threat appears. Without this structure, a threat triggers a single reflex. Something alarming is detected, and the organization alerts everyone because it has no structured way to decide who should care. The warning competes for attention with everything else in the inbox and, as Part I described, is largely ignored. The signal arrives, but it has nowhere precise to land.

In an organization that uses the Series Method, the same threat flows through a chain. The signal matches a series, which identifies the function the threat implicates. The series points to a specific requirement, decomposed to a targetable grain. That requirement already names the audience it serves and the decision that audience owns. Delivery then follows that audience and that decision, and no one else. The intelligence reaches the people who should care because the requirement defined who cares before the threat ever arrived. Targeting is not improvised in the moment of alarm. It was structured in advance, and the threat merely activates it.

Understood this way, targeting converges on three elements the paper has assembled in turn. The structured requirements of the Series Method define what intelligence is needed and for whom. The exposure map held in organizational memory, described in Part III, records who and what are exposed to which risks. The stream of live signals arriving through observation, described in Part II, reports what is happening now. Targeting occurs where these three meet. When a live signal matches a requirement, and that requirement maps through the exposure map to specific exposed stakeholders, the organization has precisely identified who needs what at this moment. None of the three is sufficient alone. Requirements without memory cannot say who is exposed; memory without requirements cannot say what a signal demands; and signals without either are merely noise. The convergence is the engine room of Just-in-Time Intelligence, and the Series Method is the component that makes it addressable.

The Series Method and Organizational Memory

The Series Method does more than route incoming signals. It also imposes order on what the organization already knows. Part III argued that the difference between storage and memory is retrieval, and that retrieval depends on knowledge being tagged along the dimensions by which it will later be recalled. The Series Method supplies exactly such a set of dimensions. Every item held in organizational memory, whether an incident record, a lesson learned, an assessment, or a prior decision, can be associated with the series it pertains to, and within that series with the strategic, operational, and tactical requirements it informs.

When memory is organized around requirements in this way, retrieval becomes dramatically easier, because the question a person asks now of need is almost always a question about a requirement. An approver facing a wire decision is, in effect, querying the Decision Support requirement that governs that decision, and everything in memory tagged to that requirement- the prior fraud incidents, the verification procedures, the outcomes of past cases- can be surfaced together. The requirement structure becomes the index to memory. Knowledge is no longer scattered across repositories waiting for someone to know the right search term. It is filed against the very requirements that will call for it.

This is also what links the two halves of the model into a single loop. Knowledge gathered and produced under the four series populates memory, and memory, organized by those same series and their requirements, feeds the targeting the Series Method performs. Requirements shape what the organization remembers, and memory, in turn, makes requirements answerable. The same architecture that tells the organization what to watch for also tells it where to find what it already knows.

The Series Method as the Foundation of Just-in-Time Intelligence

Everything to this point converges on a single shift in how requirements are understood. In traditional practice, a requirement is a question that is posed, answered, and closed. For Just-in-Time Intelligence, a requirement is a persistent, standing condition that is evaluated continuously. A tactical requirement, decomposed to its lowest grain, functions as a trigger. It describes a condition, and when observation reports that the condition has been met, the requirement fires and delivery follows. The requirement is not a question the organization asks once. It is an always-on query against the live environment.

This is what makes Just-in-Time Intelligence pull-based and threat-triggered rather than push-based and calendar-driven. The calendar does not decide when intelligence is delivered. The requirement does, by detecting that its triggering condition has occurred. Because each requirement already carries its purpose, audience, and decision, the act of

firing carries all the information delivery needs. No separate step requires someone to decide who should be told; the requirement settled that when it was written. The organization can therefore act the instant a condition that matters appears, rather than waiting for the next scheduled briefing or the next annual training cycle.

This is why the Series Method is properly understood as the foundation of Just-in-Time Intelligence rather than one technique among many. Organizational memory supplies the knowledge, and the delivery architecture described in Part V supplies the means, but neither can produce targeted intelligence without a structure that, in advance and in detail, defines what each part of the organization needs and when. Strip away the Series Method, and Just-in-Time Intelligence has no addressing system. It collapses back into the broadcast model, holding the right information but unable to say whom it is for. With the Series Method in place, the organization has a standing architecture that converts a signal and a memory into a delivered, targeted warning that is now relevant. That architecture is the spine on which the rest of the model stands.

The Series Method in Action

The wire-fraud case from earlier sections illustrates the method in action once its requirements are made explicit. Two requirements are in play. A Decision Support requirement governs the wire-approval decision, specifying that approvers need current intelligence on active fraud techniques targeting the organization at the point of approval. Decomposed to its lowest level, that requirement specifies the triggering condition (an active fraud campaign in the sector), the affected stakeholders, the current set of approvers drawn from the exposure map, and the content to deliver (the distinguishing indicators and the verification step). Alongside it sits a Threat Warning requirement that continuously monitors for sector-wide wire fraud campaigns.

When the sector ISAC report arrives, the Threat Warning requirement is triggered. Combined with the standing Decision Support requirement, the exposure map identifying current approvers, and the observation that a high-volume payment window is open, the trigger converges on a precise target. The result is the delivery described in Part II: a brief, role-specific warning to exactly those approvers, delivered immediately before approval. Strip away the structured requirements, and the same signals have nowhere to land. The organization holds the ISAC report and monitors quarter-end activity but has no clear statement of who needs what. It predictably falls back on a generic bulletin that competes for attention and loses out.

The same machinery extends naturally to other decisions. A Decision Support requirement attached to executive travel, decomposed to identify the traveling executive and the destinations that elevate risk, will trigger on the simple observation that a covered executive

has a trip booked to a covered location, delivering tailored guidance before departure. In each case, the pattern is identical. A decision is named, a requirement is decomposed to a targetable grain, a condition is monitored, and delivery occurs when memory, requirement, and signal converge. The Series Method makes that pattern repeatable across the enterprise rather than improvised, case by case.

From Targeting to Delivery

With requirements structured by the Series Method, anticipation can finally state, in operational terms, who needs what and when. The organization knows which decisions matter, which conditions trigger which requirements, and which stakeholders each requirement maps to, all from memory. What remains is the system that physically delivers the resulting intelligence into the decision itself, at the place and instant the decision is made, without depending on the recipient to go looking for it. That delivery architecture, the machinery that carries a fired requirement all the way to a changed behavior at the decision point, is the subject of Part V.

Part V: The Architecture of Just-in-Time Intelligence

From Principle to System

The preceding parts established a why, a model, a foundation, and a targeting mechanism. Part I diagnosed the utilization problem. Part II proposed the OLAA loop as a response. Part III explained how to engineer organizational memory. Part IV showed how the Series Method turns requirements into targetable, standing triggers. What remains is the hardest step of all, because a principle that cannot be operated is only a principle. Just-in-Time Intelligence demands an architecture that reliably carries a fired requirement from a detected condition to a changed behavior at the decision point, at scale, without depending on anyone to go looking for what they need.

