karmaflow

CASE STUDY · INTELLIGENCE LAYER

A PhD-level data scientist, built into the organisation

How OASIS put its network on a knowledge graph and gave it an agent that reads the whole organisation, finds risk before it surfaces, and returns decisions rather than dashboards.

CLIENT	SECTOR	WORKFORCE LAYER
OASIS	Provincial member services network	Intelligence

OVERVIEW

Executive summary

OASIS represents a province-wide network of agencies supporting people with developmental disabilities. Its leaders coordinate through shared lists where compliance questions, policy decisions, and operational pressure all surface as ordinary messages. The signal in that traffic is rich. Reading it well, continuously, is the hard part, and the skill required to do so is scarce and costly.

So OASIS put its organisation on a knowledge graph and gave it a data scientist agent: a teammate that reads the entire network of people, topics, and threads, and returns plain-language decisions instead of more reports. In one recent fortnight it proved the point. A wave of provincial compliance and coverage questions arrived across several lists at once. To anyone reading their own inbox it felt like a busy two weeks. To the agent it was something specific: an emerging, cross-functional risk forming across executive leadership, HR, and finance, detected against a quiet baseline and surfaced while it was still unfolding.

~100 hrs / week

OASIS's estimate of the specialist data-science work this capability would otherwise require. This is not hours removed from an existing process. OASIS was not doing this work before, because the capability did not exist to be staffed.

The enterprise problem

In any large organisation, the real operating picture does not live in the quarterly deck. It lives in the messages, threads, and requests moving between people every day. That is where risk first appears, where processes quietly break or get worked around, and where a few people end up absorbing far more coordination than anyone has noticed.

The raw material is already there. What is missing, in almost every organisation, is a continuous, high-quality way to read it: to see themes forming across silos, to tell a genuine shift from a noisy week, and to turn that reading into a decision about who needs support and what needs to change. A COO could hire a small data-science team to do this in occasional deep dives, but the skill is rare, the cost is high, and the work is never continuous. OASIS took a different path. Rather than commission another report, it built the data scientist into the platform.

The trigger

Over one recent fortnight, OASIS member agencies began raising a cluster of provincial compliance and coverage questions: obligations that were shifting, how they interacted with billing and liability, and the policy and training implications that followed. The questions did not arrive under a single tidy heading. They appeared as separate threads on the executive list, parallel chains in finance, and related questions in HR and training.

To each person reading their own inbox, this registered as a hectic stretch. Across the network, with no one positioned to see all of it at once, it registered as nothing in particular.

What the agent saw

To the agent, it was not more email. It was a pattern. Reading the network as a graph, the agent compared the fortnight against a quiet prior baseline and saw the compliance theme rise from near-zero to a dominant share of activity, an order-of-magnitude shift concentrated in days. More importantly, it saw that the shift was not contained in one place. The same cluster of topics was being discussed in parallel by executive leadership, finance, and HR, with shared participants and shared language tying the threads together.

The agent named it for what it was: a single cross-functional risk forming across the organisation, not three separate busy patches. It surfaced this in near real time, rather than weeks later through anecdote.

EMERGENCE OVER TIME

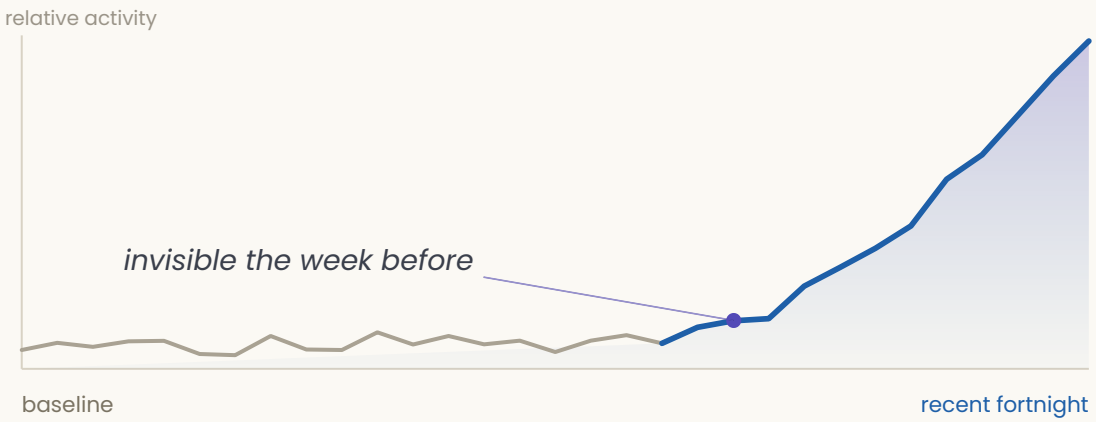

A theme that was invisible the week before becomes the dominant signal. Vertical axis is a relative index of activity, not a count of messages.

