



CMMS WITH NATIVE GIS

Functional documentation

Before a demonstration, you can read exactly what the system does.

EDITION

Public version

UPDATED

17 October 2026

WEB

maptainer.com

About this document

This is the public version of the internal functional document: no positioning notes, no development status and no internal names for fields, permissions or modules. What remains is what the system does. We have added nothing that was not already there.

How to read it

Each chapter opens with a three-line summary: which module it is available in, which surfaces it is used from (web, mobile, customer portal or public API) and what it is for. The detail follows, section by section.

Chapter and section numbering is that of the complete functional document, and it is the numbering used by the cross-references in the text (“see 5.9”).

If a question is not answered here

info@maptainer.com · (+34) 960 659 055 · maptainer.com

Contents

–	Executive summary	4
01	Authentication, users and permissions	5
02	Multi-tenant, verticals and subscription	13
03	Organisational structure	21
04	Asset management (GIS)	25
05	Work orders	31
06	Recurring scheduling of work orders	43
07	Incidents	46
08	Maintenance types, tasks, trades and skills	51
09	Staff, shifts and availability	54
10	Eligibility and smart dispatch	58
11	Multi-day work order planner	61
12	Time budget (effective duration)	65
13	Vehicle fleet	67
14	Telemetry and GPS tracking	70
15	Geofence and field presence recording	74
16	Inventory, warehouses and purchasing	77
17	Meter reading	81
18	Dynamic forms	84
19	SLA	87
20	Dashboard and analytics	90
21	Mapping and GIS infrastructure	97
22	Customer portal	100
23	Mobile app	103
24	System alerts	108
25	Public API, integrations and data governance	111
A	Annex · Modules, permissions and limits	121
A.1	How access is governed	121
A.2	Complete matrix	121
A.3	Default permissions of the system roles	123
A.4	Subscription limits	123
A.5	Subscription validity	124

Executive summary

Maptainer is a **CMMS platform with native GIS**: the map is not a view added on top of a maintenance manager, but the structure on which the asset inventory, work order creation, route planning, field work tracking and analytics rest.

It covers the full maintenance cycle for geographically distributed infrastructure: georeferenced inventory of the asset base → incident detection → work order creation and planning → assignment by skills and real road proximity → field execution with an offline mobile app → material consumption and cost → closure, validation and analytics.

Deployment model. Multi-tenant: each organisation (maintenance contractor, asset operator, network manager) operates in an isolated environment with its own hierarchy of customers, assets, staff, inventory and configuration. A global superadministrator manages the organisations and maintains sector catalogues (“verticals”) that are cloned when a new organisation is created.

Four usage surfaces:

SURFACE	USER	CHARACTERISTICS
Web application	Managers, planners, supervisors	Full-screen map with a side operations panel
Mobile application	Field technicians	App for iOS and Android, with full offline mode
Customer portal	The organisation's end customer	Separate mobile-friendly website, with its own credentials and roles: consultation, incident reporting and acceptance of completed work
Public API	The customer's third-party systems	API key per integration with scopes, incremental synchronisation, paginated reading, incident creation, signed webhooks, full data copy and its own documentation

In-house computing infrastructure. Unlike a CMMS that consumes third-party map services, Maptainer hosts its own geospatial infrastructure: a map service with caching, road routing, multi-vehicle route optimisation, its own base map and offline mapping, including the official PNOA/IGN orthophoto. Isolation between organisations is also enforced in the database itself.

Commercial modularity and configurable complexity. There is an explicit catalogue of modules that can be subscribed, **ten organisational capabilities** that the organisation itself switches on at no cost, **levels of detail** that grade the rigour within each process, **three starter profiles** that configure everything in one click, and a granular permission system resolved in four layers, with each organisation's own roles. The interface, emails, PDFs and manuals are in Spanish and English, and each organisation works in its own currency. The same installation presents itself as a product with a handful of screens to a six-person engineering firm and as a complete platform to an operator with branches.

01

Authentication, users and permissions

AVAILABILITY	Core, always active
SURFACES	Web · Mobile · Customer portal · Public API
WHAT IT IS FOR	It ensures that each person sees exactly what their organisation has subscribed to and what their role allows, with three separate identity systems (employee, end customer and machine) that do not cross-contaminate.

1.1 Three independent identity systems

The platform distinguishes three types of identity, each with its own credential store, its own login flow and its own session lifecycle:

- **Internal user** — an employee of the organisation. Accesses the management web application and the mobile app.
- **Portal user** — a person from the end customer. Accesses only the customer portal, from the browser. Completely separate credentials.
- **API key** — a machine identity for the public API. No associated user.

The separation is enforced on the server, not just in the interface: internal access (web, mobile, files and mapping) rejects portal sessions even when they are valid, and a portal user cannot register with the email address of an internal user.

1.2 Web login

Login with email and password. The login email is compared without regard to case: “Juan.Perez@...” and “juan.perez@...” are the same account.

- **Session duration:** without “remember me”, the session expires after **90 minutes of inactivity**, not 90 minutes after logging in. Only the user’s real activity —keyboard, mouse, taps— renews it, not automatic background refreshes. With “**remember me**” it lasts 30 days.
- **Single session per surface:** if the same account logs in from another browser, the previous session is invalidated with the explicit message “your session has expired because you have logged in on another device”. **Web and mobile sessions are controlled separately:** logging in on mobile does not log the user out of the web, and vice versa. The same rule applies to file and map downloads, so that a revoked session stops working everywhere at once.
- **Temporary lockout after failed attempts**, not deactivation. Consecutive failures lock the account for an increasing time (from one minute to half an hour), with additional limits per source address, and a burst of simultaneous attempts gets no more attempts than the maximum. The lockout lifts

on its own; the administrator can also view and unlock locked accounts. When an account is locked, its holder receives an email warning.

- **Uniform response:** incorrect credentials receive the same message whether or not the email exists, and a lockout states how long it lasts. Only someone who has entered the correct password is told that the account is deactivated.
- **Trusted device:** the browser or app installation from which a successful login has already been made is marked for 180 days, and from it the user can keep logging in even if a third party has caused the account to be locked from outside.
- **Operational organisation:** if the organisation has been deactivated or its subscription has expired, the login explains it (see A.5).
- **Real logout:** logging out cancels the session on the server, not just in the browser.
- The **role, permissions and language** used by the interface are always resolved on the server; none of that travels in data the browser could modify.

The same attempt controls apply to password changes (1.10) and to customer portal access.

1.3 Mobile login

The app has its own flow, independent of the web one: the session is stored securely on the device and is designed so that a field technician does not have to sign in again every working day. It does not expire while in use, but **it expires after 60 days without use**, so that a lost phone does not keep synchronising indefinitely.

The mobile session is closed:

- from the app itself, with a logout that also cancels it on the server;
- from the user's web profile (1.11), for example after losing the phone;
- when the password is changed or reset;
- when the account is deactivated or erased (1.5).

The login response already brings resolved everything the app needs to decide which screens to build: permissions, subscribed modules, organisational capabilities, levels of detail, vocabulary and language (chapter 2).

Staff record required. A user without a linked staff record cannot log in to the app, and receives a message asking the administrator to link one. This also applies to administrators: the app always works on the record.

There is also a **lightweight session check** so that the app can verify whether it is still connected without downloading data, which is useful when coverage returns.

1.4 Password recovery

A "I have forgotten my password" flow that sends a link by email, which expires after an hour. Available both for internal users and for portal users, with separate flows.

- The request always returns the same message, whether or not the email exists. If the account is deactivated, an email explaining this is sent; if it is **pending activation** (1.5), a new activation link.

- Requests are **limited per account and per source**, so that nobody's mailbox can be flooded.
- The new password must comply with the common policy (1.10), and a rejected password does not use up the link.
- Resetting the password **closes all sessions** of the account, web and mobile.
- The links sent **stop being valid** when the account is locked manually, when its login email changes and when the password is changed.

1.5 User creation, activation and removal

Creating, modifying and erasing accounts is a task for the organisation's administrator or the superadministrator. An administrator cannot create or manage superadministrator accounts.

When a user is created, an **activation link valid for 24 hours** is generated. The account is not operational until the user opens the link, chooses their password—which must comply with the common policy (1.10)—and **accepts the privacy policy and the legal notice**. The acceptance is recorded (document, version, date, source and browser) and cannot be altered. The same requirement applies to portal users.

- The account can be created from **Users** or from the **Staff** record when giving the person access to the application; in both cases an administrator decides (see 1.9).
- An account not yet activated is distinguished from a deactivated one: the staff record shows it as **"Pending activation"** and offers **"Resend activation"**, which generates a new link and cancels the previous one.
- **Deactivating** an account immediately closes its web and mobile sessions.
- **There is always an administrator left:** a change that would leave the organisation without any active administrator is rejected, asking for another one to be appointed first.
- Each new account counts towards the subscription's user limit (see A.4).
- Only the administrator sees the full list of users; everyone else receives only what is needed to choose a colleague from a drop-down list.

Deleting a user means erasing them (GDPR, art. 17). The account is not physically deleted: it is **anonymised**, so that the work orders, comments and histories in which it appears still add up without anything that identifies the person. It is irreversible, nobody can erase themselves and the erased account no longer counts towards the user limit. The administrator can first download the user's personal data (chapter 25).