This part describes the architecture in two registers. It first sets out, in vendor-neutral terms, what any system that delivers Just-in-Time Intelligence must be able to do, because the functional requirements follow directly from the OLAA loop and do not depend on any product. It then grounds those requirements in a concrete implementation: Theia, the platform built specifically to operationalize this model. The purpose of describing a real system is not to advertise it but to demonstrate that the architecture is buildable and that the loop described throughout this paper can be expressed in software and run as a continuous capability rather than admired as a diagram.

Example: Business Email Compromise

A finance employee receives an email requesting an urgent wire transfer. Traditional security awareness training may have occurred months earlier. Under a JITI model, the request itself triggers the response. Intelligence on adversary tactics, prior incidents, and organizational procedures is delivered now the decision is made. The objective is not to improve general awareness but to influence the specific decision that matters.

What a Just-in-Time Architecture Must Do

The functional spine of any Just-in-Time architecture is defined by the loop it serves. Seven capabilities are non-negotiable. The system must capture and decompose intelligence requirements into the structured, standing form defined by the Series Method. It must continuously collect against those requirements, drawing on both external sources and internal telemetry. It must preserve and retrieve organizational memory, including the exposure map linking requirements to people and assets. It must detect, in real time, when a standing requirement has been triggered by an incoming signal. It must target the specific

stakeholders to whom the triggered requirement maps. It must deliver the resulting intelligence to those stakeholders in context, at the decision point. And it must observe the outcome of that delivery and feed the result back into memory, closing the loop so the next cycle is sharper than the last.

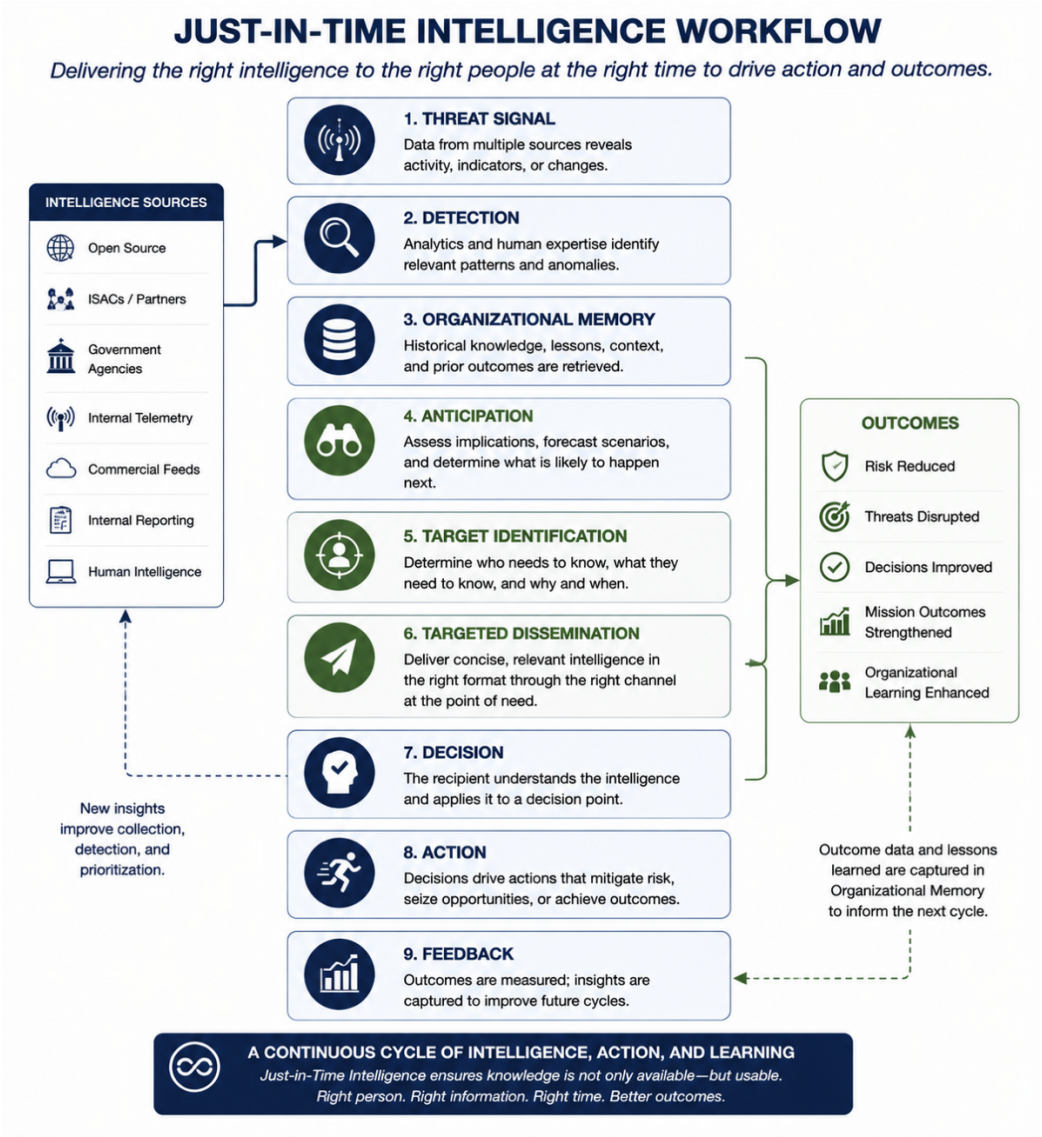


Image 5: JITI Workflow

Each of these capabilities corresponds to a stage of OLAA, and the chain is only as strong as its weakest link. A system that collects brilliantly but cannot target delivers noise. A system that targets precisely but cannot deliver at the decision moment produces reports no one reads in time. A system that delivers well but never observes its own effect cannot improve and cannot distinguish utilization from dissemination. The architecture, in other

words, must implement the entire loop, not a favored portion. Theia is organized around this chain.

Theia: An AI-Assisted Intelligence Management Platform

Theia is positioned as the first AI-assisted intelligence management platform, and the phrase is worth unpacking because each word reflects a deliberate architectural choice. It is an intelligence management platform, not a threat-intelligence feed or a security tool, because its purpose is to run the entire intelligence loop, from requirements to the delivered outcome, rather than add another stream to an already crowded inbox. It is AI-assisted rather than AI-driven, because the model treats artificial intelligence as an augmentation that accelerates collection, analysis, and orchestration while keeping human judgment in the loop where it belongs. Human-machine collaboration is most effective when artificial intelligence augments rather than replaces human judgment (Klein, 1998; Kahneman, 2011). And it is built around managing the loop itself, which is precisely the function that traditional tooling, focused on collection and analysis, has historically left unowned.

In architectural terms, Theia provides a home for each of the seven capabilities above and integrates them into a single, continuous flow. The following sections walk through that flow in order: requirements, collection, memory and analysis, orchestration and delivery, the closed loop of measurement, and finally the deployment model that lets the whole system run close to where decisions are actually made.

Requirements: The IR Wizard and the Series Method

Theia operationalizes the Series Method through a guided intelligence-requirements capability, the IR Wizard, which builds requirements across the four series and decomposes each one using the four-level hierarchy described in Part IV. This is the platform’s addressing system. The IR Wizard breaks down a broad intelligence need, such as protecting the organization from financial fraud, into discrete, answerable, targetable units that the rest of the architecture can act on, complete with the triggering conditions, affected roles, and the content each unit specifies.

