CrossConnect by Cybr IQ

HOW IT WORKS · THE DATA JOURNEY

From the wire to an answer you can trust.

Every fact inside CrossConnect makes the same journey: it enters from the network or your keyboard, gets staged as an observation, is reconciled against what you have documented, becomes part of the source of truth, is computed on, and finally leaves as a cited answer or an alert. Here is that journey, drawn out end to end, with no packet ever captured.

THE WHOLE JOURNEY ON ONE LINE

Six stages, every time.

Whether a fact arrives over SNMP or is typed in by an operator, it passes through the same six stages. Each stage has one job, and the boundary between them is where trust is earned.

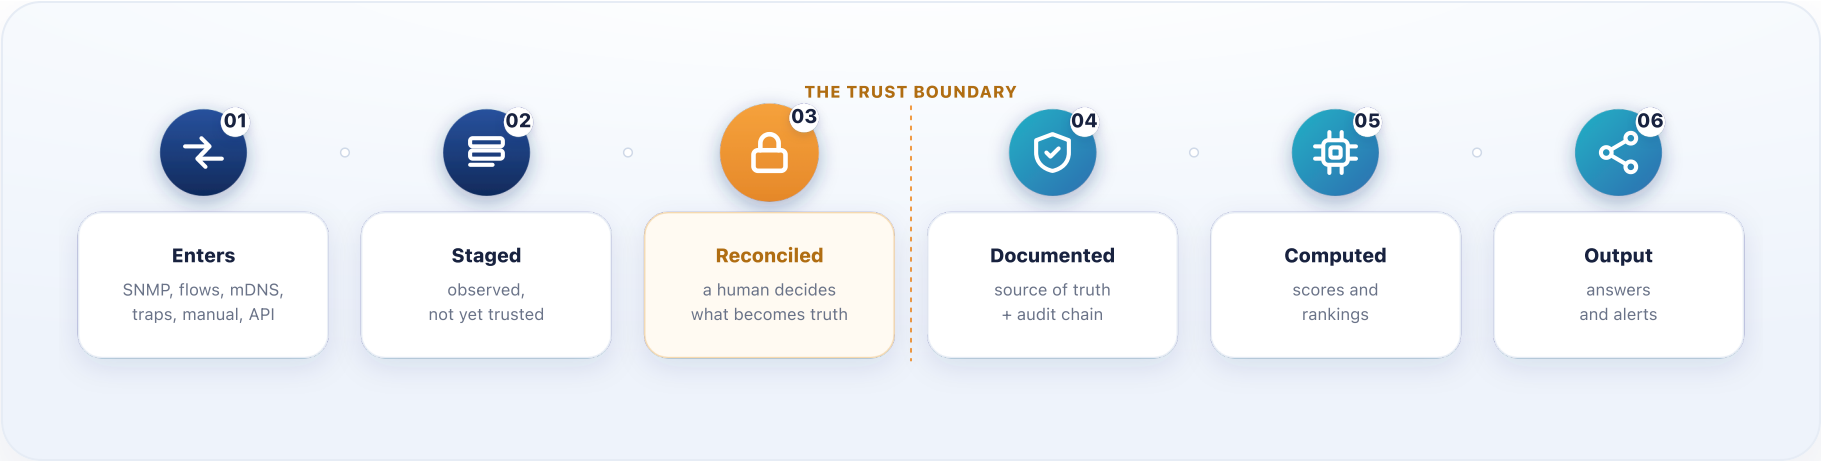

Two colors tell the whole story. Navy is what the platform has only *observed*; teal is what has become *trusted*. The single amber moment in the middle is the human trust gate, and nothing crosses it on its own.

01 ENTERS

Nine front doors, one rule: observe, never capture.

Data arrives many ways. The platform reads switch-derived signals, listens for the announcements AV gear already broadcasts, accepts traffic summaries, and takes what operators type in. It never captures a packet or inspects a payload, it reads state the network already exposes.

