CELERUM / RESEARCH

From Capability to Operationalisation

Understanding the Missing Layer in Modern Logistics Planning

Observations from a Year Investigating Planning Intelligence

November 2026

Executive Summary

Modern logistics organisations have more operational data and computational capability than ever before. Yet the decisions that determine service, cost and resilience still depend heavily on experienced planners bringing together information, constraints and judgement that their systems do not fully capture.

Why does this gap persist? Our year-long investigation into planning intelligence began with the expectation that better optimisation and artificial intelligence would be the principal answer. Instead, it revealed a more fundamental challenge: organisations must first understand and represent how planning decisions are made, which constraints matter and how competing priorities should be balanced.

We describe the missing connection as the Planning Intelligence Layer. It is not a replacement for human expertise or mathematical optimisation. It is the means by which planners' knowledge, operational realities and organisational priorities can inform computational analysis—and by which recommendations can be understood, challenged and trusted.

This matters commercially. Without a dependable representation of the planning problem, organisations risk concentrating critical knowledge in individuals, relying on manual workarounds and struggling to evaluate alternatives when conditions change. With that foundation in place, they can progressively improve consistency and control, introduce useful decision support and apply more advanced optimisation with greater confidence.

The paper sets out the observations behind this conclusion and a progressive route from existing capability to operationalisation. Its argument is deliberately practical: lasting improvement does not begin by removing planners from decisions. It begins by making their expertise part of a more explicit, explainable and continuously improving planning process.

PART ONE – THE INVESTIGATION

1. Why We Started

Like many organisations exploring optimisation and artificial intelligence, we began with what appeared to be a relatively straightforward question. Advances in cloud computing, telematics, operations research and machine learning had created unprecedented computational capability. Sophisticated optimisation techniques had become increasingly accessible and academic research had demonstrated their ability to solve highly complex combinatorial problems. It seemed reasonable to assume that the next stage in logistics planning would involve applying these capabilities more directly to operational decision-making.

The logic appeared compelling. Transport Management Systems already captured jobs and resources. ERP platforms provided commercial visibility. Telematics systems delivered real-time operational information. Mobile technologies ensured continuous communication between field operations and control centres. Increasingly powerful optimisation techniques promised the ability to evaluate millions of alternatives beyond the limits of human cognition.

We initially expected optimisation and AI to be the main route to supporting planners. We explored recommendation engines and optimisation techniques, and examined how planners combine operational data with their knowledge of vehicle and trailer suitability, customer requirements and operating constraints. Those observations informed a framework that makes planning rules and priorities explicit so alternatives can be evaluated and explained.

Over time, however, these assumptions became increasingly difficult to support.

Repeated investigations continually led us away from questions concerning computational capability and towards something more fundamental. Questions surrounding information, constraints, decision-making and trust emerged with increasing frequency. It gradually became apparent that before optimisation could become meaningful, there first had to be an understanding of how planning itself actually worked.

What began as an investigation into optimisation was gradually becoming an investigation into planning intelligence.

2. What Surprised Us

Early investigations focused primarily upon understanding how recommendation engines, optimisation techniques and artificial intelligence could improve planning quality. However, it quickly became apparent that optimisation assumes that the planning problem itself has already been understood.

Objectives need to be defined. Constraints need to be managed. Dependencies need to be recognised. Success criteria need to be understood. Recommendations need to be explainable and trusted. Without these foundations, even the most sophisticated algorithms struggle to deliver meaningful outcomes.

Repeatedly, we found ourselves asking questions that initially appeared secondary. What constitutes a better outcome? Why are planners making particular decisions? Which constraints genuinely matter? How should competing priorities be balanced? How can recommendations be trusted?

As these questions accumulated, it became increasingly apparent that existing systems already performed extremely well at recording activities and providing visibility into operations. What they provided far less support for was understanding how decisions were being made and why certain outcomes were considered preferable to others.

Perhaps most surprisingly, much of what made experienced planners successful existed outside the systems themselves. Knowledge, judgement and practical experience remained concentrated within individuals. Human planners continuously balanced competing objectives, managed uncertainty and adapted to changing conditions in ways that proved remarkably difficult to formalise.

What initially appeared to be an optimisation problem gradually revealed itself to be something very different.

