Decision Intelligence

Native to the Snowflake AI Data Cloud

Blue Yonder, a global leader in supply chain solutions, modernized its legacy supply chain management tool using RelationalAI's Decision Intelligence system running inside its existing Snowflake data environment. The result, known internally as the Supply Network Architect KG, serves as an intelligent digital twin of the enterprise supply chain, unifying planning, logistics, and execution into one governed, AI-native tool powered by RelationalAI.

The pains

Supply chains have always operated in an environment defined by fragmentation, volatility, and sustainability mandates: geopolitical shocks, climate events, fluctuating consumer demand, labor shortages. The COVID-19 pandemic exposed even more vulnerabilities in supply chains, causing disruptions on both the demand and supply sides: limited agility, no explainability, insufficient collaboration, data management issues.

To stay on top of the supply chain food chain, Blue Yonder needed to:

- Consolidate fragmented systems and data silos into a single, governed environment.
- Improve responsiveness to disruptions while enforcing policies and sustainability goals.
- Harness AI responsibly — without costly data duplication, opaque logic, or governance risk.

Blue Yonder recognized that its on-premises tool — a C++-based SCS system built in the 1990s, decoupled from BY's platform, hard to change, costly to maintain, and not AI-ready — could no longer meet these demands.

The solution

Blue Yonder and RelationalAI and launched a three-year modernization initiative covering ~4 000 enterprise customers and all major industry verticals:
— Retail, CPG, Technology, Health & Life Sciences, and Logistics.

The challenges

Legacy complexity: 472K lines of brittle code, difficult to operate on, maintain, or extend by current teams.

Slow scenario analysis: Overnight batch processes limited responsiveness.

Inflexible models and maintenance burden: Hard-coded logic hindered adaptation to new market realities; feature additions required deep code rewrites.

Disconnected data: Manual integrations and stale insights; siloed systems hindered end-to-end visibility.

AI isolation: Separate systems for rules, optimization, and analytics living outside the governed data platform created blind spots.

When fully deployed, the SNA-KG is expected to handle retail data sets > 100 million item-locations, serve 4 000 enterprise customers, and generate >$10 M in annual data-platform consumption, establishing Blue Yonder's platform as the industry standard for AI-driven supply-chain reasoning.

50%
TCO reduction

20B
AI predictions daily

8x
Faster runtime

RelationalAI

Technology overview
Supply Network Architect Knowledge Graph (SNA-KG)

RelationalAI, building **directly on Blue Yonder's Snowflake data foundation**, created an architecture that unifies millions of supply-chain entities — suppliers, facilities, SKUs, transport modes — into a single, queryable model of the business.

The Relational Knowledge Graph

The semantic model

SNA-KG unifies Blue Yonder's supply chain data — products, suppliers, facilities, shipments — into a semantic network that captures every entity and relationship in one governed model.
This active graph serves as a digital twin of the global supply chain, providing the context and structure for advanced reasoning and optimization.

The performance layer

RelationalAI's engine executes queries directly on Snowflake data.
The team implemented graph-caching with < 200 ms response time, achieved a 10× speed-up (12 h -> 1.5 h) in consolidation workloads, and scaled from 276 item/locations to 5 million with sub-linear response times.

The declarative logic and rules engine

Built on this graph, RelationalAI's declarative language replaced ~200 000 lines of legacy C++ with compact, auditable rules.
Analysts can encode and adjust business policies — like shipment priorities or emissions limits — directly in the graph without code rewrites, accelerating innovation and ensuring governance.

The multi-reasoner architecture

Graph reasoner

Traces dependencies and identifies network bottlenecks.

Predictive reasoner

Uses graph neural networks to forecast demand and minimizes manual feature engineering.

Prescriptive reasoner

Optimizes through solver integration (CPLEX, Gurobi, NVIDIA) for network-design and routing scenarios.