Crucially, requirements created this way are not static documents filed and forgotten. They become the system's standing logic, the persistent conditions against which the orchestration layer evaluates the live environment. By making the Series Method native to the platform rather than an external methodology applied by hand, Theia ensures that every downstream action, every collection task, every trigger, and every delivery traces back to an explicit, decomposed requirement. Nothing is delivered that no requirement called for, and no requirement, once defined, sits idle waiting for someone to remember it.

Collection: AI-Assisted and Standards-Based

Meeting those requirements, Theia runs an AI-assisted collection that draws on both external sources and internal telemetry. Here, AI handles the volume and velocity that overwhelm human collectors, triages incoming material, extracts what is relevant to active requirements, and surfaces the signals that matter while suppressing the noise that doesn't. Collection is requirement-driven, keeping it focused: the platform gathers data in service of defined needs rather than indiscriminately accumulating it and leaving analysts to sift through.

Because no intelligence platform operates in isolation, Theia is built to standards. It integrates with MISP and TCS via STIX 2.1, connecting the organization to the broader threat-intelligence sharing ecosystem and enabling structured indicators to flow in and out without bespoke translation. Source and information reliability are assessed using the Admiralty scale, so that what enters memory carries an explicit measure of how far it can be trusted. Standards-based integration is not a checkbox feature. It allows organizational memory to be fed continuously from the outside world in a form anticipation can use.

Memory and Analysis: The Analytics Workbench

Collected and shared material converges in the Analytics Workbench, the environment where analysts work the organization’s memory and turn information into understanding. This is where the operational memory of Part III lives in practice: structured, tagged along the dimensions of role, asset, decision, threat, and sector, and retrievable in response to an emerging situation rather than merely searchable. The Workbench is also where analytic tradecraft is applied, with hypotheses weighed and assessments formed, and where AI augmentation accelerates the mechanical parts of the work so human attention can focus on judgment.

The Analytics Workbench matters to the Just-in-Time architecture because it bridges raw collection and targetable knowledge. A trigger is only as good as the memory behind it, and the Workbench is where that memory is built, enriched, and kept current. Without an environment dedicated to working memory, an organization is left with collection feeding directly into delivery, which is to say signals pushed out before they have been understood. The Workbench is the discipline that keeps Just-in-Time from becoming just-in-time noise.

Orchestration and Delivery: Theia Flow

If requirements are the standing logic and memory is the substrate, Theia Flow is the engine that runs the loop. It is the AI orchestration layer that continuously evaluates active requirements against incoming signals and organizational memory, detects when a

standing requirement has been triggered, resolves the triggered requirement by mapping it through the exposure map to the specific affected stakeholders, and delivers the resulting intelligence to those stakeholders. Theia Flow is where the abstract idea of requirements as standing triggers, introduced in Part IV, becomes executable machinery.

This orchestration is what separates a Just-in-Time architecture from a conventional analytics platform. In a conventional tool, a human notices something, decides who should know, and sends a message. That human bottleneck is exactly where utilization fails, because attention is scarce and the right person is not always watching. Theia Flow removes the need for someone to happen to make the connection at the right moment. When a condition that a requirement cares about appears, the orchestration layer fires, targets, and delivers at machine speed, without waiting for the next briefing cycle. Human judgment governs the requirements and the analysis; orchestration governs the timing and routing, which are precisely the things humans do unreliably under load.

Threat Warning in Action: A Business Email Compromise

It is worth tracing a single threat from first signal to prevented loss, because doing so shows the architecture working as a whole and puts the Threat Warning series, Series 1 of the Series Method, at the leading edge where it belongs. The recurring business email compromise case serves well, now followed end to end through the platform rather than treated in the abstract.

The sequence begins when a threat actor opens a business email compromise campaign against the organization's sector. A standing Series 1 Threat Warning requirement, decomposed to monitor for sector-wide wire-fraud activity, is already watching for exactly this condition. When the matching signal arrives, whether from a sector ISAC report, a partner notification, or internal telemetry, the requirement's triggering condition is met and a threat warning is generated. No analyst had to be watching at that moment. The requirement was.

Theia Flow then resolves the fired requirement. It maps the warning through the exposure map held in organizational memory to the specific people the threat endangers, the finance staff currently authorized to approve and release payments, and it assembles the context those people need: the indicators that distinguish a fraudulent request, the verification step to take, and the precedent from any similar campaign the organization has weathered before. The Just-in-Time alert is delivered to exactly those approvers, in the channel where they work, immediately before the exposure window rather than weeks earlier in a general bulletin. An approver receives the warning, recognizes the pattern it describes, pauses on the suspicious request, verifies through a trusted channel, and declines to release the funds. The fraud is prevented.

Every step in that sequence is owned by a component the paper has already established. Series 1 supplies the standing trigger that turns an incoming signal into a warning. Organizational memory supplies the exposure map that turns the warning into a named audience and the precedent that gives it credibility. Theia Flow supplies the orchestration that fires, targets, and delivers at machine speed, removing the human bottleneck where utilization usually fails. And because the loop observes its own effect, the outcome, an alert that was opened, acted upon, and followed by a prevented loss, flows back into memory and makes the next warning sharper. This is anticipation made operational, with Threat Warning at the front of the loop and Just-in-Time delivery carrying it the last mile to the decision.

The Closed Loop: Measurement and Maturity

Delivery is not the end of the architecture because OLAA is a loop, and the loop only earns its name if it observes its own effect. Theia is designed to capture what happens after delivery and route it back into memory, so the system can distinguish intelligence that changed behavior from intelligence that was merely sent. This feedback is the architectural expression of the distinction at the heart of Part I, the difference between dissemination and utilization, rendered as something the platform can observe rather than assume.

Theia also makes the organization's progress visible through a Maturity Tracking capability that maps platform behaviors to the five levels of the Intelligence Maturity Pyramid. Rather than asserting that an organization has matured, the platform infers maturity from how the organization operates the loop, turning an abstract model into an observed trajectory. This measurement scaffolding anticipates Part VI, which fully addresses how to measure utilization rather than output. The platform's closed-loop and validation capabilities, including its ability to exercise defenses against simulated activity, answer the same underlying question before any real incident: when intelligence is delivered, does behavior change and does risk fall?

Deployment at the Operational Edge

An architecture for Just-in-Time Intelligence must sit where decisions are made, which directly affects deployment. Theia's current design favors a compact, self-contained footprint: a single container with an embedded database and standards-based identity via OIDC, deployable close to the environment it serves rather than only as a distant, multi-tenant service. This choice matters operationally, not merely technically. Intelligence that must traverse a remote service before reaching a decision maker is subject to an additional point of latency and an additional point of failure, in a model where timing is the whole point.