1.6 Roles

Roles belong to each organisation: changing a role's permissions in one organisation does not affect any other.

- **Six system roles:** the **superadministrator**, which is the only global one, and five that every organisation has: **administrator**, **supervisor**, **technician**, **warehouse** and **read-only**. They are created automatically when the organisation is created, with their default permissions (see A.3).
- The superadministrator and the administrator grant **everything** within their scope and do not allow their permissions to be edited. The permissions of the other three system roles can be

adjusted by each organisation.

- System roles cannot be deleted.
- Each organisation can create **its own roles**, starting from scratch or **copying the permissions of another role**.
- A custom role with users cannot be deleted: the system says which users have it, so it can be changed for them first.
- A user can only have a role from their own organisation.

1.7 Permission catalogue

Permissions are grouped by functional area. Each one grants a specific action on an area:

AREA	ACTIONS THAT CAN BE GRANTED
Maintenance	View · create · edit · delete · validate, close and reopen work orders
Assets	View · create · edit · delete · configure types, families and styles · approve change requests
Inventory	View · create · edit · delete · record movements · warehouse control (adjustments and priced receipts)
Staff	View · create · edit · delete
Branches	View · create · edit · delete
Customers	View · create · edit · delete
Sectors	View · create · edit · delete
Customer portal	View · manage
Meter readings	View · create · edit · delete · configure measurement points and rules
SLA	View · manage
Reports	View · export · costs · staff tracking
Settings	View · edit
Fleet	View · create · edit · delete
Alerts	Which families of alerts each person receives: incidents, GPS tracking, stock, fleet, assets, security, SLA and synchronisation conflicts

Some relevant separations:

- **Measurement configuration** (measurement point templates and rules) is separated from day-to-day work with readings: a technician who takes readings does not change the configuration that reaches everyone's mobiles.
- **Resolving asset change requests** is separated from editing assets. Nobody resolves their own, except the administrator.

- Within reports, the **financial** side (amounts on the dashboard, on the work order, in reports and the valued inventory) and **people tracking** (fleet, geofence and utilisation per person) are separated from the operational figures.
- **Alert** permissions do not open screens: they decide which general alerts reach each user's bell (chapter 24). The synchronisation conflicts one also gives access to its review screen. Alerts addressed to a specific person do not depend on them.

Some operations require a higher permission than that of their screen: configuring the work **catalogue** (task types, tasks, form templates and incident types) requires editing settings; schedules that delete orders already generated, the delete permission; staff rates, home address, notes and the reason for absences, the staff edit permission.

Each system role comes with a default set of permissions (A.3): a **technician** views and edits **their own** work orders (1.13), submits quick reports and creates incidents from the app, but does not close them and does not see the dashboard or reports; a **supervisor** adds closing orders, asset and measurement configuration, approval of changes, reports with export, costs and tracking, and all alerts; the **warehouse** role manages the full inventory and receives stock alerts; the **read-only** role can only consult.

1.8 Permission resolution in four layers

Every access check goes through four layers, in this order:

1. **Subscribed module.** If the module that governs the permission is not active in the organisation, access is always denied. An organisation without the fleet module will never see vehicles, even if the person's role includes the view fleet permission.
2. **Full-access role.** The superadministrator grants everything; the administrator grants everything within their organisation, once the previous layer has been passed.
3. **Per-user exceptions.** Each user can have explicit exceptions that grant or revoke a specific permission without changing their role. It covers the real-world case of "this technician can close orders" without creating a new role. Only an administrator manages them.
4. **Role permissions.** The permission must be in the role's set of permissions, and the role must belong to the user's organisation.

Two relevant design decisions: an **unknown permission denies everyone**, the superadministrator included, and an **unknown module also denies**. A typing error cannot turn into an unprotected screen or service.

Beneath these layers, which decide **what a person can do**, the database enforces **which organisation the data** seen by each operation belongs to (2.14): even if a query forgot to filter by organisation, it would not return another organisation's data.

Resources shared across the whole platform (the global SLA, the vertical catalogues or the icon library) are modified only by the superadministrator. Some organisation functions are reserved to the administrator: users, audit (1.12), API keys and full data copy, remote connections, webhooks and alert escalation rules.

The interface replicates the same logic to decide what is shown, but actual authorisation is always enforced on the server. The menu and the functions associated with a module or a switched-off capability are omitted, so that no functions are left dangling.

1.9 User ↔ staff link

A user can be linked to a **staff** record. That relationship is what allows a technician with a login to see their own assignments on mobile and on the web (1.13), to have their GPS position recorded and to have hourly cost allocated to them. There are also staff **without a user**: technicians who appear in assignments and in planning but do not access the application.

- The record is a **requirement for logging in to the mobile app** (1.3).
- Giving access to the application from the staff record, or changing from it the role, exceptions, login email or activation of the linked user, **is managing users** and can only be done by an administrator. Those changes are recorded in the audit log.

1.10 Password policy

A single rule for every path through which someone chooses their password: change from the profile, reset by email, account activation, portal registration and portal reset.

- Between **10 and 128 characters**.
- It cannot repeat almost the same character throughout.
- It cannot appear in a known **list of leaked passwords**.
- It cannot be an obvious word (the product name, “password”, a month...) surrounded by numbers or symbols.
- It cannot contain the local part of the user’s own email address.

There are no composition rules (upper case, number, symbol), following the NIST SP 800-63B recommendation: length and the leaked list do the work. The activation, reset and profile screens **tick off each rule while typing**, with the same check that is applied on saving.

Changing the password from the profile requires the current password, has its own attempt limit, closes the mobile app session and keeps open the session of the browser from which it is changed; any other copy of the previous session stops being valid.

1.11 Profile page

Each internal user has their own **Profile** page, accessible from the user menu:

SECTION	CONTENT	WHEN IT APPEARS
Personal details	Name and phone, editable. Login email and username, read-only: changing them is a task for the administrator	Always
Language	Application language (Español or English) or the company’s. Applied to screens, messages, emails and documents (2.15)	Always

SECTION	CONTENT	WHEN IT APPEARS
Your position	Company if external staff and, depending on what the organisation uses, branch, usual shift, teams, usual warehouse and vehicle. Read-only	With staff record
My calendar	Usual shift, absences and schedule exceptions from the last 30 days onwards, flagging those declared by the person themselves that are pending review	With record and working calendar active
Qualifications	Trades, skills and certifications	With record and skills active
Privacy and location	Current location consent, with the option to withdraw it; consent history; latest tracking sessions; and download of one's own location data by date range in CSV or JSON	With record and real-time field module
Security	Last login, password change (1.10) and the mobile app session — device and last use— with the option to close it	Always

1.12 Audit log

Sensitive actions are recorded, so that who did what can be reconstructed in the event of an incident:

- **Access:** successful, failed and locked logins on the web, the app and the portal.
- **Session and credentials:** logout, password change, request and reset, account activation and closing of the mobile session.
- **Accounts and roles:** creation, changes (with the before and after of role, status and exceptions), erasures, unlocks and activation resends for internal and portal users; creation and deletion of roles and permissions granted or revoked.
- **Settings and structure:** changes to the general settings, profile applied, company working hours, branches and customer life cycle.
- **Personal data:** exports and erasures (chapter 25).
- **Integrations:** API keys, full data copies, remote connections and webhooks.
- **Operations:** decisions on asset change requests, billability and customer acceptance, corrections of actual dates, alerts attended and escalated, synchronisation conflicts resolved and geofence reviews.

Guarantees: logging never causes the operation it records to fail, it is kept even if the operation is rolled back, it stores no secrets and it cannot be modified or deleted from the application. The default retention is one year and never less than 90 days.

Consultation screen. The organisation's administrator opens it in *Administration* → *Audit*: paginated list with filters by type, action, person, outcome and dates; detail of each entry with the changes shown as *before* → *after*; and the history of a specific object. Always limited to their organisation. When the person acting is platform support, the organisation sees it as "Platform support".

1.13 Technician's work order scope

Permissions decide which actions someone can perform; **scope** decides on which work orders. On the web the same criterion is applied as in the mobile app:

- Whoever **edits orders without being able to close them** (the case of the Technician role) only sees and modifies the work orders **in which their staff record appears**: list, board, record, daily pulse, map, timeline, search and multiple selection. Someone else's order responds as if it did not exist.
- Whoever **only consults** (Read-only, Warehouse) sees all of them.
- The administrator and whoever can close orders see all of them.

In an **asset's history**, the technician sees all their own interventions, but those of others appear as "By another technician", without description and without being able to open them.

For the same reason, the default Technician role does not see the dashboard or reports, which give company-wide figures. The administrator can give that permission back.

02

Multi-tenant, verticals and subscription

AVAILABILITY	Core
SURFACES	Web (global administration panel and organisation settings)
WHAT IT IS FOR	It allows a single installation to be operated for many different organisations, with completely isolated data and with preconfigured sector catalogues that reduce onboarding a new customer from weeks to minutes.