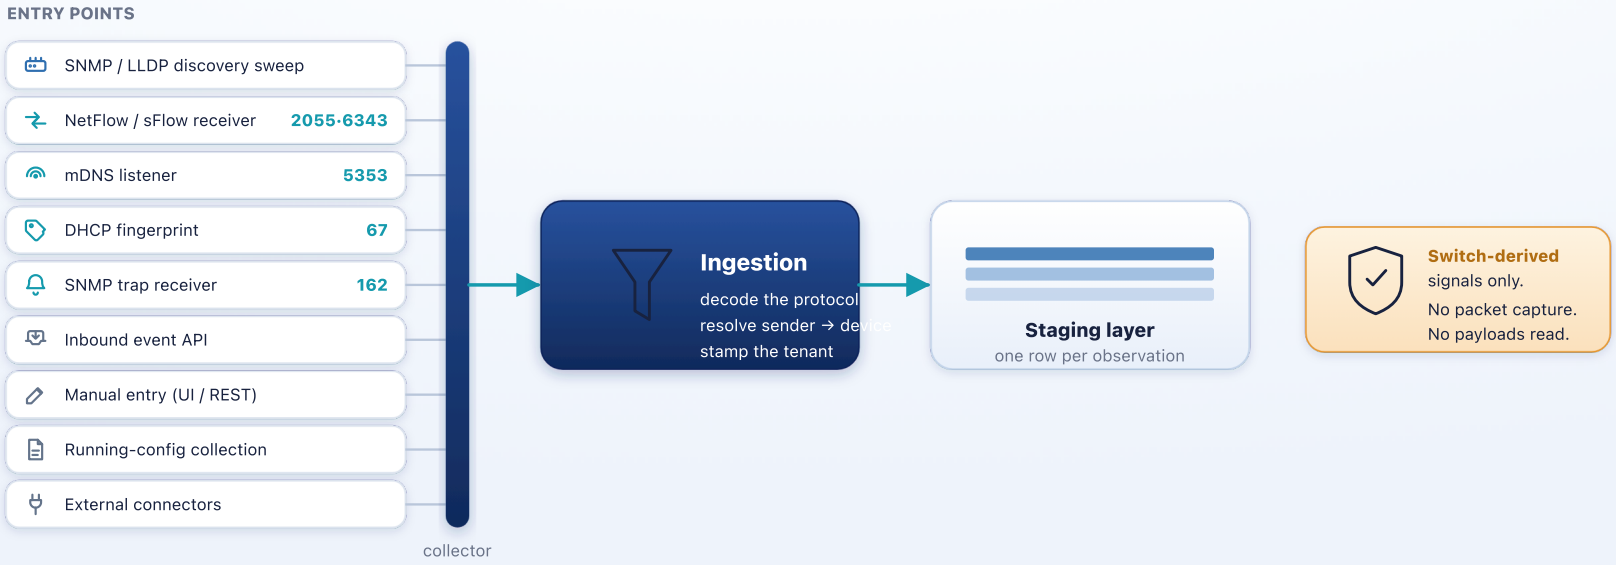

Many sources, one ingestion path. Each front door decodes its own protocol, resolves the sender to a known device, and stamps the tenant, then hands a clean observation to staging.

SNMP / LLDP sweep

The scheduled discovery worker polls each device for its interfaces, neighbors, serial, sensors, and more. The backbone of how the network is found.

NetFlow / sFlow `UDP 2055 · 6343`

Routers export traffic summaries to the built-in receiver, decoded into who-talked-to-whom, with no packet payloads ever read.

mDNS listener `5353`

Joins the multicast group and hears the service announcements AV gear already broadcasts (Dante, NDI, AirPlay), the strongest signal for typing AV endpoints.

DHCP fingerprint `UDP 67`

Reads the option-55/60 fingerprint from relayed DHCP requests to recognize a device family, corroborating AV classification.