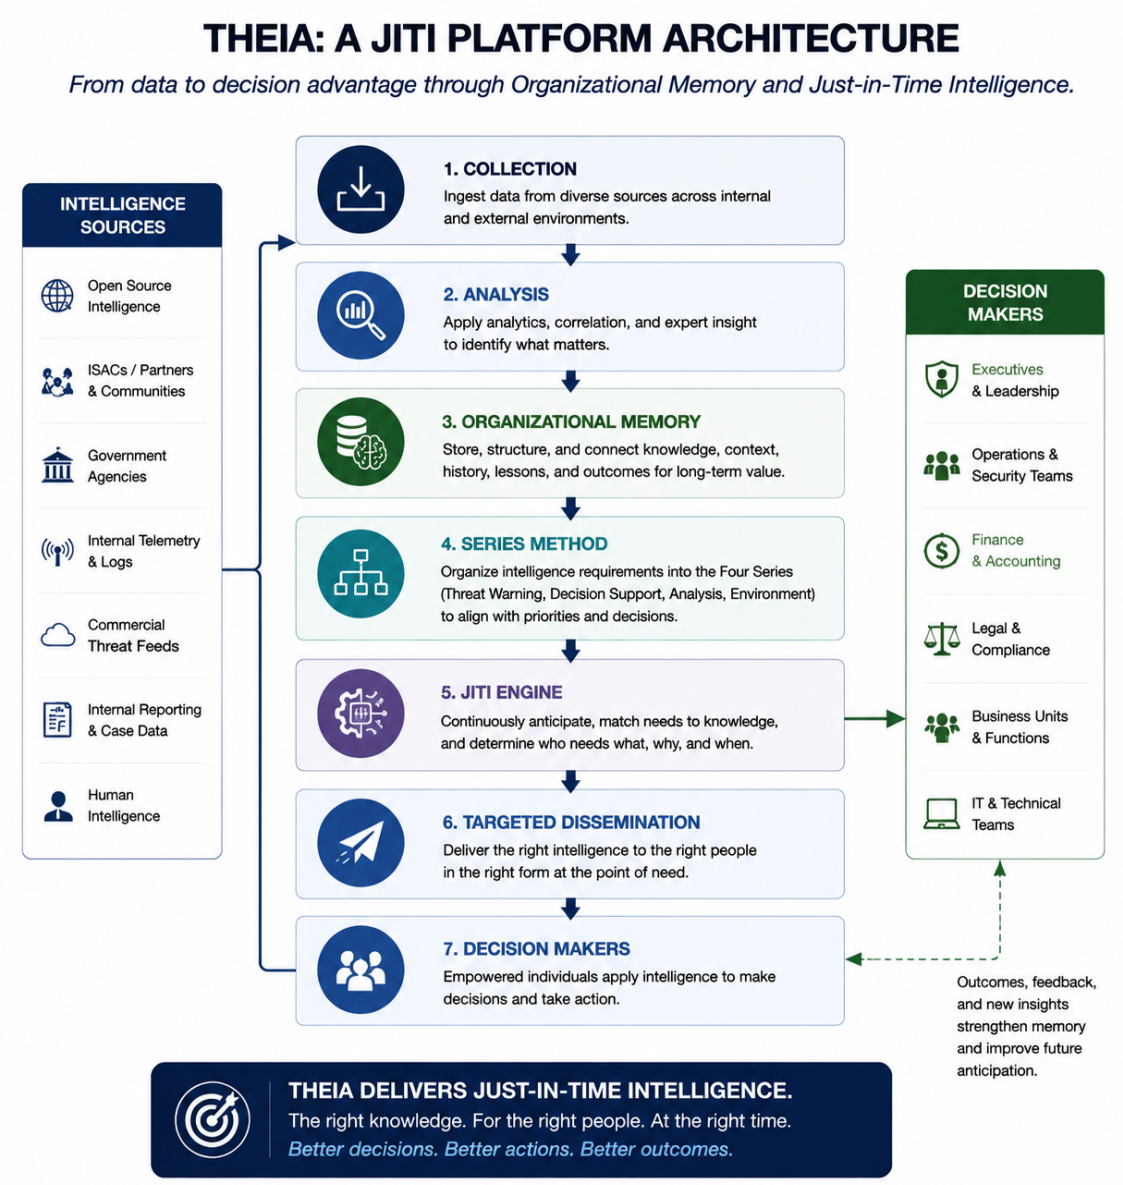


Image 6: Theia - JITI Platform Architecture

A self-contained deployment also suits the realities of the organizations that most need Just-in-Time Intelligence, including those operating in constrained, isolated, or fully remote environments where heavy reliance on cloud is impractical or unacceptable. The same design discipline is reflected in the platform's commercial posture. A thirty-day return-on-investment guarantee is, in effect, the business expression of this paper's thesis: a wager that intelligence that is genuinely used, rather than merely produced, will demonstrate its value quickly and concretely. An organization that measured only outputs could never make such a promise. One built to drive utilization can.

With the architecture in place, the loop can run. Requirements are captured and decomposed; collection feeds memory; orchestration detects triggers and delivers to targeted stakeholders at the decision point; and outcomes flow back into memory to sharpen the next cycle. Theia demonstrates that this is not a thought experiment but a system that can be built and operated. Yet building the machine and proving that it works are different tasks. The architecture produces delivery; it does not, by itself, prove that delivery changed anything. That proof requires a deliberate approach to measurement, one that counts outcomes rather than activity and treats utilization, not output, as the metric that matters. Establishing that approach is the work of Part VI.

Part VI: Measuring Utilization, Not Output

The Measurement Problem Restated

Part I observed that intelligence programs habitually measure outputs rather than outcomes, counting reports produced, indicators collected, and briefings conducted, while leaving the only question that matters, whether any of it changed a decision, unasked. Part V then described an architecture capable of delivering intelligence to the decision point and observing the result. But an architecture that can measure utilization is not the same as an organization that does. Measurement is the discipline that converts the closed loop from a technical capability into an operational truth. In the end, it also makes the loop a loop. Without measuring effect, feedback has nothing to carry, and OLAA quietly collapses back into the production line it was meant to replace. Measuring outcomes rather than outputs reflects principles found in the performance management and training evaluation literature (Kirkpatrick & Kirkpatrick, 2016; Kaplan & Norton, 1996).

Why Output Metrics Persist and Mislead

Output metrics endure for understandable reasons. They are easy to count, they rise reassuringly over time, and they let a program demonstrate visible effort to anyone who asks. The difficulty is that effort and effect are not the same thing. A team can set records for reports produced in the very year it fails to prevent the incidents those reports describe, and the output dashboard will show nothing amiss. Activity and impact have come apart, and the metrics are watching the wrong one.

The deeper problem is that output metrics distort the behavior they are meant to measure. When report count becomes the target, the rational response is to produce more reports rather than better outcomes. Once a measure becomes a target, it tends to stop measuring anything real. The result is a kind of activity theater, in which a program can appear at its healthiest precisely when it is failing at its purpose, because the dashboard rewards motion rather than results. An organization serious about utilization has to resist the gravitational pull of these numbers, not because they are false, but because they are true about the wrong thing.

