Controlant

From real-time data to real intelligence

How AI is reshaping the pharmaceutical cold chain



*This whitepaper draws on insights shared by **Gísli Herjólfsson** (CEO at Controlant), and **Saddam Huq** (Director of Cold Chain & Logistics at GSK) 15 April 2026 at LogiPharma EU, Track B: Delivering Next-Gen AI Supply Chains.*

Executive summary

Pharmaceutical cold chains have undergone a quiet revolution over the past decade. What began as a cautious move from passive temperature loggers to real-time IoT devices has matured into something more consequential: a data-rich foundation on which artificial intelligence can now operate at scale. The question facing supply chain leaders today is no longer whether to digitize, but how to convert the resulting data into decisions.

Controlant and GSK work together to operationalize real-time monitoring. This paper draws on a candid, on-stage conversation at LogiPharma EU 2026 between Gísli Herjólfsson, CEO of Controlant, and Saddam Huq, Director of Cold Chain & Logistics at GSK. Their dialogue surfaced three themes that we believe will define the next phase of pharma supply chain management.

First, the industry has crossed a threshold on data quality. With over 99.99% data availability on real-time-enabled lanes, auto-release decisions, intervention, and predictive modelling are no longer theoretical: they are routine. Operators are reporting up to 95% reductions in product waste due to temperature excursions and up to 50% reductions in lead times.

Secondly, dashboards are reaching the end of their useful life. “Big data” has become big fatigue. AI-powered interfaces where supply chain professionals ask questions of curated data and receive analysis in minutes, not weeks, are emerging as the new default.

Thirdly, the most valuable AI use cases are not glamorous general-purpose assistants. They are focused applications: generating dynamic lane risk assessments, producing daily risk scores for shipments moving through conflict zones, and optimizing thermal packaging against multi-year seasonal data.

What stands between the industry and broader adoption is not technology but validation. The path forward lies in treating AI agents as tools that extract and structure validated data, compressing the work of qualification without compromising its rigor.

For more information:
controlant.com

Introduction: the new imperative for cold chain

Every pharmaceutical supply chain leader knows the statistics by now:

- **Global spend on cold chain logistics continues to grow.**
- **Biologics, cell and gene therapies, and high-value vaccines are taking an ever-larger share of pipelines.**
- **Regulators are tightening expectations around chain of custody, temperature integrity, and audit trail.**
- **Patients are expecting shorter lead times.**
- **Boards are expecting lower cost and lower carbon per shipment.**

For most of the last two decades, the industry has tried to meet those expectations with a combination of over-engineered packaging, conservative routing, and manual after-the-fact data review. It worked, but it was expensive, slow, and fragile. Product was wasted. Lanes were qualified on historical assumption rather than real-time evidence.

And when disruption arrived — a pandemic, a conflict zone, a backlog at a transit airport — the sector reacted rather than anticipated.

Real-time IoT monitoring changed the shape of the problem. Cloud-based systems of record changed the scale of it. Now, AI is changing the speed of it.

Questions that once took weeks of analyst time can be answered in under a minute. Risk reports that once required an internal working group can be delivered as a morning briefing. Packaging decisions that once relied on a single annual qualification exercise can be revisited lane-by-lane, season-by-season.

This whitepaper examines where the industry has been, where it is now, and what the next two to three years are likely to bring. It is written for pharma supply chain leaders who are past the point of asking whether real-time data matters and are now asking how to turn it into durable operational and commercial value.



Controlant

From passive devices to real-time IoT: a decade of change

The shift from passive temperature loggers to real-time IoT devices looks obvious in hindsight. It was anything but at the time.

Controlant was the pioneer of real-time monitoring, and in 2018 one of its customers became the first pharmaceutical company to commit to real-time monitoring across its entire supply chain network.

Then COVID-19 arrived, and within two years, five additional large pharma companies had made the same switch. Few technology adoptions in pharma have moved so far, so quickly.

> **"Pharma will boldly go where someone else has gone before. It took us six years to convince the first pharma company to move from passive devices into real-time devices in their entire supply chain network."**
>
> *— Gísli Herjólfsson, CEO, Controlant*

To appreciate why the change was so consequential, it helps to remember the starting point. Before real-time monitoring with IoT devices, temperature data lived inside a logger that travelled with the product and was retrieved only after arrival, if at all. Industry estimates at the time suggested only around 50% of data made its way back to the pharma company.

Location data, chain of custody, and real-time intervention simply did not exist at any meaningful scale. Analysis meant spreadsheets.