2.1 The organisation

The root unit of isolation. All entities in the system hang from an organisation, and the database enforces that isolation beneath the application (2.14). Each one stores:

- Tax and identity data: trading name, registered name, CIF (Spanish tax ID), full address, logo.
- **Status** (active or deactivated) and subscription **validity date**, which are checked at login (A.5 and 2.16).
- **Subscribed modules**: maintenance, customer management (which includes the portal), SLA, meter readings and the fleet, real-time field and planning modules (2.2).
- **Subscription limits**: users, customers and portal users per customer (A.4).
- **Work settings**: date and time formats, default language and currency, planning horizon, organisational capabilities, levels of detail and profile, and the geofence, SLA and dashboard target settings.
- Associated **vertical**.

What constitutes the commercial relationship —modules, limits, validity and vertical— is managed by the provider.

2.2 Catalogue of subscribable modules

Three modules, deliberately few:

MODULE	WHAT IT ENABLES
Fleet and costs	Vehicles, fuel prices, technician-vehicle assignments, route legs with their cost, odometer and alerts for assignments pending review
Real-time field	GPS positions, tracking sessions, technician's route, work order geofence (chapter 15) and auto-pause alerts

MODULE	WHAT IT ENABLES
Planning	Multi-day work order optimiser and planning of a specific order from its creation (suitable teams and technicians, sorted by real proximity)

In addition, seven areas are **always active** as the core of the system: assets, staff, branches, reports, settings, sectors and alerts. An area being core does not prevent the organisation itself from hiding it with its organisational capability (2.3), as happens with branches and sectors.

Customer management is subscribed as a module of its own, which includes the customer portal, plus the number of customers. Without it, the organisation has a single customer, its own (3.2).

Withdrawing it from an organisation that had customers deletes nothing: it suspends them (3.2).

2.3 Organisational capabilities

Unlike modules, **they are not subscribed**: the organisation's administrator declares them in the general settings, because they describe **how that company works**, not what it has subscribed to. They are free by product decision: nobody should pay to simplify, nor hit a commercial wall when growing.

CAPABILITY	WHAT DISAPPEARS WHEN IT IS SWITCHED OFF
Branches	Branches, automatic branch resolution, allocation by site
Teams	Work teams, team shift, sectors assigned to a team
Sectors	Sectors, zoning types, GIS import of sectors
Warehouse	Warehouses, items, movements, weighted average price, suppliers, stock alerts and materials on the work order
Skills	Trades, skills, certifications and the "Training / Courses" module, requirements per task and expiry control
Working calendar	Shifts, public holidays, absences and one-off calendar exceptions
Incidents	Incidents as an entity distinct from the work order, and their catalogue of types
Forms	Dynamic forms, their design and their assignments
Preventive	Recurring scheduling of work orders
Costs	Staff hourly cost, work order cost, inventory value

Whatever has not been declared is considered active, so as not to alter the behaviour of organisations that have never touched their configuration: a read failure never cuts functionality, it degrades towards the complete option.

A switched-off capability is applied in three layers: the server closes all its functions, the menu and buttons are hidden, and the interface does not even load its screens.

2.4 Levels of detail

Capabilities answer “do you use this?”. Levels of detail answer something different: **how much rigour you want within a process you do use**. Nobody gives up work orders, but few companies need to distinguish “the technician has finished” from “the supervisor has reviewed it” from “it is closed and billable”.

LEVEL	VALUES	WHAT CHANGES
Work order cycle	Simple · Supervised · Complete	The order’s life cycle (5.1)
Asset record	Basic · With type · Complete	Depth of the record: type and geometry; with properties; with composition hierarchy and references
Criticality	Priority only · Complete	Hides the criticality scale, which only makes sense with SLA
Dashboard	Basic · Complete	Seven operational figures versus the full dashboard (20.1). The cut is made on the server
Mandatory creation fields	List	What creating a work order requires (5.3)

The default value of all of them is **the most complete**, for the same reason as with capabilities.

2.5 Starter profiles

Onboarding a company cannot consist of deciding fifteen things before having seen the product. A **profile** writes the organisational capabilities and the levels of detail in one go.

PROFILE	RECOGNISED IN	WHAT IT BRINGS
Essential	“I have assets and work to do on them”	Assets on the map, work orders with a three-status cycle, incidents, field forms, preventive maintenance, field app and basic dashboard. No warehouse, skills, shifts, teams, branches, sectors or costs. Creating an order does not require any optional field
Professional	“I have a team, a warehouse and cost control”	Adds inventory, working calendar and costs, with a supervised cycle
Advanced	“I have several sites, teams and certified staff”	Adds branches, teams, sectors and skills: all ten capabilities and every level at maximum
Custom	—	The state the organisation is left in when any individual switch is moved

Three deliberate properties:

- **The profile is a starting point, not a lock.** Any capability can be moved afterwards.
- **Profiles are incremental:** Essential $\subseteq$ Professional $\subseteq$ Advanced. Moving up a profile never takes anything away.

- **No profile requires more than necessary.** Profiles only relax data discipline; what is mandatory for everyone —the maintenance type— cannot be relaxed by any profile.

Switching off a capability **deletes nothing**. The recorded stock stays where it was and reappears when it is switched back on; the interface warns of this before confirming.

2.6 Verticals

A **vertical** is a business sector (water and sanitation, lighting, facilities...) with a multilingual name and description. A complete catalogue is maintained for each vertical:

- Asset families (hierarchical).
- Asset types, with their composition hierarchy, whether they are measurable and their **measurement points**.
- Asset type properties.
- Map styles.
- Maintenance types, with **mandatory signature** and **estimated duration**, and their tasks.
- Form templates, with sections (also **repeatable**), fields and assignments.

Starter verticals. The platform comes with seven complete verticals: Water and sanitation, Public lighting, Wind farm, Photovoltaic plant, Facilities and business parks, Smart irrigation and Road signage.

A vertical in use is not deleted: it is **deactivated** and is no longer offered at onboarding.

2.7 Cloning a vertical into an organisation

When a vertical is assigned to an organisation, that entire structure is **copied** into the organisation's own catalogue, respecting the order of dependencies. From that moment the organisation owns its copy and can modify it freely without affecting the global catalogue or other organisations.

This is the mechanism that allows a new organisation to be created with its complete asset typology, its inspection forms and its symbology already loaded.

2.8 General organisation settings

A dedicated screen where the administrator defines:

- **Working days and hours:** a **weekly** schedule with up to two time slots per day, which can be overnight. This is the company's schedule; each branch can inherit it or have its own (3.5). It is also the working day of anyone without an assigned shift (chapter 9), and saving it recalculates SLA deadlines.
- **Date and time format, default language and currency** (2.15). The chosen format is applied to the web, PDFs, scheduled reports, emails and the spreadsheets in data copies.
- **Planning horizon:** days to plan and how many days in advance orders are created.
- **How your company works** — the section that governs complexity, for the administrator only, with three blocks: **activation suggestions** (2.10), at the very top and only if there is something to suggest;

starter profile (2.5), as cards with the phrase in which the company recognises itself; and “**Fine-tune**”, collapsed, with the ten switches and the consequence of switching off each one.

- **Work order geofence**, with the real-time field module (chapter 15).
- **Dashboard targets**: the thresholds that colour each indicator (chapter 20).
- **Subscribed modules**, read-only.
- **SLA**, if the organisation has it: deadlines, mapping to priority, “at risk” threshold and targets (chapter 19).
- **Reading type**, with the meters module.

The time zone is not configured here: it is determined by the main branch (2.13). Changes are recorded in the audit log.

2.9 Onboarding an organisation in four questions

Whoever onboards a customer does not yet know whether they need “organisational capabilities”: they know how many people they have and what they maintain. There is a guided onboarding in four questions:

#	QUESTION	WHAT IT DETERMINES
1	How is your work organised? (<i>several customers / own asset base</i>)	Whether work is by customers or by sectors
2	How many people work in the field? (<i>up to 5 / 6-25 / more than 25</i>)	The starter profile: Essential, Professional or Advanced
3	Do you manage a materials warehouse?	The warehouse, which can contradict the profile
4	What do you maintain?	The vertical

A preview returns, in business language, what is going to be configured, readable aloud to the customer before confirming. The warehouse answer overrides the profile: someone with three people may have a depot, and someone with thirty may not. A missing answer falls back to the most conservative value.

2.10 Progressive activation suggestions

What is switched off is presented not as “you don’t have it” but as “you don’t need it yet”, and the system proposes it **when the organisation’s own data already justifies the need**:

SIGNAL DETECTED IN ITS DATA	SUGGESTION
Orders whose description mentions quantities of material	Switch on warehouse
Same type of work repeated on the same asset	Switch on preventive
Staff spread across several work centres	Switch on branches
Certificates with an expiry date already entered	Switch on skills

Registration numbers noted in orders

Find out about the **fleet** module

Only what is switched off is evaluated, each suggestion cites its evidence and its count (“in 14 of your orders you have noted materials by hand”), the free ones and the fleet one go in separate lists, and only the last 180 days are looked at.

2.11 Product vocabulary

“Work order” is the correct term in the sector, and it is exactly what makes a four-person company feel that the software does not speak its language.

In the **Essential** profile the visible labels are replaced: *Work order* → **Job**, *Incident* → **Report**, *Maintenance type* → **Type of work**, *Staff* → **Team**, *Responsible* → **Who does it**, *Planned date* → **By when**. In the other profiles the sector vocabulary is kept.