The challenge was not replacing human expertise with algorithms. Nor was it choosing between human judgement and mathematical optimisation.

The challenge was understanding how these two capabilities could be integrated in a manner that was explainable, trusted and operationally meaningful.

3. The Information Paradox

One of the earliest assumptions underpinning our work was that planning inefficiencies would be driven primarily by missing information. This appeared entirely reasonable. Modern logistics operations are increasingly complex and customer expectations continue to increase. It seemed intuitive that richer information and greater visibility would naturally lead to better decisions.

However, as the investigation progressed, this assumption became increasingly difficult to support.

Most organisations already possessed vast quantities of information. Transport Management Systems, ERP platforms, telematics solutions, warehouse systems, maintenance applications and mobile technologies had created unprecedented levels of operational visibility. Information was rarely absent. In many cases, there was considerably more information available than planners could realistically assimilate.

Yet despite this abundance, planning remained heavily dependent upon individuals.

Repeatedly, we observed planners acting as the integration layer between otherwise disconnected systems. Information originating from multiple sources was combined with experience, local knowledge and practical judgement before being translated into operational decisions. Rather than systems being designed around planning processes, planners had often absorbed new systems into their existing processes.

Perhaps the clearest symptom of this phenomenon was the prevalence of spreadsheets. Across organisations of all sizes, home-grown spreadsheets frequently emerged as indispensable planning tools. Far from representing poor practice, these spreadsheets reflected something much deeper. They represented attempts by experienced planners to bridge gaps between operational systems and the realities of day-to-day decision-making.

This observation highlighted an important distinction. Most systems had been designed primarily to record activities, manage transactions and provide visibility into historical events. Planners, however, were concerned with something entirely different. Their role was not simply to understand what had happened, but to determine what was possible.

Consequently, increasing quantities of information often increased complexity rather than reducing it. This observation aligned closely with Herbert Simon's concept of bounded rationality, whereby decision-makers seek satisfactory solutions rather than mathematically optimal ones because complexity exceeds the limits of human cognition.

Over time, we became increasingly convinced that the challenge facing organisations was not information scarcity. Nor was it information overload.

The challenge was the absence of a planning intelligence layer capable of transforming information into actionable planning capability.

4. The Constraint Management Problem

If the Information Paradox represented one of the earliest turning points in the investigation, understanding constraints ultimately became one of its defining themes.

At the outset, we assumed that the majority of effort would be spent evaluating optimisation techniques and recommendation engines. Instead, we found ourselves devoting increasing amounts of time to understanding something considerably more fundamental; namely, what experienced planners were actually considering when they made decisions and why those considerations proved so difficult to manage within traditional systems.

Initially, this appeared to be a relatively straightforward question. However, as our understanding deepened, it became increasingly apparent that planning decisions were influenced by a far richer and more interconnected set of factors than we had originally anticipated. Existing systems had been designed primarily to support operational execution. They excelled at recording activities, identifying resources and providing visibility into historical performance. Planning itself, however, required a fundamentally different perspective. Whilst operational systems focused on describing what had happened, planners were continually trying to determine what was possible.

Repeated observations demonstrated that experienced planners rarely viewed resources as simple interchangeable assets. Vehicle capabilities, trailer compatibility, customer preferences, loading restrictions, geographical positioning, future opportunities and countless practical considerations all influenced the feasibility and quality of a plan. Many of these factors were neither formally represented nor consistently visible within the underlying systems. Instead, they existed largely within the accumulated experience of individuals.

As the investigation progressed, we became increasingly aware that experienced planners were continuously managing hundreds of constraints simultaneously. Some were physical, others commercial, operational or regulatory in nature. Many interacted in subtle ways and evolved over time as circumstances changed. What had initially appeared to be a routing problem gradually revealed itself to be a problem of constraint management.

This distinction proved particularly important. Optimisation techniques assume that constraints have already been identified and understood. In reality, much of our effort was spent discovering, validating and understanding constraints rather than applying algorithms. Repeatedly, we found ourselves asking not how to optimise, but what should actually be optimised and which constraints genuinely influenced outcomes.

Over time, we became increasingly convinced that poorly understood constraints do not simply reduce the effectiveness of optimisation. They become exaggerated through optimisation, exposing assumptions and inconsistencies that may previously have remained hidden.