Flashpoints and smouldering issues

The agent then looked at how each thread behaved: how many people joined, how fast the first reply arrived, how the activity was distributed through the day. That let it separate two very different things. Some threads were flashpoints, drawing an immediate swarm of senior people and replies within minutes, often spilling into the evening. Others were smouldering, important questions about policy and training that sat for hours before anyone answered, despite carrying real operational weight.

For a COO, that distinction is the insight. It shows, at a glance, where the organisation's attention was correctly aligned to risk and where it was lagging behind it.

Workload and strain

Because the organisation was modelled as a graph, the agent could also see who was sitting at the centre of the response. Combining how central a person was to the critical threads with when they were active and how substantial their messages were, it surfaced a structural fact: a small number of people were absorbing a disproportionate share of the coordination, much of it after hours and on the most demanding questions.

Reported anonymously and structurally, that is exactly the signal leadership needs. It points to where support should be added or load redistributed, before quiet effort becomes burnout, and it makes visible the informal coordinating roles that traditional reporting never captures.

How it works, in plain language

The organisation's operating reality is modelled as a knowledge graph: people, teams, threads, topics, and time, with the relationships between them. The data scientist agent reads that graph continuously, using graph data science, the discipline of centrality, community detection, embeddings, and change-point detection, to find what is forming, what is anomalous, and who is connected to what.

It does not stop at "activity went up." It explains why, shows which parts of the organisation are pulled in, and says what to do next. Leaders stay on top of the loop, deciding and acting. The agent does the monitoring and analysing beneath them, work no one has to remember to start and no one has to learn to run.

THE CROSS-FUNCTIONAL CLUSTER

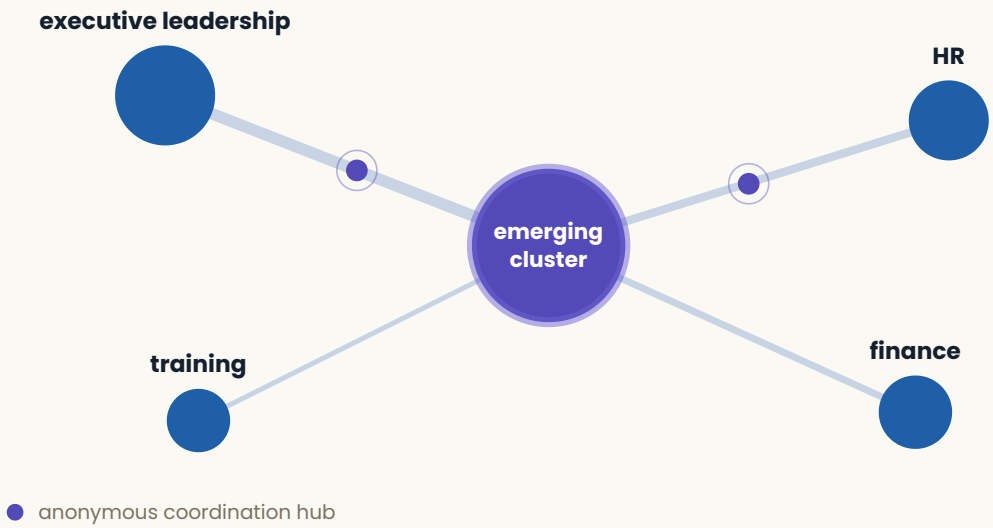

One class of issue, surfacing in parallel across functions. Node size reflects relative involvement; edge weight reflects shared participation. Hubs are anonymous.

Impact

Before this, answering the simple but vital questions ("did something actually change, or is this just a loud week?", "who is overloaded?", "is this one issue or three?") would have meant a specialist exporting data, categorising it by hand, building a one-off analysis, and explaining it, if the organisation had the people to do so at all.

OASIS does not. The capability is net-new. It estimates that putting this on the payroll, as continuous monitoring and analysis of the full communication graph, would take on the order of 100 hours of specialist data-science work every week. That skill is scarce and expensive, which is precisely why the work was never being done. Now it runs continuously inside the platform, and that time is spent deciding and acting instead. The point for a COO is structural. This is not a discount on an existing line item. It is a capability that did not previously exist to be bought.

Generalising across the enterprise

The OASIS story sits in member services, but the pattern is general. The ingredients do not change: put the operating reality on a knowledge graph, give an agent the ability to read it with graph data science, and let it monitor, analyse, and interpret patterns into recommendations.

In customer operations, the same agent detects emerging pain points across channels, routes them, and tracks frontline strain. In incident management, it maps threads and on-call activity to find systemic issues and under-resourced teams. In supply chain, it surfaces cross-functional risk, from supplier delay to quality issue, before it lands as a missed target. In any function where people coordinate through messages, the organisation gains a teammate that reads all of it.

IN CLOSING

"When your organisation lives on a knowledge graph and an agent can read it with the judgement of a data scientist, you stop guessing at what is happening. You watch the organisation think, in real time, with a data scientist beside you."