ARCHITECTURE OVERVIEW

HOW WANAWARE FINDS GHOST AND PHANTOM ASSETS

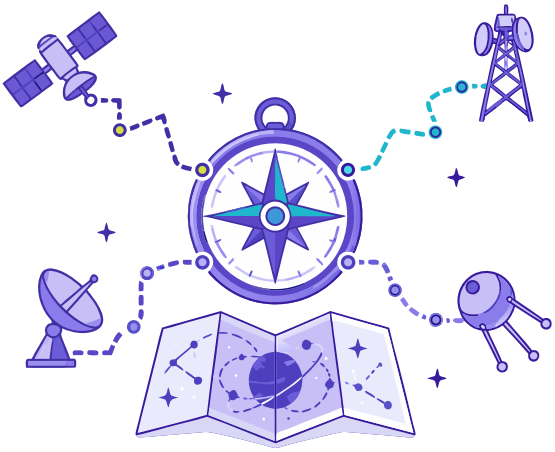

The technical architecture behind continuous ledger reconciliation, automatic ghost-asset flagging, and full visibility into untracked connected assets, without agents or scanning projects

A fixed-asset ledger is a financial record, not a technical one. It reflects what was purchased and capitalized, not what is actually deployed, powered on, and passing traffic today. WanAware closes that gap by building a live, continuously updated inventory of every **connected asset** across a financial institution's environment and reconciling it against the ledger automatically. This overview describes the architecture that makes that possible: how WanAware discovers assets without agents, how the Relationship Discovery Engine and Relationship Graph turn that discovery into a reconciled record, and how the Knowledge Discovery Engine turns a mismatch into a specific, actionable finding.

See how one financial firm stopped phantom IT leakage and reclaimed hundreds of thousands in hidden costs.

Read the case study



01. THE RECONCILIATION PROBLEM WANAWARE IS BUILT TO SOLVE

Ghost assets exist because two systems of record are never actually compared. The fixed-asset ledger is maintained by finance, updated on a depreciation schedule, and rarely touched again until an audit or a tax filing forces a review. The network, meanwhile, changes constantly: hardware is decommissioned without a ticket, licenses are reassigned without a record, and new devices are connected without ever being added to the ledger in the first place. WanAware's architecture is built to compare both realities continuously, not once a year during an audit cycle.

Integration source	What it contributes	Discovery method
Fixed-asset and GL systems	The book of record: capitalized hardware, software licenses, and their depreciation schedules.	Read-only API or file-based integration
Network infrastructure	IP subnets, VLAN membership, device population, and live traffic patterns across every site.	Passive network telemetry (NetFlow, DHCP, DNS)
Cloud and virtualization platforms	Active compute, storage, and licensing consumption across every provisioned environment.	API integration (AWS, Azure, VMware, and similar)
Identity and directory services	Which accounts and service principals are actually authenticating against which systems.	Read-only directory queries
ITSM and procurement systems	The book of record: capitalized hardware, software licenses, and their depreciation schedules.	Read-only API or file-based integration

02. DISCOVERY: AGENTLESS, PASSIVE, AND READ-ONLY

None of the discovery methods above require installing software on an endpoint, running a credentialed or active scan, or requesting standing administrative access into systems finance and IT already consider sensitive. That matters in a regulated environment, where a new scanning footprint is itself a change-control and security-review event, and where the whole point is to find assets nobody is actively managing without creating a new asset to manage in the process.

- **Passive network telemetry.** NetFlow, DHCP, and DNS signals reveal every connected asset that is actually powered on and communicating, independent of whether it appears on the ledger at all.
- **Read-only API integrations.** Cloud, virtualization, and ITSM platforms are queried directly, using access already configured for other purposes, rather than a new discovery footprint layered on top of existing systems.
- **Read-only ledger and GL integration.** The fixed-asset ledger is read, not modified, so reconciliation never puts WanAware in the position of writing back to a system of financial record.
- **No agents, ever.** Branch hardware, data-center equipment, and every other connected asset is observed from the network and from existing platform APIs, never from software installed on the device itself.

03. THE RELATIONSHIP DISCOVERY ENGINE: CONTINUOUS, NOT SCHEDULED

Discovery on its own only produces a list. WanAware's **Relationship Discovery Engine (RDE)** is the layer that keeps that list current: it continuously monitors every connected source above and updates the moment something changes, a device goes dark, a new subnet lights up, a license consumption count shifts, rather than waiting for the next scheduled scan or the next audit cycle. For ghost-asset detection specifically, that continuity is the difference between catching a piece of decommissioned hardware the week it goes dark and catching it eighteen months later when an auditor happens to ask.

- **Real-time change detection.** A device that stops responding, a license seat that goes unused, or a subnet that is reassigned is picked up as it happens, not at the next scheduled inventory pass.

- **No manual reconciliation step.** The RDE feeds the discovered state directly into the **Relationship Graph**, so there is no spreadsheet export, no manual matching, and no lag between what is discovered and what is reconciled.