Only what is written on screen changes: the API and the exported data always keep the same names. Today it is applied in the operations panel tabs, the top bar, the administration menu and the work order report; the rest of the screens keep the sector vocabulary.

2.12 What is created when an organisation is onboarded

Onboarding always creates the minimum needed to work from the first minute:

- Its **five system roles** with their default permissions (1.6).
- Its **default weekly working schedule** (2.8).
- A **main branch**, even if the company does not work with branches: it is the implicit site and the one that sets its time zone (2.13).
- Its **own customer** (3.2).
- The **system catalogue** used by the app’s quick reports.
- The **vertical’s catalogue** (2.7), if one is indicated.

2.13 Organisation time zone

Instants are stored in universal time and read in local time with each entity’s zone. The organisation’s zone is that of its **main branch**; with one branch in Madrid and another in the Canary Islands, the organisation is read in Madrid time if the main one is Madrid.

- Each **work order** is read with the zone of its branch; if it has none, that of its first technician, and failing that, that of the organisation.
- Each **person** can have their own zone; if not, they inherit that of their branch. This way, the absences and capacity of a person in Madrid are not calculated in Canary Islands time.
- Each **web user** writes and reads times in the zone of their branch.

That same rule is used by working days, the day’s due dates, SLA deadlines and reports.

2.14 Isolation by organisation in the database

In addition to the organisation filters the application applies in every query, the database enforces isolation on its own, with **row-level security** in strict mode:

- Each operation declares which organisation it belongs to, and every table with an organisation's data only allows that organisation's rows to be seen and written.
- **Without a declared organisation, nothing is visible.** Scheduled processes work organisation by organisation.
- The superadministrator, when acting on a specific organisation, works inside it.

The consequence is that an oversight in an application query cannot return another organisation's data: the database hides it anyway. There are automated tests that fail if a new table is left without this protection.

2.15 Languages and currency

The application is translated into **Spanish and English**.

- **Each person's language.** The company sets a default language; each user can choose another in their profile. The login screen and the portal have their own selector, and the mobile app uses the device language.
- **What is translated.** The web and portal screens, error messages, bell alerts (which are read in the language of whoever is viewing them), catalogue names, PDFs, scheduled reports and **emails**, which go out in the recipient's language.
- **Formats.** The company's date and time format and the decimal separators of the user's language are applied throughout the web.
- **Currency.** Each organisation chooses the currency in which its amounts are shown: euro, pound, US and Canadian dollar, Swiss franc and the main Latin American currencies. It is presentation only: no amount is converted.

Some automatic history or stock notes are stored in Spanish whatever the company's language.

2.16 Organisation life cycle: deactivation and deletion

An organisation goes through three states: **active** → **deactivated** → **deleted**, and each step is a separate action by the provider.

Deactivating asks for a reason and notifies the organisation's administrators by email. From that moment nobody gets in —web, app, portal or API— and scheduled processes skip it. **Nothing is deleted:** reactivating it returns the organisation exactly as it was.

Deleting erases every trace of the organisation: data, uploaded files and caches. It is only possible for organisations deactivated at least 7 days earlier, with express confirmation, and a **full copy** with its files can be downloaded beforehand. Deletion is a phased process that resumes if interrupted, ends by verifying that no data from the organisation remains and issues a **certificate** in PDF.

2.17 Help and user manuals

The user menu has a “**Help**” entry that opens the **user manuals**: twenty-one, one per area of the application (getting started, map, assets, layers, work orders, preventive, incidents, planning, staff, working calendar, branches and sectors, customers and SLA, customer portal, inventory, fleet, reading batches, forms, costs, dashboard, administration and ERP integration), with screenshots, an index with search and a complete English version.

The manuals are **public**: they can be read without logging in. Some administration screens also carry collapsible help, “**What is this screen for?**”, with a link to the manual section.

03

Organisational structure

AVAILABILITY

Core; customer management (with its portal) is a module, and branches, teams and sectors depend on the corresponding organisational capability

SURFACES

Web · Mobile (consultation) · Customer portal (implicit)

WHAT IT IS FOR

It models two very different organisational realities —the company that provides services to several customers and the operator that manages its own asset base— without forcing either of them to use concepts that do not suit it.

3.1 The two models: customers or sectors

The system supports two alternative territorial organisations:

- **With customers** (service company, contractor, maintenance contractor): work is organised by customer, and each customer has its zones.
- **Without customers** (operator managing its own network): work is organised directly by geographic sectors.

The selector at the top of the map changes from “Select customer” to “Select sector” depending on the model.

The chosen customer scopes the whole map. In the customer model, the drop-down does not just frame the map: it filters everything shown —assets, orders, incidents, readings, planning, heat map and the day’s counters— until “All” is chosen.

Everything belongs to a customer (3.2), also in the model without customers: there, only one exists, the organisation’s own.

3.2 Customers

Full record: name, contact person, email, phone, full postal address, its own logo.

- The customer’s **geographic area**, which can be drawn on the map, and a calculated **bounding box** of all its installations to centre the map when it is selected, recalculated automatically.
- **Own SLA that can be enabled** per criticality and type (chapter 19).
- Its own **signature policy** (5.9), **customer acceptance at closure** and **justified SLA pauses**, which can be enabled per customer (chapters 19 and 22).
- Vertical and responsible branch.

Own customer. Each organisation has exactly one **internal** customer, “Own”, which gathers the work on its own asset base. It cannot be deactivated or deleted, and it does not count towards the customer limit.

Mandatory customer. Assets, work orders, incidents, readings and the rest of the work always carry a customer, which the server resolves on every entry path (web, app, synchronisation, uploads):

- Without the customers module, always the organisation’s own.
- With the module, the one indicated or the one inherited—from the parent order, the incident or the assets—; if there is nowhere to take it from, it is requested.
- The assets must belong to the same customer as the work, and a deactivated or suspended customer does not accept new work.
- The only exception is **unclassified incidents**: they can come in without a customer, but they must be classified before being converted into an order or marked as resolved.

Bulk customer assignment. Anything left without a customer can be distributed from the web.

Life cycle: active → deactivated → deleted.

- **Deactivating** asks for a reason and first shows a summary of what is affected: pending and in-progress orders, open incidents, preventive, reading schedules and routes, portal users and assets. For each block a decision is made, with a default value: pending orders are cancelled, those in progress are left to finish, open incidents are closed, preventive and scheduled readings are switched off. The portal is always closed and its assets leave the mobile app (they are kept on the web).
- **Reactivating** returns the customer to active and returns the list of what needs reviewing; each schedule is reactivated separately. A customer’s assets can be transferred to another.
- **Delete completely** is only offered for deactivated customers and requires typing their exact name. A full copy can be downloaded beforehand. It is a phased process that resumes if interrupted and ends with a downloadable **certificate**.

Customer limit. The number of customers is a subscription limit (A.4). A deactivated customer **still occupies its slot**; only deletion frees it.

Customers module suspended. If the module stops being subscribed, nothing is deleted: customers are **suspended**, the portal is closed and the web shows the customers read-only until it returns.

3.3 Sectors and zoning types

A **sector** is a geographic work area with its own geometry, code, name, colour and level. Sectors are **hierarchical**, without cycles.

- They are classified using **zoning types** (service zones, inspection zones, districts...).
- They can be **imported from GIS files**, with a prior analysis step that shows the fields before confirming.
- A sector is assigned to a branch and can have **teams** assigned, one of them the main team.
- Free-form metadata in key-value format.

Sector of each asset. The system assigns each asset the sector that contains it. When a linear or areal asset crosses several, the one that takes the greater part of its length or area wins, with a margin so that small edits do not make it jump sector. Work orders also store the distribution across the **affected sectors** (5.22).

Geometry changes and deactivations. Drawing or moving a sector only reassigns **open** orders; finished ones keep the sector with which they were done. A sector with orders or reading schedules is not deleted: it is deactivated.

3.4 Customer zones

Geographic subdivisions within a customer, with their own zoning type and geometry editable on the map. They are used to filter assets and work by area within the same customer, and to **limit what each portal user sees** (chapter 22).

- Hierarchical and without cycles.
- **Import from file** (GeoJSON, zipped Shapefile or GeoPackage), with the same prior analysis as sectors.
- A zone with subzones or with portal users is not deleted: it is deactivated.

3.5 Branches

Internal organisational unit: regional branch, office, work centre. It is the piece that connects the administrative structure with geography and with accounting.

- Name, address, city, region code.
- One branch marked as **main**, created automatically when the organisation is onboarded.
- **Location on the map**, with **its own layer** in the layers panel: they are the bases from which technicians and teams set out, and seeing them explains the journeys the planner proposes.
- **Base radius or polygon**, to know when a technician is at their branch.
- **Time zone**, selectable in the form. That of the main branch is that of the whole organisation (2.13), and changing it recalculates open SLA deadlines.
- **Working hours**: inherits the company's or has **its own**.
- Its own **public holidays** (chapter 9).
- **Default warehouse**.
- Active/inactive.

The branch propagates to work orders, incidents, sectors, customers, staff, vehicles, warehouses, reading batches and alerts, and is the axis of cost accounting allocation and of the division of responsibility. A general alert affecting a branch is seen only by the users of that branch, those without a branch and administrators (chapter 24).