SNMP traps `UDP 162`

Devices push real-time events (link down, PSU fail, lamp hours). Each becomes a confidence-scored observation for review.

Manual & API

Operators document gear through the UI or REST, and external systems push assertions to the inbound event API. Both are stamped with who and when.

02 STAGED

Observed, but not yet believed.

Raw observations land in append-only staging tables, one family per kind of fact. Nothing here is treated as truth. Each row carries when it was seen, the newest one per natural key is the operative one, and old rows are purged automatically. This is the platform's short-term memory of reality.

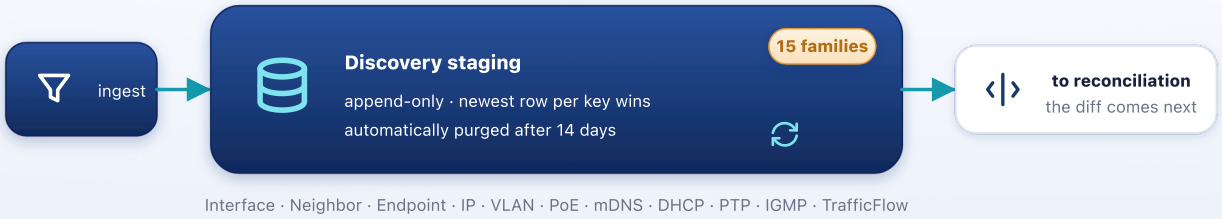

Short-term memory, kept honest. Fifteen families of observation, append-only so the diff against your records has history to show, and self-cleaning so the tables never bloat.

Append-only by design

A new sweep does not overwrite the last one, it adds rows. That history is what powers the operator-facing “what changed since we last looked” view, and the audit trail behind it.

Self-purging

A daily sweep drops observations past the retention window (14 days by default), so staging stays a rolling picture of recent reality rather than an ever-growing log.

03 RECONCILED

The trust gate: a human decides what becomes truth.

This is the most important boundary in the system. CrossConnect compares what it observed against what you have documented, scores how believable each observation is, and presents the differences as proposals. Nothing crosses into the source of truth on its own, an operator promotes it.

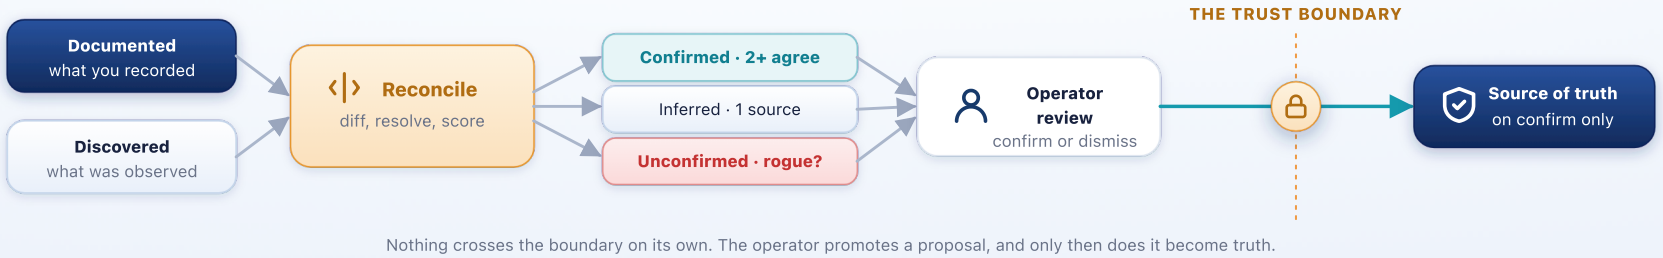

Nothing crosses the boundary on its own. The operator promotes a proposal, and only then does it become truth.

Confidence, then consent. Two independent sources that agree score Confirmed; a single source is Inferred; something that resolves to nothing is Unconfirmed and flagged as a possible rogue. The operator has the final say.