04. THE RELATIONSHIP GRAPH: WHERE THE LEDGER MEETS THE NETWORK

Every asset the RDE discovers, and every asset line on the fixed-asset ledger, is resolved into the same live, queryable Relationship Graph. This is the architectural step that makes reconciliation possible at all: the ledger and the network stop being two separate records compared by hand and become two views into one graph, continuously checked against each other.

- **Ghost assets flagged automatically.** A ledger line with no corresponding live asset in the graph, hardware still being depreciated and insured that the network has not seen in months, is surfaced as a discrepancy the moment the graph confirms the absence, not assumed from a missing scan result.

- **Untracked connected assets surfaced too.** The reverse case matters just as much: a live, active asset in the graph with no corresponding ledger entry is just as much a blind spot, and just as visible, as a ghost asset going the other direction.

- **Schemaless by design.** New asset types, new leasing arrangements, and new vendor formats from either the ledger or the network become graph nodes the moment they appear, without a data-modeling project to define how they should be reconciled first.

05. THE KNOWLEDGE DISCOVERY ENGINE: TURNING A MISMATCH INTO A FINDING

A raw discrepancy between the ledger and the graph is not yet an actionable finding on its own. WanAware's **Knowledge Discovery Engine (KDE)** is the reasoning layer that evaluates each mismatch against the graph's context, how long an asset has been dark, what it was originally provisioned for, whether a related asset was recently decommissioned, and produces a specific, explainable conclusion rather than a bare list of exceptions for someone to investigate by hand.



STEP 1

Discover

The Relationship Discovery Engine continuously observes every connected asset across network, cloud, and identity sources, agentlessly.



STEP 2

Reconcile

Discovered assets and ledger line items are resolved into the same Relationship Graph, continuously, not on an audit schedule.



STEP 3

Flag

The Knowledge Discovery Engine evaluates each ledger-to-graph mismatch and flags ghost assets and untracked connected assets as specific findings.



STEP 4

Resolve

Finance and IT receive an explainable finding, what the asset is, why it was flagged, and since when, instead of a raw exception to research from scratch.

ZERO AGENTS INSTALLED

Every branch, data center, and cloud workload stays untouched throughout discovery.

CONTINUOUS, NOT ANNUAL

Reconciliation runs every day the graph updates, not once at audit time.

ONE GRAPH, TWO RECORDS

Ledger and network are checked against each other automatically, not by hand.

06. WHAT THIS ARCHITECTURE PREVENTS

Integration source	What it contributes	Discovery method
Ghost assets on the ledger	Discovered during an annual audit, if at all, months or years after the equipment went dark.	Flagged automatically the first reconciliation cycle after the asset stops appearing on the network.
Untracked connected assets	Never surfaced, because reconciliation only checks the ledger against itself, not against live network reality.	Surfaced with the same automatic confidence as a ghost asset, in the same graph.
Overpaid tax and insurance	Continues indefinitely on equipment nobody has verified is still in service.	Stopped as soon as the underlying asset is confirmed absent from the network.
Unpatched, unknown hardware	Remains invisible to security because it was never on a scan target list to begin with.	Visible in the same Relationship Graph used for reconciliation, not a separate security-only inventory.
Audit preparation	A multi-week manual reconciliation project, run from spreadsheets.	A standing, continuously current record that is ready to export the day the audit starts.

07. THE TAKEAWAY

A financial institution's ghost-asset problem is not a discipline problem inside finance or IT. It is an architecture gap: no system has ever continuously compared what the ledger says exists against what the network can actually confirm. WanAware's **agentless discovery**, Relationship Discovery Engine, Relationship Graph, and Knowledge Discovery Engine close that gap as one continuous pipeline, so every connected asset, and every asset that no longer is one, is accounted for the same way, every day, not once a year.

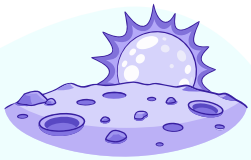

NEW HORIZONS

None of this architecture requires a scanning project, a credentialed sweep, or a rip-and-replace of existing ITAM or GL tooling. WanAware ingests the ledger and the network as they actually exist today and lets the Relationship Graph and Knowledge Discovery Engine do the reconciliation a manual process was never built to do continuously.

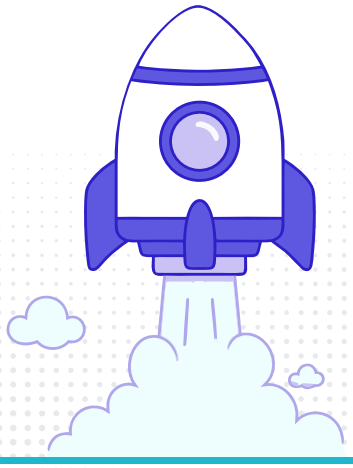

READY FOR LIFTOFF?

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

FREE TRIAL SEE A DEMO