Automatic branch resolution. When an order, an incident or a reading batch is created, the system determines the responsible branch through an explicit cascade: assets → sector → most frequent branch; otherwise, the work's own location → the sector that contains it; otherwise, customer; otherwise, the assigned staff; otherwise, the record of whoever creates it; and as a last resort the main branch.

Protected deletion. A branch cannot be deleted while anything references it, even if inactive or closed. The rejection states what and how many, and the web offers to deactivate it.

3.6 Work teams

Groups of staff, with:

- **Members**, each with their role within the team.
- The team's **default shift**, which the planner uses as the time window for the whole unit.
- **Branch** it belongs to. All members belong to the same branch as the team.
- **Assigned sectors**, with a main one.

For the planner, a team is a **single route unit**: work is assigned to it as if it were a single resource, and its daily capabilities are the union of those of the members present that day.

3.7 Warehouses

Logistics nodes that belong to a branch. Three types:

- **Physical** — the branch's depot or warehouse, with a mandatory location.
- **Mobile** — the van. It is linked to a **vehicle** and/or a **person**; a vehicle and a technician each have at most one active van.
- **Virtual** — for movements without a physical location.

Each with its own stock per item (chapter 16). In the app, "**My van**" resolves each technician's warehouse: the van of their vehicle, their own, the one on their record, that of their branch or the main one.

A warehouse can only be deactivated without stock, and it is not deleted if it has history.

04

Asset management (GIS)

AVAILABILITY	Core, always active
SURFACES	Web · Mobile · Customer portal (consultation) · Public API
WHAT IT IS FOR	It is the inventory of the asset base to be maintained. Everything else — orders, incidents, readings, routes— hangs from here. An asset can be a streetlight, a valve, a network section, a building, a green area or an inverter.

4.1 Families and asset types

- **Families:** first-level grouping, **hierarchical**, with a display order that can be reordered in bulk and a tree view. A family in use is not deleted: it is deactivated.
- **Asset types** (the map layers): the operational classification. Each type defines:
- **Multilingual** name and description.
- **Layer use:** **operational** (generates assets that are worked on) or **cosmetic** (display only).
- Expected **geometry type** (point, line, polygon, their multi-part variants or generic).
- **Identifying fields:** which properties identify the asset. Their values form the **asset code**, which is calculated on creation, edits, uploads and synchronisations and can be regenerated in bulk. They are also the key with which a reimport recognises existing assets (4.4).
- **Display order**, default visibility and whether the layer is cached.
- Whether it is **measurable** (accepts meter readings), whether the reading **requires a photograph** and which trade is assigned by default to its readings.
- **Hierarchy rules** (4.2).
- Whether it is **active**. A layer with history is not deleted: it is **deactivated**, and a deactivated layer does not accept new assets by any route.

Families and types exist on two levels: **global** (within a vertical) and **organisation-specific** (the cloned, editable copy).

4.2 Asset hierarchy

Rules per layer. The **Hierarchy** block of each layer indicates which layers its assets can hang from (an electrical panel contains switches, a pumping station contains pumps). Each rule accepts a **maximum quantity** of children per parent and whether the parent is **mandatory**.

The parent of each asset. When creating or editing an asset, the “**Hangs from**” field searches for candidates by code or by any field, sorts them by proximity, allows **choosing it on the map** and proposes the nearest. The server validates the parent on every route: it must exist and be live, belong to a permitted parent layer, belong to the same customer, not create cycles and not exceed the maximum quantity.

Record. The **Hierarchy** section shows the parent and the children grouped by layer, and allows **adding a child**.

Bulk assignment. To attach an already loaded layer in one go, the **Assign parents...** window offers three methods, with a **preview** before applying:

METHOD	HOW IT FINDS THE PARENT
By field	A field of the child contains the parent’s code; useful for converting reference fields into hierarchy
By proximity	The nearest parent, within a maximum distance
By containment	The parent polygon that contains the child

It respects the maximum number of children, can be limited to assets without a parent and to the same customer, and is designed for large volumes. In a layer’s list it is also possible to **change the parent** of a selection of up to 1,000 assets, and file uploads accept a column with the parent’s code.

What it does not yet do. A work order on a parent asset does not include its children, and the mobile app has no screen for choosing the parent.

4.3 Custom properties

Each asset type defines its technical record: a set of fields with a **multilingual label**, type, default value, whether it is mandatory, **validation rules**, width in the form, visibility (visible, hidden, read-only) and a reorderable order.

Field types: text and text area, number, integer, decimal, **auto-number**, date, date and time, time, single selection, multiple selection and radio buttons, checkbox, yes/no, email, URL, phone, colour and **file**. The auto-number is numbered by the server without duplicates, also when it is added to a layer that already has assets. The file field links a document from the asset’s documentation.

In addition:

- **Key field:** which properties identify the asset.
- **Cascading options:** the options of one field can depend on what was chosen in another (for example, model dependent on manufacturer).
- **Fields referencing another asset type:** the property takes its values from the assets of another layer.
- **Customer-specific properties,** in addition to those common to the type.
- **Hidden and read-only fields:** the server keeps their saved value even if the form sends something else.

- **Renaming a field** carries the new name over to the values, the styles, the documentation and the references from other layers; a field used by measurement rules, documents or references cannot be deleted.

4.4 Asset creation

- **Manual**, drawing the geometry directly on the map, with customer or sector, branch, status, parent and properties. A creation that fails keeps the drawing.
- **Bulk import from GIS files**: GeoPackage, zipped Shapefile, GeoJSON and DXF, up to 150 MB per file, in the background with **progress tracking** and **automatic creation of properties** from the file's fields when the layer does not yet have any. The source reference system can be specified and is mandatory for DXF.
- **Bulk import from CSV**, also in the background. Each column becomes a property; **assets are created without geometry**, which is drawn afterwards or brought in with a GIS import.
- **Synchronisation from external REST APIs** (21.10 and 25.3): a layer can be kept up to date from a remote service, with **incremental** synchronisation at the times chosen. Requests only go out to public Internet addresses.

Add or replace. An upload **adds** assets or **replaces the data** of the layer. Replacement validates the whole file before touching anything, **deactivates** (does not delete) the previous assets, and is rejected in full if any live asset has work, readings, photos, documents or other data depending on it.

What the import does with problematic rows.

SITUATION	GEOJSON AND DXF	GEOPACKAGE AND SHAPEFILE
Geometry of a different type from the layer	Nothing is imported; the result says which rows	Same
Invalid geometry	It is repaired	It is repaired; if it ends up empty, it is skipped
Row without geometry	Discarded with its reason	Skipped and counted
Row missing a mandatory property	Not imported, and it says which one is missing	Same
Reimporting the same file	Rows whose key already exists update the asset and only the new ones are created	Creates new assets: duplicates

4.5 Asset statuses

Catalogue of statuses with multilingual name and description, colour and icon. The status is changed individually from web and mobile, is reflected on the map and can be seen across all layers with the thematic by status (4.11).

Permanent decommissioning. There is **only one way to do it**, the same for deleting the asset, changing its status, approving a decommissioning request and rejecting a creation request. A decommissioned asset is not drawn on the map or sent down to the app, cannot be a parent and does not return to another status. It is never physically deleted from the application.

4.6 Photographs and documentation

- Multiple **photographs** per asset, automatically optimised with a **thumbnail**.
- **Documentation** attached per asset, and also **per asset type** or per value of a property, which makes it possible, for example, to attach the technical data sheet common to all equipment of a given model.
- **Control of what is uploaded:** each file is checked internally according to its format, not by its extension. Images must open as images; PDFs, without executable actions; modern Office documents, without macros or embedded objects; old formats and those that allow macros are rejected. Limit of **20 MB per photo** and **50 MB per document**.
- Downloading requires a session, and each organisation can only access its own files.

4.7 Spatial queries and lists

- **Selection by polygon** drawn on the map, and **selection by point with radius**.
- **Layer list:** paginated, with **Code**, **Status** and **Hangs from** columns, sorting by any column or property, search, scoped by customer and with selection to **change parent**. **CSV export** with the same filters, up to 50,000 rows, protected against formulas and recorded in the audit log.
- **Bounding box of the whole asset base**, stored per organisation to load the map instantly.
- **Locating an asset by its code**, for integration or quick search.

4.8 Asset geofence

The technician's arrival at the work location is detected with a **fixed radius of 200 m** around the order's assets, the same for the whole application. Departure is counted, with the default settings, at 300 m. The details are in chapter 15.

4.9 History and measurement points

Measurable assets accumulate a reading history that can be queried with statistics, and have **measurement points** with validity thresholds and frequency (chapter 17).

4.10 Asset work history

The asset record shows its code and its sector and, in its tabs:

- **Work orders:** all those that have had it as their object, with status and dates. A technician also sees those of others, but summarised as "By another technician" (1.13).
- **Incidents:** open and closed ones from the **last 12 months**, including those that include it among several assets.

4.11 Map styles

Each asset type has one or more styles. The editor covers the full cartographic vocabulary:

- **Points:** colour, opacity, radius, outline.
- **Lines:** colour, opacity, width, join and cap, dash pattern, offset, blur.

