ARCHITECTURE OVERVIEW

THE UNIFIED DATA MODEL

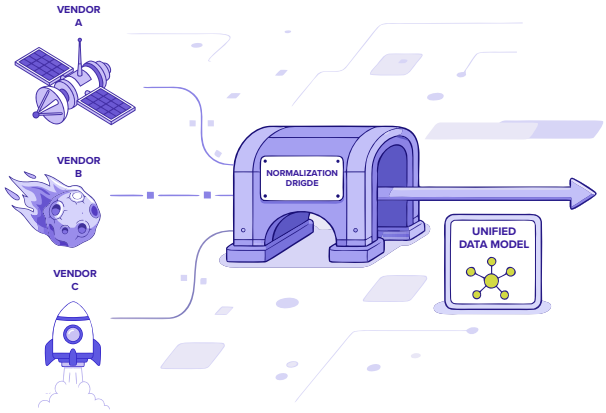

The unified data model

A single pane of glass is only as real as the model underneath it. This overview describes how WanAware ingests proprietary, vendor-native data from Ciena and Nokia optical gear, Cisco routers, Juniper edge, plus a database that documents all of it in whatever way the acquired company happened to use, and normalizes it into one canonical, cross-domain topology, spanning Layer 1 through Layer 3, rather than presenting a dashboard that merely aggregates six disagreeing sources side by side.

01. THE ABSTRACTION PROBLEM

Every vendor in a multi-vendor network exposes its own management plane, its own data model, and its own definition of what a circuit, a port, or a cross-connect actually is. Optical and transport equipment frequently still relies on TL1, a decades-old but still active management protocol. Modern routing and switching platforms increasingly expose NETCONF over YANG data models, often alongside legacy SNMP. Where streaming telemetry exists, it is commonly carried over gRPC using gNMI. None of these systems describe a cross-connect, a VLAN, or a BGP session the same way twice.

02. THE UNIFIED DATA MODEL

WanAware's platform is built around a canonical, vendor-neutral data model that spans Layer 1 through Layer 3. Every vendor-native object, regardless of the protocol it arrived on, is translated into this model before it becomes part of the topology. It is never stored as a vendor-specific exception to be reconciled later.

Want the full mission briefing?

Read the white paper on The Invisible Drain



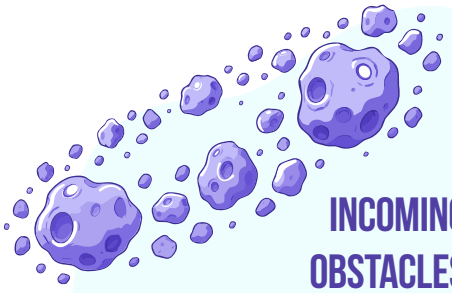

INCOMING OBSTACLES

A unified topology cannot be built by asking six different systems to answer the same question in six different languages and hoping the answers line up. That is a translation problem before it is ever a visualization problem.

Layer	What it represents	Example objects
Layer 1 (Physical)	Physical connectivity and topology	Fiber strands, ports, patch panels, cross-connects
Layer 2 (Data link)	Segmentation and switching topology	VLANs, LAGs, Ethernet trunks
Layer 3 (Network)	Routing and session topology	IP interfaces, routing adjacencies, BGP sessions

03. SPEAKING EVERY VENDOR'S NATIVE LANGUAGE

Normalization only works if ingestion meets each vendor where it already is. WanAware's discovery layer is protocol-fluent by design. TL1 mediation covers legacy optical and transport gear that has never been retrofitted with a modern management interface. NETCONF and YANG cover platforms that expose structured configuration and operational state. gRPC-based streaming telemetry, using gNMI, covers real-time interface counters, adjacency state, and session status where it is available. WanAware does not require a vendor to modernize its management interface before it can be discovered, and it does not require your team to standardize on one protocol before visibility becomes possible.

Equipment category	Native protocol(s), typically	Normalized to
Optical and transport (e.g. Ciena)	TL1, and increasingly NETCONF/YANG	Layer 1 physical topology
Nokia	NETCONF/YANG, SNMP	Layer 1 through Layer 3
Cisco	NETCONF/YANG, gRPC (gNMI), SNMP	Layer 2 through Layer 3, streaming state
Juniper	NETCONF/YANG, gRPC (gNMI), Junos XML	Layer 2 through Layer 3, streaming state

04. LAYER 1 TO LAYER 3 CORRELATION

A circuit ID in the BSS, a port assignment in the OSS, a cross-connect in a patch panel, a VLAN tag, and a BGP session are, in the network's reality, the same path described five different ways at five different layers. WanAware correlates these layers directly. Physical continuity at Layer 1 establishes what is actually connected. Data-link membership at Layer 2 establishes how it is segmented. Routing and session state at Layer 3 establishes what is actually passing traffic, and to whom. The result is a single, traceable path from patch panel to BGP session, built once and kept current continuously, rather than reconstructed manually every time someone needs to trace one circuit.

05. FROM NORMALIZED DATA TO A SINGLE PANE OF GLASS

Normalized, cross-domain data only matters if it is queryable as one thing. WanAware stores every discovered object, and every relationship between objects, physical, logical, and network-layer, in a single Relationship Graph. A query that traces a customer circuit does not need to know that it crosses a Ciena optical hop, a Cisco aggregation router, and a Juniper edge device using three different management protocols. It returns the full path as one answer, because the underlying model was never vendor-segmented to begin with.

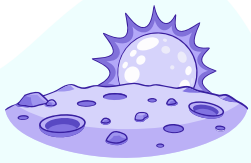

NEW HORIZONS

The distinction that matters: Aggregation asks 'where can I see all my vendor data at once.' Abstraction asks 'what is this, regardless of which vendor reported it.' WanAware is built to answer the second question, so cross-domain queries are native, not stitched together after the fact.

06. WHY ABSTRACTION, NOT AGGREGATION

Many multi-vendor management tools describe themselves as a single pane of glass while actually functioning as a single pane of many glasses, one dashboard tile per vendor, still siloed underneath. Aggregation collects vendor data in one place. Abstraction translates it into one model.

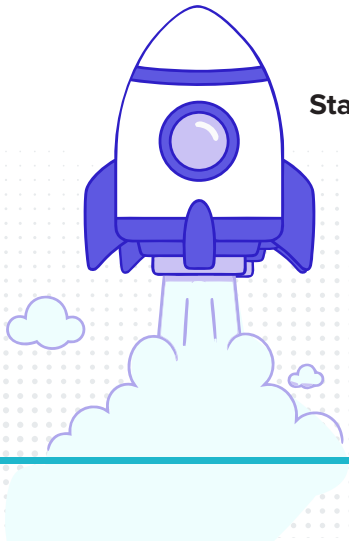

READY FOR LIFTOFF?

Start a free trial or see a demo at wanaware.com.

FREE TRIAL SEE A DEMO