Perhaps the most important conclusion from this phase of the investigation was that the challenge facing organisations was not a lack of mathematical capability. Rather, it was a failure to adequately capture the complexity of the decisions being made and the constraints being managed.

5. Why Human Planners Still Matter

One of the more interesting observations to emerge from the investigation concerned the continuing importance of human planners. At the outset, it was tempting to assume that advances in computational power, optimisation techniques and artificial intelligence would gradually reduce dependence upon individual expertise. In many respects, this appeared to represent the natural progression.

Repeated observations suggested something rather different.

Experienced planners were not simply allocating jobs to resources. They were continuously balancing uncertainty. Service commitments, customer expectations, operational risk, future opportunities and countless practical considerations had to be considered simultaneously, often under significant time pressure. Planning represented a continuous process of compromise between competing objectives, many of which were difficult to define precisely and continually evolved as circumstances changed.

This observation aligned closely with Herbert Simon's concept of bounded rationality and Daniel Kahneman's work on heuristics and decision-making. Experience and intuition emerged not as poor substitutes for optimisation, but as entirely rational responses to operating within complex adaptive systems.

However, these same characteristics created fragility. Operational knowledge became concentrated within individuals, making decision-making difficult to explain and even more difficult to scale. Furthermore, adapting to changing business drivers proved challenging because many constraints and assumptions remained implicit rather than explicit. Experienced planners often continued to operate according to rules that had evolved over time, some of which no longer reflected current priorities.

Perhaps the most important conclusion to emerge from this phase of the investigation was that increasing complexity had not made planners obsolete. Quite the opposite. Complexity had made human expertise even more valuable.

At the same time, however, it became increasingly apparent that neither humans nor optimisation algorithms alone represented the answer. Human judgement provides context, experience and adaptability. Mathematical optimisation provides the ability to evaluate vast numbers of possibilities beyond the limits of human cognition.

The challenge, therefore, is not choosing between human expertise and mathematical optimisation.

The challenge is integrating them.

Over time, we became increasingly convinced that a missing planning intelligence layer was required; one capable of capturing those aspects of decision-making that represent human judgement and allowing them to guide computational optimisation in a focused and controlled manner.

6. What We Concluded

By this stage of the investigation, many of the assumptions that had existed at the outset had been challenged. What began as an exploration into optimisation gradually evolved into something considerably broader.

Repeatedly, we found ourselves returning to the same underlying questions. Did sufficient information exist to support planning? Which constraints mattered? Why were planners making particular decisions? How could expertise itself be represented and trusted?

Increasingly, it became apparent that these questions formed a hierarchy.

Information alone proved insufficient. Constraints needed to be managed. Decisions needed to be understood. Human expertise needed to be represented. Recommendations needed to be trusted. Only then could optimisation and artificial intelligence be operationalised effectively.

Over time, these observations led us to an unexpected conclusion.

Organisations do not lack data. They do not lack systems. Nor do they lack optimisation capability.

What they lack is a planning intelligence layer capable of integrating human judgement with computational optimisation.

This layer sits between human planners and, at the other end, data and optimisation algorithms. It captures those aspects of decision-making that represent human judgement and provides the mechanism through which experience, constraints and practical realities can direct the planning process. In doing so, it allows the computational power of optimisation and artificial intelligence to be deployed in a focused and controlled manner.

Perhaps the most surprising aspect of this conclusion was its simplicity. Looking back, many of these observations appear obvious. Yet they were far from obvious at the outset.

Like many others, we had initially assumed that algorithms represented the primary challenge.

It was only through repeated investigation that we came to understand that the real challenge lay in creating the missing planning intelligence layer.

This conclusion would fundamentally alter the direction of our work and ultimately lead to the development of a very different approach.

PART TWO – THE FRAMEWORK

7. Why Existing Systems Are Not Enough

One of the more interesting observations to emerge from the investigation was that organisations rarely suffered from a lack of technology. In fact, most organisations possessed considerable capability. Transport Management Systems, ERP platforms, telematics solutions, warehouse systems and reporting tools all performed valuable functions and represented years of investment and operational development.

However, repeated investigations demonstrated that these systems had largely been designed to support execution rather than planning. Their strength lay in recording activities, managing transactions and providing visibility into historical events. Whilst these capabilities were essential, they offered limited support in understanding how decisions were being made or why experienced planners consistently produced outcomes that were difficult to replicate.