Documented vs discovered

The platform always knows which facts you recorded and which it observed, and shows the difference rather than quietly merging them.

Confidence scoring

Believability is earned by corroboration: agreement across sources lifts an observation toward Confirmed; isolation keeps it low and visible.

Promotion is a choice

An undocumented link or a new device waits as a proposal until an operator promotes it. No silent writes to the source of truth, ever.

04 DOCUMENTED

The source of truth, and a history that cannot be rewritten.

Once promoted, a fact becomes part of the canonical model: devices, interfaces, cables, addresses, circuits, and the services they deliver. And every change to that model is written into a tamper-evident, hash-linked chain, so who changed what, and in what order, can always be proven.

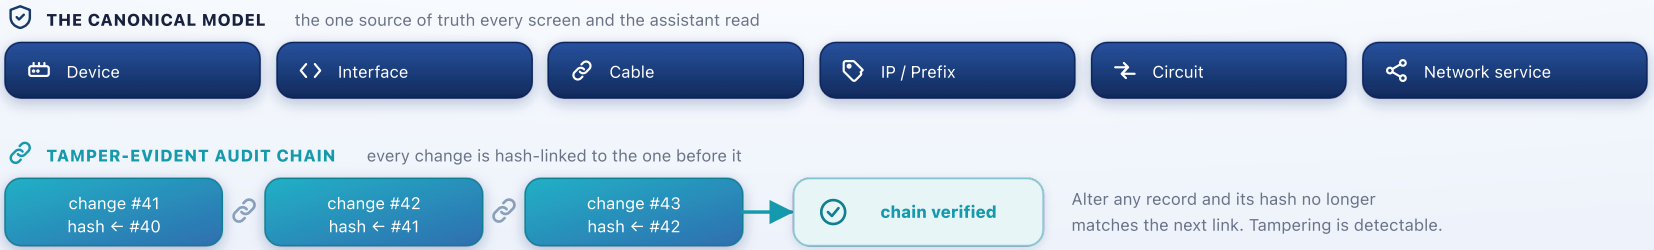

One model, one ledger. The canonical entities are what every screen and the assistant read; the hash chain underneath makes the history provable, which is what lets an auditor take the platform's word for it.

05 COMPUTED

One snapshot, many answers.

The intelligence layers do not store new facts, they compute from the source of truth. Each reads a snapshot and returns a score or a ranked list as a pure function, which is why the same evidence can power data quality, compliance, maturity, and the “what should I fix first” queue all at once.