- **Polygons:** fill, opacity, outline, pattern.
- **Labels:** field, size, colour, halo, anchor, offset.
- **Icons:** image, size, opacity and **icon colour**.
- Layer **minimum and maximum zoom** and **geometry type filter**, so that the same type can render its points and its lines differently.

Styles can depend on the **values of a property** (thematic styling: colouring by material, by year of installation), and the system lists the available values to build the classification. From this, each layer's **legend** is derived automatically.

Thematic by status. All asset layers on the web map offer, without any configuration, the “**Asset status**” thematic, with its own legend.

Changing a style is applied on open maps without reloading. As with families and types, there are global styles per vertical and organisation-specific styles.

4.12 Icons

System icon library, organised by themes, with a default colour and preview. It brings together some three hundred freely licensed icons in ten themes aligned with the verticals, presented with an almost square proportion on the map. The colour is set by the style. They are used in layer styles, statuses, criticality levels and catalogues.

4.13 Copy of the assets for mobile

In addition to the mapping for drawing the map (21.7), each organisation has a **copy of its assets for mobile** (type, customer, sector, status, parent, code, properties and geometry) that the app downloads on its first launch.

- It is generated **every night** and, **on demand**, from the app.
- It is built separately and only replaces the published one when it is complete and verified: the app never downloads a half-finished copy.
- **Assets that stop being needed in the field.** The assets of a deactivated customer and deleted assets do not go to mobile, and synchronisation tells the app which ones to remove. The web keeps showing them with the deactivated customers filter.

4.14 Asset position history

Moving an asset does not overwrite its position without a trace: each geometry change leaves a record of the **previous and new position**, when, who and **from where**:

ORIGIN	WHEN
Web	Editing in the record or in the map's edit mode
Mobile	App synchronisation with geometry
Request	Approval of a change request (the author is whoever requested it)

ORIGIN	WHEN
REST synchronisation	Synchronisation from an external connection
Other	Any other route

The record is made by the database itself, so that no route bypasses it. The record has the collapsible section “**Location history**”, with the distance between positions and buttons to see each one on the map. The history is also offered in the public API and in the full data copy (25.17).

05

Work orders

AVAILABILITY	Maintenance module
SURFACES	Web · Mobile · Customer portal (viewing and acceptance) · Public API (indirect)
WHAT IT IS FOR	The work order is the unit of execution and of cost. It brings together what needs doing, where (on which assets or at which point or area), who does it, when, with which materials and with what documented outcome.

5.1 Life cycle: seven statuses

STATUS	MEANING
Pending	Created and not started, whether or not staff are assigned
In progress	Work under way
Paused	Interrupted (manually or by end of working day)
Completed	Finished by the technician
Validated	Reviewed by the supervisor
Closed	The only successful closure; billable
Cancelled	Unsuccessful closure, before completion

Transitions permitted in the full cycle, enforced by the server:

FROM	CAN MOVE TO
Pending	In progress · Cancelled
In progress	Paused · Completed · Pending ("Return to pending") · Cancelled
Paused	In progress · Cancelled
Completed	Validated · In progress (can be reopened)
Validated	Closed · Completed (can be sent back)
Closed	Terminal
Cancelled	Terminal

A transition that is not in the table is rejected. Two simultaneous changes to the same order cannot take a cancelled order to completed.

Permissions for each status change: a single rule for the web and the app.

CHANGE	PERMISSION
Start, pause, return to pending, complete	Edit orders
Validate or close	Close orders
Reopen a completed, validated or closed order	Close orders
Cancel	Delete orders, with a mandatory reason

It is applied by the form, the “...” menu, the Kanban board, bulk actions and mobile synchronisation. An organisation that wants the technician to close orders gives the close permission to that role.

Every order is created pending, in every cycle and for any user. Orders generated by a schedule (chapter 6) are too. The only exception is the app’s **quick report**, which is work already done and is created completed.

Assigning is not a status. An order is **assigned** when it has at least one person, whatever its status: that is how the technician sees it in their app. Assigning, changing or removing staff does not change the status and is recorded in the order’s timeline (5.16). An order is **planned** when it has operational staff and a window (planned start and end); that is what the planner (chapter 11) looks at so as not to propose it again.

The separation between completed, validated and closed is deliberate: it distinguishes “the technician has finished” from “the supervisor has reviewed it” from “it is closed and billable”. Every status change is recorded in a **history** with timestamp, author and notes —and, for pauses, their cause (5.20)—, which feeds the dashboard indicators, the work segments (5.18) and the timesheet (5.19).

What can be changed in each status. Closed and cancelled are terminal. A completed or validated order only allows a status change; to edit it, it has to be reopened.

What blocks closure. Completing, validating or closing requires, in addition to the permission:

- the customer’s **signature** or a signature exception when the policy requires one (5.9);
- the **mandatory forms** completed (5.7);
- for closing, the **customer’s acceptance** when the customer requires it (5.21 and chapter 22).

If the closure comes from the app and any of these checks fails, the order returns to its previous status with its previous actual dates.

Reduced cycles. A small company has no supervisor to validate, and forcing it through three statuses to close a job makes the product feel as if it belonged to another league. The cycle detail level (2.4) chooses between three flows, which differ only in **how many steps it takes to close**:

LEVEL	STATUSES OFFERED	HOW IT IS CLOSED
Simple	Pending · In progress · Paused · Closed · Cancelled	In progress → Closed, in one step

LEVEL	STATUSES OFFERED	HOW IT IS CLOSED
Supervised	Adds Completed	Completed → Closed, without validation
Full	All seven	Completed → Validated → Closed

When statuses are removed the cycle is reconnected, so that at every level there is a path to closure; pause exists in all of them, because the end-of-day auto-pause (14.6) needs it. The “...” menu and the Kanban board only offer the changes valid for the organisation’s cycle and the user’s permissions. The history does not change, so analytics keeps reading the same thing.

5.2 Priorities

Four levels: **Low, Medium, High, Urgent**. Independent of **criticality**, which is derived from the priority for the SLA. An urgent order alerts its persons in charge as a critical order.

5.3 Creating a work order

Full form in the side panel, over the map, in six blocks (main information, location, work type, planning, team and resources). It is opened from the orders panel, by right-clicking on the map (“New order here” or “New order for this asset”), from an incident or from a slot in the calendar.

- Title, description, text address.
- **Customer**: mandatory with the customers module; it is inherited from the parent order, the incident or the assets, and all the assets must belong to the same customer (3.2).
- **Maintenance type** —always mandatory— and specific **tasks** within that type. Inactive types and tasks are not offered.
- **Where**: assets selected on the map (by clicking or by drawing a polygon) or a **location of its own** —a point or an area— for work that does not fall on an inventoried asset (5.14).
- Sector and **branch**, resolved automatically (3.5 and 5.22).
- Priority, **planned start**, **planned end** and an optional **deadline** (5.25).
- **Estimated duration**, calculated automatically (chapter 12) and overridable.
- **Assignment** to staff, with one or more persons in charge, and help for adding an entire team (5.5).
- **Associated forms** proposed by type, tasks and assets; mandatory ones cannot be unticked (5.7).
- **Resources**: what the tasks call for —material, tools and PPE— and **planned materials** with quantity and warehouse (5.6 and 5.23).
- **Documentation** attached from creation.
- Free-form properties.

Creation returns, without preventing the order from being created, two kinds of warning: **stock** (the planned quantity exceeds what is available or leaves the material below minimum) and **qualification** (none of the assigned staff covers a current mandatory skill or a required trade).

There is also creation **from an incident** (7.6), which inherits its customer, its assets and its location. An incident can only have **one live order** resolving it, unless the previous one ended without solving the problem (5.24). And there is the **quick report** from the mobile app (23.8).

What is mandatory. The **title** and the **maintenance type** are mandatory for everyone: without a type the estimated duration cannot be resolved, forms and tasks cannot be attached, and the dashboard cannot separate planned from corrective work. The rest is decided by the organisation among four fields: **assigned person in charge**, **planned date**, **at least one asset** and **location**. By default only the person in charge is required; in the Essential profile, none. The error speaks the language of the form: “The work order needs an assigned person in charge”.

Editing. Editing keeps what remains on the order (assignment date, status and notes of each task, **form responses**) and only removes what has been unticked; a form with responses cannot be unticked. If another person or another tab has saved the order in the meantime, the save is rejected so as not to overwrite their changes.

5.4 Planning assistants at creation

Three aids built into the form, all based on real data:

- “**Plan this WO**” — finds the first window in which there is an **available and qualified unit** —team or technician— and ranks them by **real road proximity**. Described in detail below.
- “**Suggest next slot**” — finds, within the next 30 days, the first date and time at which the specified staff are free, respecting shifts, weekly working hours, public holidays, absences and orders already assigned.
- “**Check availability**” — validates before saving that the chosen people are available in that window, with the readable reason for each block (outside shift, public holiday, absence, overlap). It does not prevent saving.

“Plan this WO”: qualified units and real proximity

It answers the complete operational question at creation —“*who do I give this to, and for when?*”—:

1. It evaluates teams as well as people. Many companies do not send Juan: they send the crew. A **team is a single unit** and its capabilities for the day are the **union of those of its members present**. A crew with an electrician and a labourer can carry out a job that requires an electrical certificate even though the labourer does not hold it. A team with no defined shift uses the company’s working hours.