What Utilization Metrics Actually Measure

Utilization metrics address the questions posed in Part I: Did the intelligence reduce uncertainty for the decision-maker? Did it improve their awareness, enable preventive action, influence behavior, and reduce risk? Within a closed-loop architecture, these questions translate into observable, countable measures. Was the delivered intelligence received and acknowledged at the decision point? Did the recommended action occur? Did

the targeted behavior change? Was the risk event averted? How quickly did action follow delivery? Acknowledgment at the moment of decision, time to action, behavior change, and averted loss are all measurable, and each speaks to effect rather than activity.

These quantities are harder to capture than report counts, and that difficulty is precisely why they were historically avoided. But the architecture of Part V is what makes them capturable in the first place. A system that ends at dissemination cannot observe any of them because it loses sight of the intelligence the moment it is sent. A system that delivers at the decision point and monitors the result can observe all of them. Utilization measurement is therefore not a separate initiative bolted onto the program. It is the payoff of having built a closed loop, the reason the loop was worth closing, and the evidence that the architecture is doing what it was designed to do.

From Metrics to Maturity: Why a Pyramid

Individual utilization metrics tell an organization whether a particular delivery worked. They do not, on their own, describe how capable the organization is at utilization as a whole or reveal where that capability is weakest. For that, a single number is worse than useless because it hides the very variation that matters. What is needed is a model of progression and a way to see capability across its component parts at once.

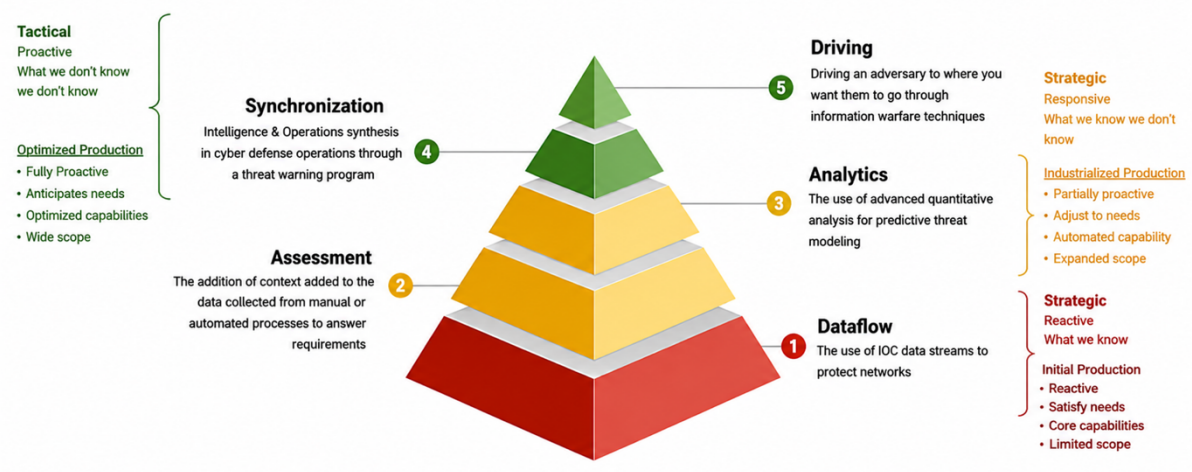


Image 7: The Intelligence Pyramid

The Intelligence Maturity Pyramid serves this purpose. It describes the trajectory of an intelligence operation as it moves from passively handling raw data to actively shaping an adversary, and it breaks that trajectory into five distinct levels of maturity. Although the model is expressed in the language of cyber defense, the progression it captures generalizes to any intelligence enterprise. Its decisive feature, in the context of the architecture described in Part V, is that maturity is inferred from how the organization operates rather

than asserted through self-assessment. An organization does not declare itself mature. Its maturity is read from what it does.

The Five Levels of the Intelligence Maturity Pyramid

Level One, Dataflow. Uses indicator-of-compromise data streams to protect networks. At the base of the pyramid, the organization consumes and reacts to raw indicators. Production is initial, posture is reactive, immediate needs are met with core capabilities within a limited scope, and the knowledge horizon extends no further than the organization's existing knowledge.

Level Two, Assessment. The addition of context to data collected from manual or automated processes to meet requirements. Here, raw data becomes assessed intelligence. Collection is tied to defined requirements, and context is added so that information answers a question rather than merely arriving. This is the level at which the discipline of requirements, the subject of Part IV, first becomes load-bearing.

Level Three, Analytics. The use of advanced quantitative analysis for predictive threat modeling. The organization shifts from describing what has happened to modeling what is likely to happen. Production becomes industrial and predictive, posture becomes partially proactive, capabilities become automated, scope expands, and the knowledge horizon reaches what the organization knows it does not know.

Level Four, Synchronization. The synthesis of intelligence and operations in cyber defense forms a threat warning program. Intelligence and operations stop running in parallel and begin to act as one. The organization anticipates needs and operates with a wide scope. This is the level at which Just-in-Time Intelligence becomes structural, because a threat warning program is, in essence, the institutional commitment to delivering the right warning to the right operator at the right time.

Level Five, Driving. Driving an adversary to where you want them to go through information warfare techniques. At the summit, the organization is no longer merely reacting to or anticipating the adversary but actively shaping the adversary’s behavior. Production is optimized, posture is fully proactive, and the knowledge horizon extends to what the organization does not know it does not know.

The shape of the climb is significant. The lower three levels build a substrate of capability, turning data into context and context into prediction. The upper two levels are inherently acts of utilization performed at operational speed. An organization cannot synchronize intelligence with operations, much less drive an adversary, if its intelligence does not reliably reach and influence decisions in time. The ascent up the pyramid is therefore gated by utilization, and Just-in-Time Intelligence is the mechanism for that ascent. The constructs

developed in this paper, including the OLAA loop, engineered organizational memory, the Series Method, and the delivery architecture, are precisely what carry an organization upward from Dataflow toward Synchronization and Driving.

Reading Maturity Across Dimensions

Maturity is not monolithic, and the pyramid does not represent a single ascending score but a progression visible across several attributes at once. Production maturity advances from initial, through industrial and predictive, to optimized. Posture shifts from reactive, through partially proactive, to fully proactive. The organization's relationship to need moves from satisfying needs, to adjusting to them, to anticipating them. Capabilities progress from core, through automated, to optimized, and scope widens from limited, through expanded, to wide. Across it all runs a knowledge horizon that climbs the familiar ladder from what we know, to what we know we do not know, to what we do not know we do not know, alongside a consumer focus that spans the strategic, operational, and tactical.

Because these attributes advance unevenly, an organization can sit at different levels across dimensions, predictive in its analytics yet reactive in its delivery, and broad in its collection scope yet narrow in its synchronization. Just as Part V argued that the architecture is only as strong as its weakest link, maturity is capped by the lagging dimension. Assessing each dimension on its own exposes the binding constraint rather than averaging them into a single, misleading grade. In Theia, the Maturity Tracking capability formalizes this assessment, scoring the organization across the model's attributes based on observed behavior rather than self-report, so that the pyramid functions as a live diagnostic rather than a static picture.