Increasingly, we observed that planners themselves had become the mechanism through which information from multiple systems was integrated and transformed into operational decisions. Human judgement compensated for gaps in representation, managed competing objectives and adapted to circumstances that existing systems struggled to capture.

Perhaps the most revealing aspect of this phenomenon was that organisations frequently viewed these limitations as normal. Over time, workarounds, spreadsheets and tacit knowledge had become accepted components of the planning process. However, these same characteristics also introduced risk. Expertise became concentrated within individuals, improvement opportunities remained hidden and adaptability became increasingly difficult.

Repeatedly, we found that the issue was not that existing systems were inadequate. Rather, they had never been designed to provide planning intelligence.

Organisations did not necessarily need more systems.

They needed a missing layer capable of connecting existing capability with planning intelligence.

8. Where We Went Wrong

Looking back, many of the conclusions that emerged now appear obvious. However, they were far from obvious at the outset.

Like many others, we initially assumed that optimisation represented the primary challenge. Advances in computational capability and the increasing accessibility of artificial intelligence created the impression that better algorithms would naturally result in better decisions. Early investigations therefore focused heavily upon recommendation engines, optimisation techniques and the mechanics of applying increasingly sophisticated mathematics to operational problems.

In many respects, this seemed entirely logical.

However, repeated investigations gradually led us back to more fundamental questions. We found ourselves spending considerably more time understanding constraints than evaluating algorithms. Questions concerning information and decision-making proved more important than expected. Understanding why planners made particular decisions became more valuable than understanding how those decisions might be automated.

Several areas that initially appeared promising ultimately proved less significant than originally anticipated, whilst seemingly secondary questions concerning trust, explainability and representation became increasingly central to our thinking.

Perhaps the most important lesson to emerge from this phase of the investigation was that optimisation itself was never the problem. Indeed, mathematical optimisation represented an

essential component of the solution. The challenge lay in attempting to apply computational capability before understanding the planning intelligence required to direct it.

Over time, we came to understand that optimisation does not create poor assumptions. It exaggerates them.

This shift in perspective proved to be one of the defining moments of the entire exercise and fundamentally altered both our thinking and the direction of our work.

9. The Missing Planning Intelligence Layer

Over the course of the investigation, one observation emerged repeatedly.

Between human planners and, at the other end, data and optimisation algorithms, there exists a missing layer of planning intelligence.

This layer captures those aspects of decision-making that represent human judgement and provides the mechanism through which experience, constraints and practical realities can direct the planning process. In doing so, it allows the computational power of optimisation and artificial intelligence to be deployed in a focused and controlled manner.

Perhaps the most important characteristic of this layer is that it is neither purely human nor purely mathematical. Rather, it provides the means through which human expertise and computational capability can complement one another.

Human planners provide experience, adaptability and contextual understanding. Mathematical optimisation provides the ability to evaluate possibilities beyond the limits of human cognition. Neither capability alone is sufficient. The value emerges through integration.

Repeatedly, we found that planning intelligence consisted not merely of information, but of relationships, constraints, trade-offs and priorities. These characteristics evolve over time and frequently differ between organisations. Consequently, planning intelligence cannot simply be purchased or imposed. It must be progressively developed, refined and trusted.

This observation ultimately became the foundation upon which the framework itself was constructed.

10. From Capability to Operationalisation

As the investigation progressed, it became increasingly apparent that capability alone does not guarantee better outcomes.

Most organisations already possess considerable capability. Systems exist. Data exists. Human expertise exists. Optimisation techniques are readily available. Yet despite this, operational performance frequently remains heavily dependent upon individuals and difficult to improve systematically.

This led us to view transformation through a different lens.

Rather than beginning with algorithms, we began with understanding. Rather than asking how operations could be optimised, we first sought to understand how capability could be progressively transformed into operational intelligence.

Over time, a pattern emerged.

Capability enables understanding.

Understanding enables planning information.

Planning information supports planning decisions.

Planning decisions enable planning intelligence.

Planning intelligence enables recommendation.

Recommendation enables optimisation.

Optimisation enables operationalisation.

Operationalisation enables continuous improvement.