The digital twin visualization dashboard

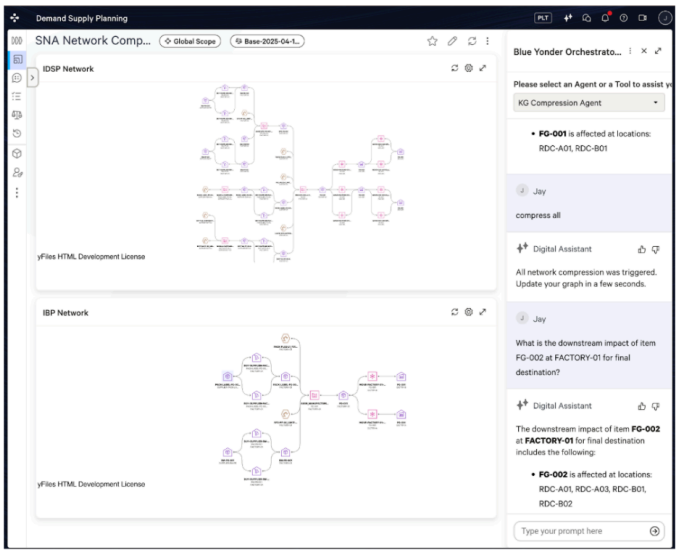

An interactive UI shows a dual-panel view of the Integrated Demand and Supply Planning (IDSP) network and the Integrated Business Planning (IBP) network. Planners can collapse thousands of detailed nodes into a higher-level view via more than 400 customizable business rules, simplifying analysis without losing traceability.

It's perfectly oriented towards supply chains, because supply chains are all about relationships between product, supply, brand, and store. It's well tuned to that, giving us really really strong AI capabilities. For language models and agents, it brings critical context and semantic meaning to the underlying data.

Duncan Angove, CEO

RelationalAI

Measurable results - BY outcomes

	LEGACY INFRASTRUCTURES	RELATIONALAI
Business growth potential	Constrained by legacy stack	**$10 M+ annual data-platform consumption potential if deployed across ~4 000 enterprise customers**
TCO	High infrastructure and maintenance costs	**Down 50% through simplified architecture and faster innovation cycles**
Scalability	Hard to extend beyond current capacity	10× improvement, supporting complex global retail and manufacturing networks
Governance	Manual processes, no lineage or auditability	Policy-aware, auditable decision logic with full traceability
Architecture	On-premises system built in Oracle and C++	Decision Intelligence embedded in Blue Yonder's platform, fully native to Snowflake
AI enablement	None—limited to static analytics	**AI-ready and able to transform reactive operations into proactive intelligence**
Performance	~12 hours per supply-chain consolidation	1.5 hours runtime with 10× scalability and real-time responsiveness
Operational efficiency	Manual scenario analysis, reactive planning	25% fewer delays, 15% lower operational cost, dynamic simulation of disruptions and routing
Visibility	Fragmented data and inconsistent reporting	Unified semantic knowledge graph providing a single source of truth across inventory, logistics, and suppliers
Sustainability	Limited visibility into carbon or policy impacts	Embedded carbon metrics, optimized for both cost and environmental goals
Deployment speed	Multi-year development cycles	**24× code reduction; fully modernized in 16 weeks with faster bug resolution and iteration**
Maintenance	Costly, fragile, and code-intensive	Low-code, declarative logic enables rapid updates and self-service extensibility

End customer outcomes

	LEGACY INFRASTRUCTURES	RELATIONALAI
Decision speed	Manual analysis and delayed reactions	**Dynamic simulations and automated recommendations cut response times dramatically**
Decision quality	Static dashboards and guesswork	**AI-driven reasoning delivering governed, explainable, and repeatable decisions**
Scalability	Hard to adapt to new markets or data volumes	10× scalability supports complex, multi-enterprise networks and global expansion
Collaboration	Fragmented teams and inconsistent data	Shared, trusted data foundation improves cross-functional alignment and supplier collaboration
Innovation pace	Long release cycles, high change risk	Low-code extensibility accelerates feature rollout and experimentation
Operational performance	Reactive workflows and frequent delays	Proactive disruption management
Visibility	Disconnected data, siloed reporting	Unified view of inventory, demand, logistics, and suppliers through a single semantic graph

Learn more at relational.ai