Maturity as a Diagnostic, not a Scorecard

The pyramid is not a vanity score or a badge to display. It is a diagnostic that identifies the binding constraint and, therefore, the next worthwhile investment. Read correctly, a maturity profile is a map that shows where the loop breaks and what to repair first, which is far more actionable than a grade. The aim is not to reach Level Five in the abstract, but to find, at any given moment, the single dimension whose improvement would do the most to convert produced intelligence into used intelligence.

This is also why inferring maturity from observed behavior matters so much. A self-assessed scorecard invites exactly the gaming that plagues output metrics, with organizations rating themselves generously and mistaking aspiration for capability. Maturity inferred from how the loop is operated grounds the assessment in reality. This is the role of the Maturity Tracking capability described in Part V: by observing how requirements, collection, memory, delivery, and feedback are genuinely exercised, the platform infers the maturity level and

dimension profile from behavior, turning maturity from a claim the organization makes about itself into an observation drawn from what it does.

Closing the Loop on Measurement

Measuring utilization and tracking maturity enable an organization to improve deliberately rather than by accident. With outcomes observed and the binding constraint identified, improvement becomes directed rather than hopeful, and the organization can invest where the return is real. In this sense, measurement is not the end of the work but the steering mechanism for all of it.

Yet measurement and architecture, however sound, do not run themselves. People ultimately sustain a loop: the roles that own requirements, curate memory, and govern delivery, and the cadence and culture that keep the loop turning when attention is scarce and operational pressure is high. Translating the architecture and metrics into a durable daily practice, one that survives staff turnover, competing priorities, and the natural entropy of any standing process, is the operating model's work. That is the subject of Part VII.

Part VII: The Operating Model

From Capability to Practice

The architecture of Part V can deliver intelligence to the decision point, and the measurement of Part VI can prove whether that delivery changed anything. Neither, by itself, sustains the loop. A capability that exists is not the same as a practice that runs, and people and disciplined routines bridge the gap. This part addresses the operating model: the human roles, cadence, governance, and culture that keep the loop turning when attention is scarce, staff changes, and operational pressure is high. It is the least technical part of the paper and, in many organizations, the one that most often determines success or failure.

The operating model matters so much because a loop, left to itself, runs down. The machinery described in earlier parts is necessary, but without ownership and rhythm, it decays. The operating model is what turns a well-designed system into a living capability that improves rather than erodes.

The Risk of Entropy

Every standing process is subject to entropy, and intelligence loops are no exception. Unreviewed requirements drift out of alignment with the decisions they were meant to support. As Part III argued, organizational memory has a half-life. Without curation, the exposure map ages and the threat history goes stale. Feedback that is not deliberately captured quietly stops flowing, and the loop, no longer observing its own effect, collapses back into the production line it replaced. Left unattended, an organization that has reached Synchronization can slide back toward Dataflow without anyone deciding it should.

The operating model counters this entropy. Its purpose is not to add bureaucracy but to assign ownership and establish rhythm, so that the maintenance the loop requires happens rather than being perpetually deferred until the next incident forces it. Entropy is the default. A deliberate operating model is the only thing that keeps it at bay.

Roles That Own the Loop

A sustainable loop requires that each function be owned, even if, in a small organization, one person owns several. The point is accountability, not headcount. Five ownership functions recur. Requirements ownership keeps the Series Method current by decomposing new requirements and retiring obsolete ones, ensuring the addressing system remains accurate. Memory stewardship curates organizational memory by maintaining the exposure map, refreshing threat history, and converting tacit knowledge into explicit form before it walks out the door. Analytic ownership operates the assessment and analysis layer, turning

collected material into understanding. Delivery and orchestration governance owns the triggers and routing logic, tuning what fires and to whom. Feedback ownership tracks utilization, runs the maturity diagnostic, and identifies the binding constraint that should guide the next investment.

The augmentation described in Part V changes how many people these functions require, not whether they require owners. Artificial intelligence can handle much of the mechanical load of collection, triage, orchestration, and routing, which allows a small team to run a loop that once required a large team. But judgment about what the organization needs to know, what belongs in memory, and whether delivery is working remains human work. The operating model names those who hold that judgment, so that no function of the loop is left as everyone's responsibility and therefore no one's.

Cadence: The Rhythm of the Loop

The loop runs at two very different tempos, and conflating them is a common mistake. Some functions are continuous and event-driven, operating at machine speed: observation, the evaluation of standing requirements, and Just-in-Time delivery that triggers as soon as a triggering condition appears. These are not scheduled. They happen whenever the environment warrants, which is the entire point of a pull-based, threat-triggered model.

Other functions are periodic and human, requiring a deliberate cadence. Requirements should be reviewed on a defined schedule to keep pace with evolving decisions and exposures. Curate memory on a schedule to prevent drift. Review utilization and run the maturity diagnostic regularly to identify the binding constraint before it festers. The distinction is important enough to state plainly: in this model, intelligence delivery is event-driven, but the governance of the system that delivers it is calendar-driven. The calendar that generalized awareness programs misapply to delivery, sending training in January for a threat that arrives in September, belongs instead to the maintenance of the machine, not to the dissemination of its output.

Governance and the Discipline of Restraint

A Just-in-Time program has a characteristic failure mode that its operating model must actively manage. A system built to deliver the moment a condition is met can, if left ungoverned, deliver too often. Over-triggering destroys utilization as surely as silence does. Alert fatigue is not a minor irritation; it is the reintroduction of the attention-scarcity problem from Part I through the program's own front door. A recipient buried in warnings stops reading them, and the most important warning is lost in the volume of the merely possible.

Governance is, therefore, in large part, the discipline of restraint. It sets the thresholds that determine which fires to fight, ensures that every delivery traces to a genuine requirement

and includes a genuinely actionable recommendation, and suppresses the noise that would otherwise erode the channel's credibility. A Just-in-Time program must protect the scarcity of attention it depends on, treating each delivery as a withdrawal from a limited account of recipient trust. Governance also extends to what enters memory, applying source and reliability discipline so that the knowledge feeding anticipation is sound, and to the accuracy of targeting, so that intelligence reaches those who are genuinely exposed, not those who merely might be.

Culture and Adoption

The most carefully engineered loop fails if the people on the receiving end do not act on what it delivers. Ultimately, it is a matter of human behavior, and it depends on a culture in which delivered intelligence is trusted and acted upon rather than merely noted and ignored. That trust is earned, not assumed. A program earns it by being consistently right and relevant, by delivering intelligence that recipients come to recognize as worth their attention because past deliveries have proven worth acting on. Each well-targeted, genuinely useful delivery deposits trust; each false alarm or irrelevant broadcast withdraws it.

Adoption also depends on the feedback the loop requires. Recipients must engage sufficiently for the system to observe whether intelligence changed their behavior, which means acting on intelligence and acknowledging it must become a normal part of how decisions are made. This is as much a leadership matter as a technical one. When leadership visibly values utilization over production, asks whether intelligence changed decisions rather than how many reports were filed, and acts on what it receives, the culture follows. When leadership rewards the appearance of activity, the organization optimizes for that appearance, and the loop starves for the engagement it needs.

