

FOUNDATION — LAYER 2

L2 The Foundation Layer

WHERE RAW SIGNAL BECOMES CANONICAL TRUTH

L2 is where Converge360 turns raw signal into trusted truth. The layer below pours in overlapping, vendor-specific records; Foundation **resolves and de-duplicates** all of it into a single canonical set of entities. It answers two questions with authority — **who** is in your environment, and **what** is in it — handing one clean, trusted record per thing to every layer above.

WHERE L2 SITS

- L5 Experience
- L4 Operations
- L3 Intelligence
- L2 Foundation**
- L1 Ingestion & Discovery

COMPONENTS

2

Identity Fabric & Asset Registry — the who and the what.

CONSUMES

L1 raw & candidates

Normalized feeds and discovered candidates.

PRODUCES

Trusted entities

Canonical identity and asset records.

FEEDS INTO

L3

Intelligence reasons on entities, not raw logs.

Two components, one source of truth



03

Identity Fabric

THE CANONICAL "WHO"

One trusted, de-duplicated record for every human, service and machine identity — resolved from raw L1 feeds.



04

Asset Registry

THE CANONICAL "WHAT"

One trusted record for every asset — its owner, exposure and business criticality — built from feeds and discovered candidates.

Why the foundation matters

Everything above L2 reasons on **entities**, not raw logs. If those entities are duplicated, fragmented or missing context, every judgment, correlation and action inherits the error. Foundation exists to prevent that: it collapses many noisy records into one authoritative entity, attaches the context that makes it meaningful, and becomes the **single source of truth** the rest of the platform trusts without question.



One record, no duplicates

Fragmented signals from many tools resolve into a single canonical entity per thing.




The who and the what

Identity and assets are modelled side by side — so people and machines connect to things.



Trusted by everything above

Layers 3 to 5 build on these entities, never on the raw feeds beneath them.



COMPONENT 03 · FOUNDATION

Identity Fabric

The canonical "who" — one record per identity

Identity Fabric is the canonical "who" — **one trusted record for every human, service and machine identity.**

A single identity rarely lives in one place. The same person, service account or machine appears across IAM, endpoint, cloud and a dozen other feeds, each with its own identifier and spelling. Identity Fabric consumes those raw L1 feeds and resolves the duplication away — matching, linking and merging fragmented records into one authoritative identity per actor. The result is a set of trusted identity entities: de-duplicated, attributed and ready for every layer above to reason about who did what.

■ KEY CAPABILITIES



Resolution & de-duplication

Matches and merges fragmented records into one canonical identity per actor.



Human, service & machine

Models every identity type — employees, service accounts and machines alike.



Cross-source correlation

Links the same identity across IAM, endpoint, cloud and more.



Authoritative records

Produces one trusted, attributed identity that downstream layers can rely on.

■ IDENTITY TYPES IT RESOLVES

HUMAN · SERVICE · MACHINE

Human users

Service accounts

Machine identities

Privileged identities

Federated / external

Non-human / workload



From many records to one identity


A single human can appear as a dozen records across your tools. Identity Fabric collapses them into **one canonical identity** — so Intelligence never double-counts, and Operations always acts on the right actor.

← CONSUMES

- Raw normalized feeds from L1
- Identity signals across many vendor sources

→ PRODUCES

- Trusted identity entities — de-duplicated & attributed
- Canonical "who" for L3 Intelligence



COMPONENT 04 · FOUNDATION

Asset Registry

The canonical "what" — one record per asset

Asset Registry is the canonical "what" — **one trusted record for every asset, with its owner, exposure and business criticality.**

Assets are scattered and short-lived: a server in one feed, a discovered host from a scan, a cloud workload that spun up an hour ago. Asset Registry consumes both raw L1 feeds and the discovered candidates surfaced by Discovery Tools, then resolves them into one trusted record per asset. Each record is enriched with the context that makes it actionable — who owns it, how exposed it is, and how critical it is to the business. The output is a set of trusted asset entities the rest of the platform can prioritise against.

■ KEY CAPABILITIES



Asset resolution & de-duplication

Reconciles feeds and discovered candidates into one record per asset.



Ownership mapping

Attaches an accountable owner to every asset on record.



Exposure tracking

Captures how reachable and at-risk each asset is.



Business-criticality scoring

Grades each asset by its importance to the business.

■ WHAT EACH RECORD CAPTURES

Owner

Exposure

Business criticality

Type / category

Environment

Lifecycle state



The other half of the entity graph

Identity Fabric answers **who**; Asset Registry answers **what**. Modelled together, they let every layer above connect people and machines to the assets they touch — the backbone of meaningful detection and response.

← CONSUMES

- Raw normalized feeds from L1
- Discovered candidates from Discovery Tools


→ PRODUCES

- Trusted asset entities — owned, scored & resolved
- Canonical "what" for L3 Intelligence

HOW L2 WORKS AS ONE FOUNDATION

Two questions, one trusted answer.

Identity Fabric resolves the who; Asset Registry resolves the what. Both consume the raw material from L1, collapse the duplication, attach the context — and hand a single, trusted set of canonical entities to Intelligence above. From here on, every layer reasons on clean entities, never on raw feeds.



CROSS-CUTTING · AGENTIC AI / LLM FABRIC

Resolution is hard — the fabric makes it sharper.

The agentic fabric threads through all five layers, including this one. At L2 it helps **match and merge duplicate records**, **infer ownership** and **score business criticality** — turning judgement calls that once needed an analyst into fast, consistent entity resolution.

COMPONENT	CONSUMES	PRODUCES	DOWNSTREAM
Identity Fabric	Raw normalized feeds from L1	Trusted identity entities	→ L3
Asset Registry	Raw feeds + discovered candidates	Trusted asset entities	→ L3



Single source of truth

One canonical record per identity and asset, trusted by every layer above.



Resolved & de-duplicated

Overlapping, vendor-specific records collapse into one clean entity.



Context-rich entities

Ownership, exposure and criticality travel with every record.



Foundation turns raw signal into trusted truth — one canonical record for every who and what the platform reasons about.