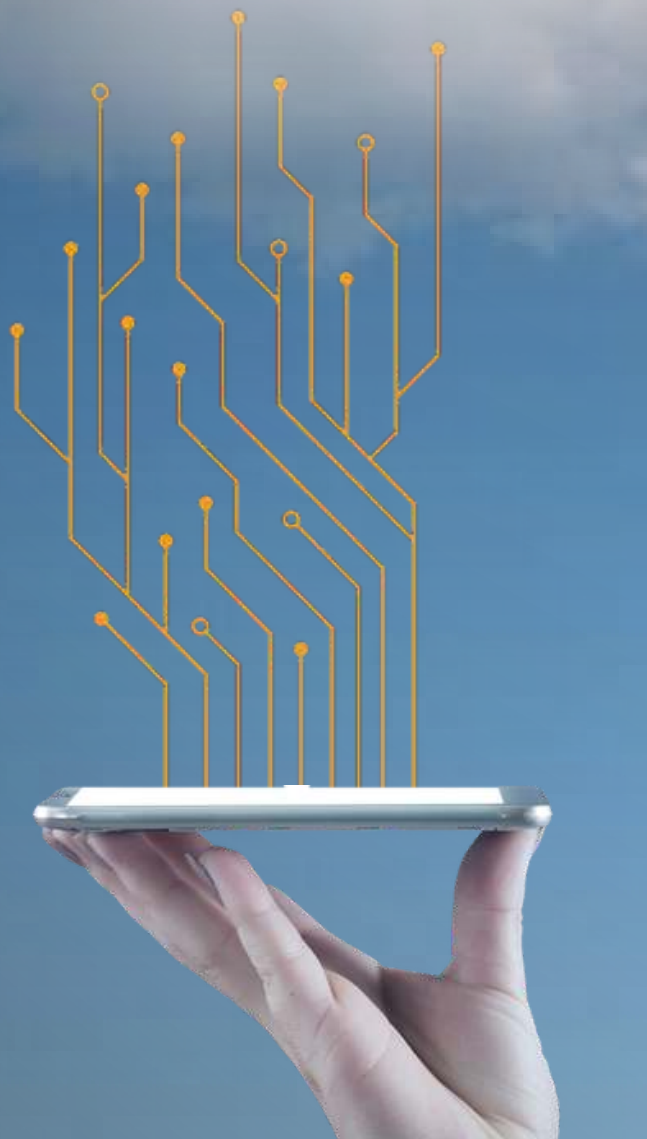
Real-time IoT devices, streaming into cloud-based systems of record, changed all of that. Data was no longer dependent on a human to gather it. Dashboards became active rather than retrospective.

During the pandemic, the most mature operators were able to see airport congestion, dwell-time spikes, and lane-level performance as it happened — and adjust routing while shipments were still in flight.

The operational consequences were significant. When data availability crossed the 99.99% threshold, previously aspirational capabilities became routine: temperature-based auto-release, real-time intervention on excursions, and continuous lane qualification.

On fully real-time networks, operators today are reporting up to 95% reductions in product waste due to temperature issues during transportation, and up to 50% reductions in lead times where auto-release has been embraced.

Those numbers matter. But they also set up the next problem.



From data overload to intelligence

The more sensors a supply chain deploys, the more data it produces. The more data it produces, the more dashboards it builds. The more dashboards it builds, the harder it becomes for any one person to actually make a decision from what they are seeing.

> **"I've seen more dashboards than I can count. The most common problem isn't the lack of data, but the lack of an efficient way to cut through the noise."**
>
> *— Gísli Herjólfsson, CEO, Controlant*

This is the paradox of cold chain digitalization. The more successful the data capture, the more acute the cognitive burden. A supply chain with five hundred active lanes, each producing continuous telemetry, simply cannot be managed through human inspection of dashboards.

Yet for many organizations, that is still how critical decisions are escalated and made.

The response from leading operators has been to invest heavily in the layer between raw data and human decision-making.

Three characteristics define that layer:

Validated data: Every location reported is qualified. Every device is calibrated. Every telemetry stream is reconciled against expected behavior. AI on top of unvalidated data is worse than no AI at all, as it simply accelerates the rate at which bad recommendations reach the operator.

Curated context: Lanes, airports, hubs, carriers, and packaging configurations are described with enough semantic context that AI can reason about them, not just pattern-match on them. This is the unglamorous work of turning time-series data into a supply chain model.

Actionable insights: Outputs are framed as recommendations or decisions, not as raw numbers. 'Downgrade packaging on lane X during months Y–Z' is useful. 'Here is a chart of temperature excursions by quarter' is not.