Sustaining the Loop Through Change

An operating model is finally tested by its resilience to change. People leave, organizations reorganize, and priorities compete. A loop that depends on the heroics of a few key individuals is fragile, which is why the memory stewardship described above matters beyond its day-to-day function: by converting tacit knowledge into explicit, retrievable form, it makes the loop survivable when those individuals depart. A practice that lives only in people’s heads leaves with them. A practice encoded in requirements, memory, and governed routines endures.

There is also a subtler threat to sustainment, namely success itself. A well-run anticipatory program prevents incidents, and prevented incidents are invisible. The better the loop performs, the less necessary it can appear, until the program that quietly averted a dozen crises is questioned because no crisis is visible. This prevention paradox is why the

measurement discipline of Part VI is not merely an internal tuning tool but a survival mechanism. An organization that measures averted risk, time to action, and behavior changed can demonstrate the value of prevention even when prevention's product is the absence of disaster. An organization that measures only incidents responded to will, perversely, find its most effective intelligence program hardest to justify.

Bringing It Together

The operating model turns the architecture of Part V and the measurement of Part VI into a living practice. With each function owned, the right rhythms established, governance protecting the scarce attention the model depends on, and a culture that trusts and acts on delivered intelligence, an organization can operate the OLAA loop continuously and climb the Intelligence Maturity Pyramid deliberately rather than drifting up and down it by accident. The pieces are now assembled: a diagnosis, a model, a memory, a targeting method, an architecture, a means of measurement, and an operating discipline to sustain it all. What remains is to step back from the mechanics and consider what this approach means for the practice of intelligence itself, which is the subject of the conclusion.

Part VIII: The Future of Intelligence

The Argument in Brief

This paper began with a paradox. Organizations have never held more information, yet they continue to suffer failures that the information they already possess should have prevented. The cause, it argued, is not a shortage of collection, analysis, or warning, but a shortage of utilization. Intelligence fails not because it is wrong, but because it does not reach the right person in the right form when a decision is made. Production has been mistaken for effectiveness, and dissemination for use.

Everything that followed was an attempt to close that gap. The OLAA model reframed intelligence as a continuous loop of observation, learning, anticipation, and action, anchored in organizational memory rather than in a production line that ends at dissemination. Memory requires deliberate engineering, not passive accumulation, because anticipation and action can only be as good as the memory they draw on. The Series Method gave requirements a structure precise enough to target an individual at a specific moment. An architecture, expressed concretely in Theia, demonstrated that the loop can be built and run, rather than merely diagrammed. The Intelligence Maturity Pyramid provided a way to measure progress by behavior rather than by assertion, and the operating model supplied the roles, cadence, governance, and culture that keep the loop alive against the constant pull of entropy. Just-in-Time Intelligence is the name for what all of this produces: the right intelligence, delivered to the right decision maker at the precise moment it can change what they do.

A Change in What Counts as Good Intelligence

Beneath the frameworks and architecture, the paper advances a single, simple argument. It asks the intelligence profession to shift its measure of success from production to utilization. The metrics that have long defined a healthy program, such as reports produced, indicators collected, and briefings delivered, describe activity rather than impact. An organization can excel at all of them while failing at the only thing that matters. The principle that should replace them was stated at the outset and has run through every part since: the true measure of intelligence is not whether it was produced, but whether it changed something.

Taking that principle seriously has consequences that reach into every corner of an intelligence operation, from how requirements are written to how memory is preserved, from how delivery is engineered to how success is measured and how leadership evaluates the program. The whole of this paper is, in a sense, an extended working-out of what it means

to take that one sentence seriously. The frameworks are instruments. The measure is the point.

Why This Matters Now

The utilization gap is not new, but the conditions that make it urgent are intensifying. Information abundance is accelerating, attention remains as scarce as ever, and artificial intelligence now generates analysis, summaries, and alerts faster than any human audience can absorb. As production accelerates, the gap between what is produced and what is used does not narrow on its own. It widens because more output competes for the same fixed supply of human attention, and the likelihood that any single product is noticed and acted upon falls.

Artificial intelligence cuts both ways in this story, which is why its application matters so much. Used to generate ever more output, it deepens the flood and worsens the very problem it appears to solve. Used as augmentation within a disciplined loop, to accelerate collection, orchestrate triggers, and route the right knowledge to the right person at machine speed, it becomes the enabling technology of Just-in-Time Intelligence. The same capability that can drown an organization in signal can, directed differently, deliver exactly the signal that matters at exactly the right moment. The choice between those outcomes is organizational, not technical, and it is being made, explicitly or by default, right now.

The Climb Ahead

For most of its history, the intelligence profession has judged itself by its output, and that habit has shaped how programs are funded, staffed, evaluated, and defended. Collection programs were measured by the information they gathered, the sources they accessed, and the coverage they achieved. Analytical organizations were measured by the quality, volume, and sophistication of their assessments. Warning organizations were measured by their ability to identify threats before those threats materialized. Each of these measures reflected a real advance in the profession, and none should be dismissed. Collection, analysis, and warning remain essential pillars of intelligence work. Yet each also carries the same inherited assumption: that intelligence has value once produced.

Modern organizations suggest this assumption is no longer sufficient. Intelligence does not create advantage simply because it exists, and it does not reduce risk simply because it has been written, briefed, or stored. Its value depends on whether it reaches the point of decision in time to shape action. A report that arrives after a decision has been made, a warning that reaches people who cannot act on it, or a lesson that remains buried in an archive may be accurate and professionally produced, but it will not change the outcome. In practical terms, unused intelligence is indistinguishable from absent intelligence.

This is why organizations keep failing even when intelligence is available. They fail because intelligence is unavailable when it is needed. They fail because they forget lessons before similar conditions return. They fail because warnings are disconnected from the decisions they are meant to inform. They fail because knowledge remains trapped in reports, briefings, databases, ticketing systems, and individuals' minds rather than becoming part of the organization's behavior. The decisive issue is not whether the organization knows something somewhere; it is whether that knowledge becomes visible, credible, and actionable when a person must choose what to do.

The next evolution of intelligence will therefore not be defined by better collection, more sophisticated analysis, or faster warning alone. Those capabilities will continue to matter, but they will be judged by a stricter standard: whether they are converted into use. The organizations that gain advantage in the coming decade will be those that can systematically convert knowledge into action, memory into anticipation, and intelligence into decisions. They will treat intelligence not as a reporting function that ends when a product is published, but as a continuous capability for organizational learning, adaptation, and timely intervention.

Just-in-Time Intelligence is one approach to achieving that future. The OLAA Model, Organizational Memory, the Series Method, and the supporting architecture described in this paper are not ends in themselves. They are mechanisms designed to answer a single question: did the intelligence lead to any meaningful change? That question ultimately measures every intelligence program. It asks whether uncertainty was reduced, whether a decision improved, whether behavior shifted, whether risk declined, and whether the organization acted sooner or better than it otherwise would have. The answer isn't in the number of reports produced or the number of briefings delivered. It must be found in the observable effect intelligence has on decisions and outcomes.