2. Proximity is measured from where each unit actually sets off: the technician’s home if they have declared it so, their branch if not. The destination is the centre of the order’s assets or its own point or area. It also works in an organisation that does not work with branches.

Windows. The candidates are the start of each window in the branch’s weekly working hours, and the job has to **fit entirely** within that window: with a split working day there are two candidates per day and lunchtime is never proposed. The search looks over a 45-day horizon.

Output. Units are ranked by suitability, score (desirable skills and proximity) and travel time.

Unqualified units are also returned, with their reason. For each candidate it also checks whether they can reach **their next order** within the following 12 hours, with a 20-minute margin; if not, it is flagged with a warning.

If route calculation does not respond, units are returned without proximity instead of failing: **proximity improves the ranking, it does not decide suitability**. If the order has nowhere to measure from, it is planned by suitability alone, and the response says so.

5.5 Assigning staff and teams

- Assignment to operational **staff** (with or without an application user), marking one or more **persons in charge**.
- **Teams**: the form lets you add a team in one go; its active members are added as assigned people.
- Later reassignment from the form, from **Assign staff...** in the “...” menu, from the Kanban board, from bulk actions (5.15), from dispatch and from “Assign to the nearest...” (chapters 10 and 20). In no case does it change the status.
- If none of the assigned staff covers a mandatory skill or a required trade, the order is saved anyway with a **qualification warning**.

Each user’s scope: the same on the web and in the app. A technician only sees and can act on the orders in which their staff record appears (1.13).

5.6 Consumed materials

Each order carries a list of materials with **planned quantity** and **quantity actually used**. Each line also sets the **warehouse** it will come from and keeps track of what it has taken out of stock at any given moment, so that no route counts twice.

Which warehouse it comes from. On each line, “**Comes from**” offers **Automatic** or a specific warehouse with how much it has **free**. In automatic mode the server decides: the van of whoever does the work — or that of the vehicle they are driving that day—, the warehouse on their record, that of the branch or that of the main warehouse; never “whichever has the most”. What is chosen by hand is not changed by staff changes or by the planner. What open orders still have to take out of each warehouse counts as **committed** (chapter 16).

Material is **deducted from the warehouse when it is consumed, not when it is planned**:

WHAT HAPPENS	STOCK MOVEMENT
The office plans 5 units	None: the line sets its warehouse and commits that quantity
The technician reports from the app that they used 3	Consumption of 3
The order is completed without anyone reporting what was used	None: what was planned is not taken as used
The line is deleted or the order is cancelled	Return of what was taken out and not used

It makes no difference whether the consumption comes from the web, the app, offline synchronisation or the quick report: every route moves stock in the same way and records the movement with its cost, which feeds the order’s actual cost (5.10).

5.7 Field forms

Forms are linked to an order according to the maintenance type, the task or the asset type. Each link declares whether the form is filled in **once per order** or **once for each asset**, and whether it is **mandatory**. Details in chapter 18.

Mandatory forms block closure. An order cannot be completed, validated or closed with a mandatory form not completed, and the system says which ones are missing. A response marked as completed with empty mandatory fields does not count. It is checked on the web, in bulk actions and in the app; cancelling is not blocked.

Tasks are completed with their forms. A task becomes completed when all the forms linked to it have a completed response, and goes back to pending if any of them stops being so.

The **Forms** tab of the record shows the responses, including photos, and warns of those still pending upload; the PDF record includes signatures and form photos.

5.8 Attachments and evidence

Files linked to the order: photos, videos, PDFs, drawings. Each attachment stores its **thumbnail**, file type, description, author, **GPS coordinates of capture** and **capture date** (that of the photo, not of the upload), so that photographic evidence is georeferenced and dated. Photos are deleted from the app as a soft deletion, and on a validated, closed or signed order they can no longer be deleted.

There is also order **documentation** with title, description and file, with protected download and a time-limited link.

5.9 Customer signature

When closing the work in the field, the customer's **signature** can be collected: name, DNI (Spanish national ID number), role (customer, technician, supervisor), signature image, **GPS location**, date and device information.

When a signature cannot be obtained, an **exception** is recorded from a typed catalogue (customer absent, customer refuses...), where each reason declares whether it **requires a photo as evidence**.

Which orders require a signature is decided by each customer's **signature policy**: always, never, or depending on the maintenance type. **The policy is enforced on completion**: without a signature or an exception, the order does not move to completed, validated or closed, from the web, bulk actions or the app. Quick reports are exempt.

Each signature also carries an **integrity seal** of what was signed and a **timestamp** from an external authority, which the record shows alongside the signature (25.16).

5.10 Order cost

Report that calculates the actual cost of an order, for those who can see amounts and in organisations that work with costs. It aggregates three components:

- **Labour**: each person's **timesheet** (5.19) at their **rate in force on the day worked**: changing a rate does not rewrite past periods. Hours pending review do not count until they are approved. Those of people without a rate count in hours but not in amount, and flag the cost as **incomplete**.
- **Consumed materials**, valued at the cost of the movement, net of returns.
- **Vehicle**: with the fleet module, the **journeys** charged to the order and their fuel at each day's price, with a breakdown by vehicle and "View route on the map" (chapter 13).

Review orders arising from “not satisfied” and those marked as non-billable are flagged in the cost and in the reports (5.21, 20.10). Amounts never appear in the order PDF.

5.11 Display and lists

- **Work order layer** on the map, with filters by status, priority, type, person in charge, staff, customer and dates, and its own legend. Orders with a location of their own are drawn at their point or area (5.14).
- **Paginated list** in the right-hand panel, with search, sorting by due date, date, priority or “unassigned first”, restricted to the map area if desired and with counts by status. Each card shows the staff, the planned start and end, the **deadline**, the overdue and SLA warnings and the incident it resolves.
- **Daily counters** on the map (20.14.1).
- **Timeline per technician**: what each person has assigned and when.
- **PDF report** of the order: work sheet before completion, report of what was done afterwards and a record of work not carried out if it was cancelled; with planned and actual dates, time in progress, signatures with their seal and form responses, without amounts.

All these views respect the technician’s scope (1.13).

5.12 Planning views

The Planning tab offers six views of the same work:

VIEW	WHAT IT SHOWS
Day / Week / Month	Calendar of planned orders, with preventive maintenance forecasts (chapter 6). Only open orders can be dragged, and they cannot be dropped on a public holiday or on a day the company does not work
List	List of the week
Kanban	Board with the columns of the organisation’s cycle, with drag and drop between columns (those not permitted appear dimmed; dropping on Cancelled asks for the reason) and an Assign button on each card
User view	Workload per technician, with the dispatch queue (20.14.4)
Optimiser	Multi-day planner (chapter 11)

Moving an order in the calendar changes its planned window, never its deadline (5.25).

5.13 Orders pending review

A daily task detects orders that were left paused and have gone too long without being resumed (3 days by default, configurable) and flags them as “**needs review**” **without changing their status**, so that the supervisor is the one who decides.

5.14 Own location: a point or an area

Not all work falls on an inventoried asset: pruning in a street, a spill on a plot, an inspection of an area. An order can be located, in addition to or instead of its assets, at a **point** or in an **area** of its own, and that location takes precedence over that of the assets.

- The area is drawn on the map and is rejected if it has no surface or exceeds 50 km².
- With its own location, the sector and the branch are derived from that location (5.22).
- The point of an order located by its assets falls **on one of them**, not at a centre that could end up outside an L-shaped pipe.
- The own location is used by the geofence (chapter 15), “Plan this WO”, the app and the portal.
- The dashboard and the lists distinguish work **with an asset** from work **without an asset** (chapter 20).

5.15 Bulk actions

On a selection of orders —ticked in the list or with a rectangle on the map— the same action is applied to **up to 200 orders** at once:

ACTION	WHAT IT DOES
Assign	Sets the staff and, if specified, persons in charge and dates. It does not change the status
Priority	Changes the priority
Reschedule	Changes the start, the end or both, in each order’s local time
Cancel	Cancels with a mandatory reason

Each order goes through the same rules as individual editing and **is saved separately**: one that fails does not prevent the others. The result says which were applied and, for those that were not, why; on the web they remain selected.

5.16 Order timeline

The record has an **Activity** tab that brings together in chronological order what has happened to it: status changes (pauses, with their cause), arrivals and departures recorded by the geofence, automatic pauses, review alerts and work segments.

It includes **assignment events**: each person who joins or leaves the order leaves an event with who did it, when and **from where** —form, list or board, bulk action, dispatch, planner or schedule—.

The **Info** tab also brings together the schedule, the SLA pauses, the signatures with their seal, the tasks, the staff with their **timesheet**, the assets, the materials, **PPE and tools** (5.23), the forms, the documentation, the **actual cost** and, for reviews, the originating order (5.21).

5.17 Dates as instants and the order's time zone

Every date on an order is stored as an **instant** and displayed in the **order's time zone**: that of its branch; failing that, that of its first technician; failing that, that of the organisation (2.13). They are presented in the organisation's format.