Perhaps the most important lesson to emerge from this progression was that successful transformation should be viewed not as an event, but as an evolutionary process.

Confidence develops incrementally.

Trust develops incrementally.

Intelligence develops incrementally.

And perhaps most importantly, operational risk is reduced incrementally.

The Capability to Operationalisation Framework emerged not as a theoretical construct, but as the consequence of repeated observations and practical experimentation.

PART THREE – THE OUTCOME

11. What Organisations Gain

One of the more interesting observations to emerge from the investigation was that value itself exists at several levels. Whilst optimisation and artificial intelligence often attract the greatest attention, our experience suggests that the most immediate benefits are frequently realised much earlier in the journey.

Perhaps the greatest challenge facing many organisations is not the absence of sophisticated algorithms, but the difficulty of maintaining consistency, confidence and control within increasingly complex operations. Expertise becomes concentrated within individuals, decisions become difficult to explain and standards often rely upon memory rather than systematic processes.

By introducing a planning intelligence layer, organisations immediately gain something far more fundamental than optimisation.

They gain confidence.

Even at its most basic level, planning intelligence provides flexibility and control through the management of known capabilities and constraints. Vehicles, drivers, locations, regulations, customer requirements and operational policies become represented explicitly rather than existing solely within the experience of individuals. Expertise that was previously concentrated within a small number of planners becomes accessible across the wider organisation.

This creates consistency. Processes become repeatable. Standards become embedded within operations rather than relying upon memory. Compliance becomes part of the planning process rather than an afterthought. Perhaps most importantly, organisations gain confidence that decisions are more consistently feasible and operationally sound.

Over time, this confidence creates the conditions necessary for a second level of capability.

12. Assisted Intelligence and Operational Insight

As planning intelligence matures, organisations begin to benefit from a level of insight that is difficult to achieve through traditional approaches alone.

Human planners remain central to decision-making. However, increasing volumes of information and growing complexity mean that opportunities frequently remain hidden. Alternative scenarios cannot realistically be evaluated and improvement opportunities often go unnoticed.

Rather than replacing human expertise, it augments it. Recommendations, insights and alternative scenarios become visible, allowing planners to make better informed decisions whilst retaining control. Importantly, these recommendations operate within the framework of known capabilities and constraints, helping ensure suggestions are feasible, compliant and operationally meaningful.

Operators are no longer relying solely upon experience and intuition. They are supported by insights that would previously have been impossible to obtain. Improvements become continuous rather than occasional. Decisions become increasingly informed. Confidence grows because recommendations can be trusted and understood.

Perhaps most importantly, every insight delivered represents an improvement relative to the information and circumstances that exist at that particular moment in time. Perfection is neither expected nor required. Instead, organisations benefit from a continuous process of learning and refinement.

13. Future Capability and Continuous Evolution

Although the introduction of planning intelligence delivers immediate value, the journey does not end there.

Artificial intelligence and machine learning represent exciting opportunities, but they should not be viewed as destinations in themselves. Rather, they represent further stages in the progressive

operationalisation of planning intelligence.

Perhaps one of the most important outcomes of establishing a planning intelligence layer is that organisations become ready for these technologies. The information, constraints, decisions and expertise required to support advanced capabilities already exist within the framework. This creates a foundation upon which increasingly sophisticated techniques can be safely introduced.

Making constraints and priorities explicit can improve planning consistency before advanced optimisation is introduced. As planning intelligence matures, organisations can explore recommendations, optimisation, machine learning and autonomous decision support where those techniques fit the operation.

It creates future readiness.

Businesses become better positioned to respond to technological uncertainty because they possess the structures required to adopt new capabilities as they emerge. Artificial intelligence becomes an evolutionary extension of planning intelligence rather than a disruptive event.

In many respects, this may represent the greatest long-term benefit.

The objective is not simply to optimise today's operation.

It is to ensure that the organisation remains capable of exploiting the opportunities of tomorrow.

Conclusion

The investigation began with a question about the application of optimisation and artificial intelligence. It led instead to a question about how organisations understand, represent and improve the decisions on which operations depend. The missing planning intelligence layer connects human experience, operational constraints and computational capability. Developing that connection progressively offers a route to greater confidence, more useful recommendations and continuous improvement without treating human expertise as something to be replaced.