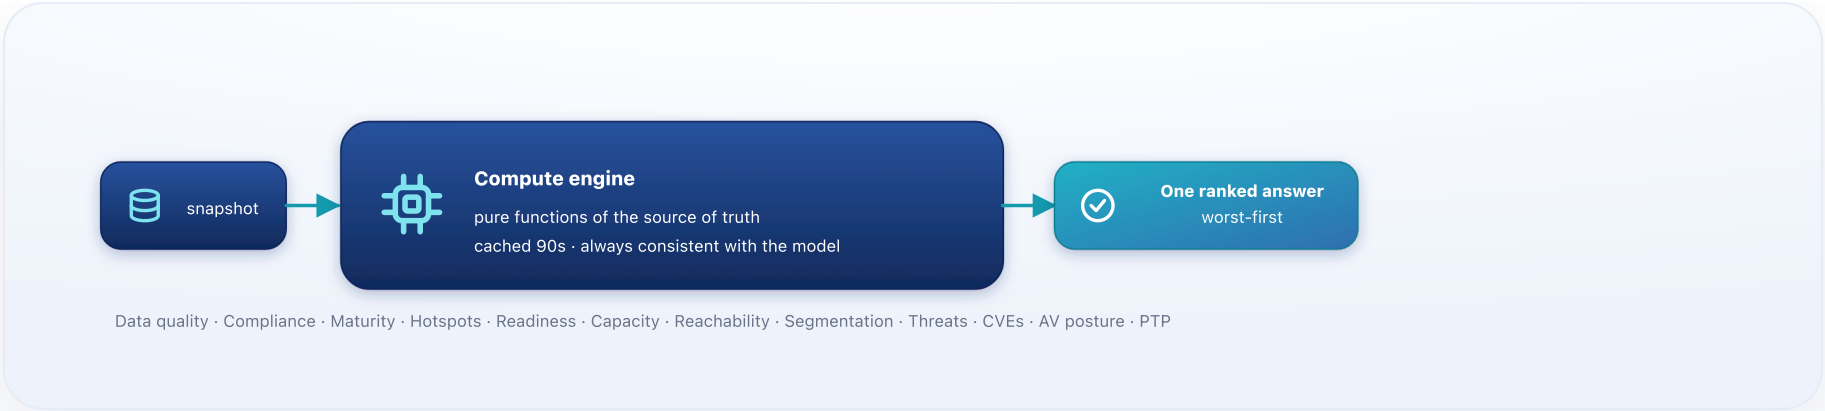

Compute, do not duplicate. Because the analytics are pure functions of the snapshot, they stay consistent with the model, and a short cache means one computation serves every viewer for the window.

06 OUTPUT

It leaves as something you can act on, and check.

At the end of the journey, data becomes an answer. The assistant replies in plain language with the records it used. Webhooks and SIEM sinks carry changes to your other tools. Reports and the API export the model. And every screen reads the same single source of truth.

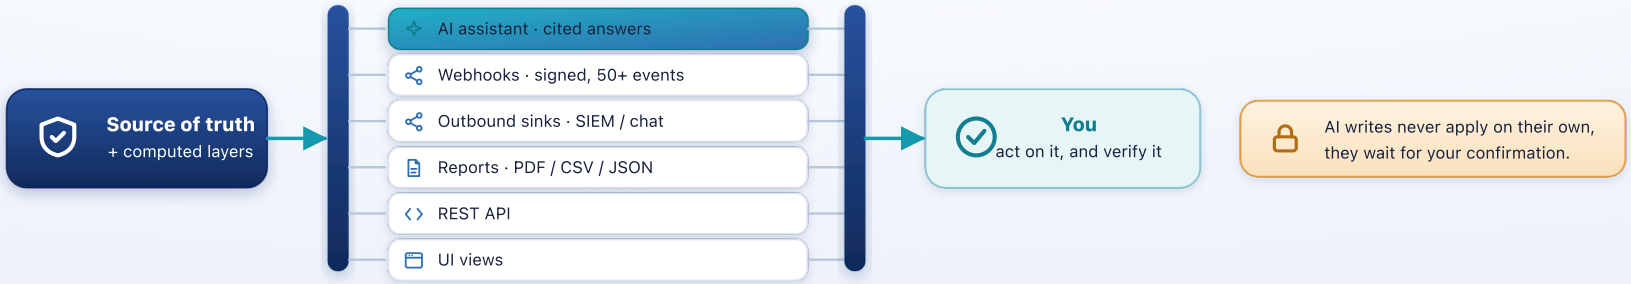

Cited, signed, exportable. The assistant proves its answer with records; webhooks and sinks are signed and safe-by-default; and the AI never changes the network without a human pressing confirm.

End to end, in one breath: a switch is polled over SNMP, its new neighbor is staged as a DiscoveredNeighbor, reconciliation flags an undocumented link, an operator promotes it to a Cable, the audit chain records the change, data quality and the topology recompute, and the assistant can now answer “what is plugged into this switch?” and cite the cable it just learned.

THE JOURNEY, IN ONE IDEA

Observe, earn trust, then answer.

CrossConnect treats every fact as an observation until a human makes it truth, records the moment it does, and computes everything else from that one trustworthy model, so the answer you get at the end is one you can check all the way back to the wire.

CrossConnect