That standard changes what the profession must value. Not what was collected, unless collection fed memory and enabled action. Not what was written, unless writing clarified a decision and reached the person responsible for making it. Not what was briefed, unless the briefing altered behavior or prepared the organization to act. The central question is what changed. What decision was improved? What loss was prevented? What risk was reduced? What lesson was retained and reused? What future action became possible because intelligence arrived in time?

The future of intelligence belongs to the organizations that can answer those questions with evidence. They will be able to show not merely that intelligence was produced, but that it was used; not merely that warnings were issued, but that people acted on them; not merely that lessons were captured, but that those lessons shaped later decisions. In an

environment defined by information abundance, attention scarcity, and accelerating artificial intelligence, that ability to demonstrate utilization will become a defining source of credibility and advantage. Intelligence will reach its full value not when it is finished, but when it changes what the organization does.

Conclusion

For generations, intelligence professionals have sought better ways to understand uncertainty and reduce the risk it poses to decision makers. They built collection systems to gather information adversaries, competitors, or complex environments would otherwise conceal. They developed analytic methods to transform scattered information into coherent understanding. They created warning methodologies to identify threats before they materialized and to give leaders time to act before events overtook them. Each advance made intelligence more capable, more disciplined, and more valuable. Yet throughout this evolution, one assumption remained largely unchallenged: that intelligence creates value only after it has been produced.

Modern organizations suggest otherwise. The defining challenge of the information age is no longer access to information, because most organizations already have more information than their people can absorb. The defining challenge is applying information when decisions are made. Organizations routinely hold the knowledge needed to avoid failure, reduce risk, and seize opportunity, but that knowledge often remains disconnected from the moment of action. It may be stored in reports, preserved in briefings, buried in databases, or remembered by individuals, yet it still fails to influence the person who must decide what to do. Organizations lack not knowledge itself, but a reliable mechanism for converting knowledge into action.

This paper has argued that the future of intelligence lies in closing that gap. Just-in-Time Intelligence is one approach to achieving that goal. Organizational Memory, the OLAA Model, the Series Method, and the supporting architecture described throughout this paper share a common purpose: ensuring that intelligence reaches the right person, in the right context, at the moment it can influence behavior. Together, they shift intelligence from a product-centered activity to a decision-centered capability. They make memory operational, anticipation actionable, and delivery measurable.

Ultimately, however, the specific frameworks matter less than the principle behind them. Intelligence exists to improve decisions. Everything else is a means to that end. Collection matters because it gives decision makers access to knowledge they would not otherwise have. Analysis matters because it clarifies meaning and reduces uncertainty. Warning matters because it creates time to act before harm occurs. Dissemination matters because intelligence must reach those who can use it. But none of these functions is complete unless intelligence changes what someone understands, chooses, or does.

When intelligence stops influencing decisions, it stops fulfilling its purpose. A report that is never read, a warning that is ignored, a lesson that cannot be retrieved, or a briefing that produces no change in behavior may still reflect professional effort, but it does not confer an organizational advantage. The standard must therefore shift from whether intelligence was produced to whether it was used.

Did it reduce uncertainty?

Did it change behavior?

Did it prevent loss?

Did it improve the timing, quality, or confidence of a decision?

These questions define intelligence utilization.

The organizations that will thrive in the coming decade will not necessarily be those that collect the most information, employ the largest analytic teams, or produce the most reports. They will be the organizations that learn faster, adapt sooner, and act more effectively than their competitors. They will preserve what they have learned, recognize when that knowledge becomes relevant again, and deliver it to those who can use it before the decisive moment passes. In doing so, they will transform intelligence from a product into a capability and from a record of what the organization knows into a mechanism for shaping what the organization does.

Just-in-Time Intelligence offers one path toward that future. Its promise is not simply better information but better decisions, made sooner, with greater confidence, and with observable impact.

The future of intelligence will not be determined by what organizations know.

It will be determined by what they do with what they know.

References

Intelligence Theory and Warning Intelligence

Boyd, J. (1987). *A discourse on winning and losing*. Unpublished briefing.

Grabo, C. M. (2004). *Anticipating surprise: Analysis for strategic warning*. University Press of America.

Heuer, R. J. (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence.

Kent, S. (1966). *Strategic intelligence for American world policy*. Princeton University Press.

Lowenthal, M. M. (2022). *Intelligence: From secrets to policy* (9th ed.). CQ Press.

Warner, M. (2002). Wanted: A definition of intelligence. *Studies in Intelligence*, 46(3), 15–22.

United States Department of Defense. (2020). *Joint Publication 2-0: Joint Intelligence*. Department of Defense.

United States Department of Defense. (2013). *Joint Publication 2-01.2: Counterintelligence and Human Intelligence Support to Joint Operations*. Department of Defense.

Organizational Learning

Argyris, C. (1977). Double loop learning in organizations. *Harvard Business Review*, 55(5), 115–125.

Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.

Senge, P. M. (2006). *The fifth discipline: The art and practice of the learning organization* (Rev. ed.). Doubleday.

Schön, D. A. (1983). *The reflective practitioner*. Basic Books.

Garvin, D. A. (1993). Building a learning organization. *Harvard Business Review*, 71(4), 78–91.

Knowledge Management and Organizational Memory

Davenport, T. H., & Prusak, L. (1998). *Working knowledge: How organizations manage what they know*. Harvard Business School Press.

Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company*. Oxford University Press.

Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91.

Alavi, M., & Leidner, D. E. (2001). Review: Knowledge management and knowledge management systems. *MIS Quarterly*, 25(1), 107–136.

Jennex, M. E. (Ed.). (2016). *Knowledge management, organizational memory and transfer behavior*. IGI Global.

Decision Science and Human Judgment

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Klein, G. (1998). *Sources of power: How people make decisions*. MIT Press.

Simon, H. A. (1997). *Administrative behavior* (4th ed.). Free Press.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.

March, J. G. (1994). *A primer on decision making*. Free Press.

Learning Science and Human Performance

Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology*. Teachers College, Columbia University.

Knowles, M. S., Holton, E. F., & Swanson, R. A. (2015). *The adult learner* (8th ed.). Routledge.

Sweller, J. (1988). Cognitive load during problem solving. *Cognitive Science*, 12(2), 257–285.

Clark, R. C., Nguyen, F., & Sweller, J. (2011). *Efficiency in learning*. Pfeiffer.

Mayer, R. E. (2020). *Multimedia learning* (3rd ed.). Cambridge University Press.

Intelligence Utilization and Performance Measurement

Kirkpatrick, D. L., & Kirkpatrick, J. D. (2016). *Kirkpatrick's four levels of training evaluation*. ATD Press.

Moore, M. H. (1995). *Creating public value*. Harvard University Press.

Rummler, G. A., & Brache, A. P. (2012). *Improving performance*. Jossey-Bass.

Kaplan, R. S., & Norton, D. P. (1996). *The balanced scorecard*. Harvard Business School Press.

Deming, W. E. (2000). *Out of the crisis*. MIT Press.