The transition from *data* to *intelligence* is sometimes described as a technology problem. It is really a discipline problem that requires supply chain teams to agree, up front, what *good* looks like, and to invest in the data infrastructure that allows them to measure against it continuously.

The evolution of data insights:
From static files to autonomous agents 2018 → Future


Offline temp logger • USB pull only


IoT • condition + location



Airports, hubs & lane segments


Data infrastructure: ready

Hindsight

~50% of data ever captured

- Data pulled on request
- Answers were days old
- Insights already stale
- No live view of anything

2018

Insight

99.99% data availability

- Live visibility by role
- Instant deviation alerts
- KPIs tracked in real-time
- Still reactive by nature

2023

Foresight

Built for big data & AI

- Ask in plain language
- Query validated data
- Predict issues early
- Forecast before failure

NOW

Action!

Organizational transformation: frontier

- Runs validated protocols
- Pre-approved, zero delay
- Corrects within limits
- Audit trail, always on

The end of the dashboard era: AI-powered interfaces

If the last decade was about building dashboards, the next will be about replacing them.

The interface through which a cold chain professional engages with their supply chain is in the middle of a fundamental shift. In place of static dashboards, operators are increasingly able to ask questions directly of their data in natural language, and receive structured analysis, recommendations, or forecasts within minutes.

This matters for three reasons. First, it collapses the distance between question and answer. An analysis that might once have required a data engineer, a business analyst, and three working days can be produced while the user drinks a coffee. Second, it democratizes access. The supply chain manager who used to wait for a report now runs their own. Third, it changes what gets asked.

When the cost of a question falls to near zero, curiosity scales.

But AI-powered interfaces only work when the substrate underneath them is sound. The value is not in the language model; it's in what the language model has access to.

A general-purpose assistant with access to a warehouse of unvalidated telemetry is little more than a confidence-inflating hallucination engine. A focused agent with access to a curated, qualified cold chain data model is something else entirely: a tool that compresses expertise and frees the human operator to focus on judgement rather than aggregation.

The use cases that follow illustrate what that looks like in practice.



Three use cases for AI in the pharmaceutical cold chain

Use case 1: **Lane risk assessment**

Lane risk assessment has historically been a prospective exercise. Before a lane is opened, a team gathers carrier documentation, climatic data, and assumptions about handling, produces a risk score, and, if the score is acceptable, ships product. Once the lane is operational, the assessment rarely changes.

> **"You may start off with a very robust solution at your first launch. And then, based on your data, you can possibly come down to a solution that is more fit for purpose."**
>
> *— Saddam Huq, GSK*

AI changes the economics of this in two ways. First, it compresses the prospective exercise: assumptions can be tested against comparable lanes in the operator's own history or across a wider industry dataset.

Second, and more importantly, it makes dynamic assessment practical. Instead of a static risk score that ages out the moment a shipment leaves the dock, lane risk becomes a living metric, updated every time a new shipment completes, seasonality shifts, or a carrier changes hub.

In practice, this allows supply chain leaders to graduate lanes safely downward in packaging specification over time as the dynamic risk picture clarifies, and to escalate packaging or routing decisions quickly when the signal deteriorates.

The first-launch conservatism of traditional lane qualification does not go away. It simply stops being a permanent ceiling.



Use case 2: **Proactive risk monitoring in disrupted environments**

Most supply chain organizations still manage risk reactively. An excursion alert fires, a deviation is logged, a corrective action is taken. The process works, but it is, by definition, behind the event.

The more consequential use of AI is proactive:

scanning the entire in-flight shipment portfolio every morning and flagging the subset that warrants attention. Given the current global environment, that subset is rarely empty. Conflict zones, climate events, industrial action, and infrastructure failures are appearing in supply chain news with uncomfortable regularity.

A morning risk brief, generated automatically and delivered before the team arrives at their desks, changes the shape of the day.

Rather than a manager scanning a hundred shipments for anomalies, they are presented with the handful that score highest on a composite risk metric: temperature, delay, dwell time, geography, and escalation potential combined. The brief explains why each shipment is flagged and recommends where to focus attention.

> **"If this kind of report comes to our desk, it helps us to focus and prioritize where the issues are. We don't need to look at 100 shipments; maybe just look at some of the red ones."**
>
> *— Saddam Huq, GSK*

This has two downstream consequences. The first is operational: the team stops processing dashboards and starts processing decisions. The second is strategic: the organization moves from reacting to excursions to pre-empting them.

