

CONCEPTS

Database Schema Overview

The IDN platform is built on a deeply normalized relational schema designed for Computer-Aided Facilities Management. It models every operational entity a facilities organization touches — from portfolios and buildings down to individual work-order line items.

At a glance

11 Modules	333 Tables	12,000+ Fields	8 Core tables
----------------------	----------------------	--------------------------	-------------------------

Three layers

The schema is organized in three conceptual layers:

- **Modules** — functional domains (Portfolio, Work, Equipment, EHS...). A grouping, not a database object.
- **Tables** — the entities themselves (e.g. `building`, `work_order`, `asset`). 333 in total.
- **Fields** — typed columns on each table; 12,000+ across the system with consistent naming and relationship conventions.

The 11 modules

Portfolio	Sites, buildings, floors, rooms
People	Employees, partners, contractors
Work	Work orders, PM, SLAs, dispatch
Equipment	Assets, components, meters
Space	Allocation, occupancy, moves
MAC	Moves, adds, changes
EHS	Health, safety, environment
Materials	Inventory, parts, procurement
Documents	Drawings, manuals, contracts
Reporting	KPIs, dashboards, exports

System	Config, security, audit
--------	-------------------------

Why normalization matters

- **Reliable analytics** — every metric has one source of truth.
- **Predictive maintenance** — asset, work-order, and sensor history are linked, not duplicated.
- **AI-native queries** — the AI Assistant reasons over a clean graph instead of guessing across spreadsheets.
- **Compliance** — EHS, audit, and document trails stay attached to the entities they describe.

Naming conventions

Tables and fields use snake_case. Foreign keys end in _id. Lookup tables are singular (building, not buildings). Date fields use _at for timestamps and _on for date-only values.