Over time, that shift compounds. Fewer shipments are at risk because more shipments are routed away from the risk in the first place.

Use case 3: **Packaging optimization**

Thermal packaging is one of the largest controllable costs in a pharma cold chain. It is also one of the most over-engineered. For good reason: the consequence of under-specifying packaging is a temperature excursion and, potentially, a destroyed batch.

So packaging decisions tend to default towards the conservative, and to stay there. Once a lane is qualified with a given solution, the incentive to revisit that decision is weak.

The result, across the industry, is a long tail of lanes running packaging specifications that are fit for purpose but expensive, heavy, and carbon-intensive. Most supply chains know this is true. Few have the bandwidth to address it.

This is one of the cleanest AI use cases in pharma. Given multi-year data from both inside and outside the thermal packaging (the former for product release, the latter to characterize what the logistics service provider actually delivered on that lane) an AI agent can identify where packaging can safely be downgraded, where it needs to be upgraded, and where a hybrid seasonal approach is justified.

At the LogiPharma keynote, a working example using real GSK data shows the pattern. A question is posed to an agent: identify lanes where packaging may be downgraded to reduce cost, or upgraded to reduce risk, taking seasonality into account and ignoring patterns shorter than three consecutive months.

Within roughly a minute, the agent returns a lane-level analysis with explicit recommendations. The operator is not being asked to trust the model blindly. They are being handed a shortlist that a human analyst might otherwise take weeks to produce. The final qualification trial still happens in the real world.

> **"If I were to crunch those numbers over hundreds of shipments of data, it would take me weeks and months. But now, if I have the data at my fingertips, it helps me to do that job in a more efficient way... This gives us the direction of where those potential opportunities are."**
>
> *— Saddam Huq, GSK*

The commercial case is compelling. Moving from an active container to a passive advanced solution, or from passive advanced to a cold corridor with phase change material, typically reduces freight cost (lighter packaging, lower chargeable weight), reduces CO_2 per shipment, and unlocks hybrid seasonal strategies that would be impractical to manage without continuous data.



What's holding pharma back and how to move forward

The industry has had the technical capability to apply AI to cold chain for several years. Adoption has lagged for reasons that are well understood to anyone who has worked in pharma: validation, change control, and regulatory caution. None of these are bad things. They are what keeps patients safe.

The practical question, therefore, is not whether to accept AI on faith, but how to operate AI within the quality framework the industry already has.

> **"I really see from a validation perspective that we spend a lot of time and effort on validating something. So, if you are able to use an AI agent to extract the data and give it to you in a format that is usable, that will replace some of your validation efforts."**
>
> *— Saddam Huq, GSK*

Huq's reframing is important. Much of what is described as "AI" in the cold chain context is, at its core, automation of data extraction, summarization, and comparison against pre-qualified reference sets. That is well within the scope of existing validation practice, provided the underlying data and models are themselves qualified.

Supply chain leaders evaluating where to start can look at these three steps:



1 Find a high-value use case

Begin with a narrow, high-value use case such as a daily risk brief for a specific trade corridor, or packaging optimization on a defined set of lanes, and measure outcomes rigorously against the status quo.

2 Prioritize data quality above AI

Invest in data quality first, AI models second. A validated, curated cold chain data model is the hardest part of the problem; once it exists, the AI layer on top of it becomes replaceable and upgradeable.

3 Position AI as an analyst, not an authority

Treat AI output as decision support, not decision authority. The value is in compressing the analyst's workload, not in replacing the operator's judgement. This framing keeps the approach consistent with pharmaceutical quality expectations.

For more information:
controlant.com

Conclusion: toward value-driven supply chains

In pharma, “value-driven supply chain” has a specific and testable meaning: a supply chain in which every shipment, every lane, and every packaging choice is continuously optimized against cost, risk, carbon, and patient outcome, and in which the evidence for that optimization is visible, validated, and auditable.

That kind of supply chain is no longer theoretical. The data substrate exists. The AI interfaces are emerging. The use cases of dynamic lane risk assessment, proactive risk monitoring, and packaging optimization are producing measurable results.

What remains is the work of implementation: defining the narrow use cases that matter, investing in data quality, and bringing validation and quality functions into the design of AI workflows from the beginning.

The companies that do this well over the next two to three years are unlikely to be the ones with the most exotic models. They will be the ones who treated AI as the final layer on top of years of disciplined work on data, and who moved first on the unglamorous use cases where the value was largest and clearest.

The next generation of pharma supply chains will not be distinguished by how much they know. They will be distinguished by how quickly they can act on what they know.