- **Planned dates** entered as a day only: the start is the **start-of-day** time for that day and the end the **end-of-day** time, according to the shift of the person in charge, that of the team or the weekly working hours (9.8).
- **Days and durations** are calculated as the difference between instants: clock changes do not throw anything out.
- **From the app**, dates arrive with their time zone.

5.18 Actual dates, work segments and time in progress

Actual dates. The actual start and end of each order are stored together with **where they come from**: the geofence, a status change, the app, a manual correction. When something does not add up —end before start, a continuous segment of more than 24 hours, completed without an end, start before creation, impossible dates or dates far from the evidence—, the order is flagged with the reason.

The **Actual dates to review** screen lists those orders with a summary by reason, and lets you **mark them as reviewed** or **correct** the start or the end with a mandatory reason, which is recorded in the audit trail.

Work segments. An order that is started one day, paused and finished on another has not lasted “end minus start”, which would count the nights. Each move to in progress opens a **segment** and the pause or closure ends it. The sum of the segments is the order's **time in progress**.

The record shows actual dates, total duration and time in progress, and the Activity tab lists the segments. The same figure reaches the PDF, the portal, the export and the app, and is the basis for the time budget (chapter 12) and the timesheet (5.19).

5.19 Timesheet per person

An order's duration is not its labour. Each order has a **timesheet**: one row per person, day and hour type, which is what is read by the cost (5.10), hours per technician, staff capacity (20.11) and the export.

How it is calculated. For each person who worked on the order —operational staff always; a person in charge only if there is a trace that they worked on it—:

- **Work**: the order's segments or, if the app sends it, **each technician's stopwatch**, so that a colleague's pause does not cut their time.
- **Normal hours**: work that falls **within their working day** on that day, excluding public holidays, days off and absences.
- **Overtime**: work outside the working day **with evidence** of work (a start or end from the field, presence in the geofence or active tracking). It is created **pending review** and does not count until it is approved. Work outside the working day without evidence —the nights of an order forgotten in progress— is not charged.

- In a job lasting more than **7 days** in a row, the days without evidence come out pending review.

The timesheet is recalculated when the order's history, dates or staff change, and every night to pick up changes to shifts, public holidays and absences.

Review. The record shows the timesheet per person and day. Whoever manages staff **approves** or **rejects** what is pending, **corrects** a time slot, **removes** hours or **adds** hours by hand (for example, for someone who worked without the app). What has been reviewed stays fixed and recalculation does not touch it. The **Hours to review** inbox brings together what is pending across all orders.

Hours per technician. Screen and CSV export, in ranges of up to one year, that adds up the timesheet per person and day, with overtime shown separately.

5.20 SLA deadlines and justified pauses

Each order stores its **response** and **resolution** deadlines, where they come from (the customer's, the organisation's or the global SLA) and since when it has been **at risk**. Deadlines are counted in **working hours**: they only consume the windows of the weekly working hours, skip public holidays and use the branch's time zone. The deadline counts from a **fixed baseline** —creation or, for preventive maintenance, its scheduled date—: moving the order in the calendar does not move its deadline. Details in chapter 19.

Justified pauses. When pausing an order, the **cause** is given —attributable to the customer, a third party or supply, force majeure or own—. The first three **stop the resolution deadline** if the customer's contract allows it; own causes do not. The web always asks for the cause, because it also feeds the dashboard's breakdown of time to closure. The record lists the pauses with their cause and what was deducted, and the cause of a pause already recorded can be corrected, which is recorded in the audit trail (19.8).

5.21 Review following a customer's "not satisfied"

When the customer answers **Not satisfied** to a finished order in the portal (chapter 22), the system opens a **review order** at that moment instead of reopening the original:

- **Pending, high** priority and **without a technician**, with the customer, branch, assets, location, type and tasks of the original, and the customer's comment in the description.
- It is created **non-billable**; whoever can close orders can change this, and each change is recorded in the audit trail.
- A single review per completion cycle of the original: another "not satisfied" is added to the existing one.
- The original does not change status and can now be closed; acceptance is requested on the review.
- If the original was resolving an incident, the incident goes back to open and the review takes over resolving it.
- The "**Customer not satisfied**" alert is raised, which closes on its own when staff are assigned to the review.

The review's record links to the original and to who worked on it; costs separate what was spent on reviews and on non-billable work.

5.22 Responsible sector and affected sectors

An order has a **responsible sector** and can **affect several**:

- **Located by its assets**: the responsible sector is the one that occurs most among them.
- **With its own point or area**: the distribution of its location across the sectors, with a margin so that small changes do not make it jump.
- **Affected sectors**: all those its assets or its location pass through, with their share. They are used to **filter and count** work by area (an order crossing two sectors counts in both); **costs** go only to the responsible sector, so as not to count them twice.

When the geometry of the sectors changes, only open orders are reassigned.

5.23 Material, tools and PPE for tasks

Each task in the catalogue can declare what is needed to carry it out (8.8). On the order:

- The form shows “**What the chosen tasks call for**”. **Add material from the tasks** copies it as normal lines, multiplying the quantity per asset if so defined. It is a proposal, not a reservation.
- **Tools** appear on the record and in the app as a list for preparing the job; they do not deduct stock.
- **PPE** must be confirmed in the app when starting the order. The record shows what each technician confirmed: whether they have it or not, with their reason. Starting without some item of PPE generates a **safety alert** for the supervisor, which closes on its own only if the technician corrects it afterwards. **Starting is never blocked**, because the app works offline.

5.24 Order that does not resolve its incident

When finishing an order that resolves an incident, the web and the app ask “**Is the problem resolved?**”:

- **Yes** (or no answer): the incident follows the order’s status as usual.
- **No, another intervention is needed**, with a note of what is missing: the incident **remains open** with its SLA running, this is noted in its history, the new intervention alert is raised and it **accepts another order**, which carries in its description what was missing.

This way, a diagnosis or a temporary fix does not count the incident as resolved. Reopening the order clears the answer, which is asked again.

5.25 Planned window and deadline

The order distinguishes its **planned window** from its **deadline**:

- **Planned start** and **planned end** are the planning: they are changed by the form, the calendar, dispatch, bulk actions and the optimiser.
- **Deadline**, optional: the one requested by the customer or the contract. It is only changed by whoever edits it, never by planning. Preventive maintenance is created with a deadline equal to the generated end.

The **due date** used by alerts, counters, dashboard, map, lists, dispatch, portal and planner is, in this order, the deadline, the SLA resolution deadline and the planned end. Replanning never falsifies a contractual due date. Preventive compliance is measured against the deadline and SLA compliance against its deadline.

06

Recurring scheduling of work orders

AVAILABILITY	Maintenance module, with the preventive capability
SURFACES	Web
WHAT IT IS FOR	It turns preventive maintenance into something that happens on its own. The quarterly inspection of the electrical panels is defined once and the system generates the orders with the configured lead time, warning if the warehouse cannot cover the material.

6.1 Recurrence rules

On a “template” work order a rule is defined in the **Recurrence** block of the form, which appears with periodic maintenance types:

- **Start date** and end date (or indefinite recurrence).
- **Frequency**: daily, weekly, monthly or yearly, with an **interval** (every N periods).
- Specific **days of the week** (for example, Monday and Wednesday), with weekly frequency.
- Specific **days of the month** (for example, the 1st, 15th and 30th), with monthly frequency. A day the month does not have falls on the **last day** of that month, instead of skipping the month.
- Specific **months of the year** (for example, March and September), with yearly frequency.
- **Time of day** of each occurrence, in local time: a 10:00 preventive task stays at 10:00 after a clock change. Without a time, it takes that of the template order.
- **Next visit: fixed or floating** (6.5).

Working days. An occurrence that falls on a non-working day is moved to the next working day, according to the branch’s weekly working hours and its public holidays.

End and deadline of each generated order. They keep the duration of the template order, counted in working days.

6.2 Automatic generation

Generation works with **two horizons** from the organisation’s settings:

HORIZON	DEFAULT	WHAT IS CREATED
Orders	7 days	Real orders , ready to execute
Planning	90 days	Forecasts : entries that can be seen in the calendar but are not yet orders

It runs when the schedule is created or edited and every night. The window **slides**: every night the forecasts that enter the first horizon become orders and the horizon extends by one day.

- **It never generates past days.**
- **It can be repeated without duplicating.** Each order remembers the day of the occurrence it covers: a day that already has its order does not receive another, and **moving a generated order to another day** does not cause it to be regenerated on its original day.
- **Excluded days** (forecasts deleted by hand, 6.4) are not generated.

Each generated order is created **pending**, with its own code and the template's title with its **sequence number** (“#017”); it is created with a **deadline** equal to its generated end, and its **SLA deadline counts from the scheduled date**, not from when it was generated days earlier; it copies assets, tasks, staff, materials and forms **without** what is no longer valid (decommissioned assets, deactivated staff, inactive tasks); and it alerts its persons in charge. If the template order is left without assets or without active staff, it is generated anyway and an alert is raised.

6.3 Material warning

When **the schedule is created**, the response warns if the warehouse on each line cannot cover the template order's material or leaves it below minimum: this is known before orders start coming out. Generated orders **do not move stock when they are generated**: the material is committed and is deducted when consumed (5.6).

6.4 Query and clean-up

- **Query** of forecasts by date range and by customer.
- **Create order** from a forecast turns it into a real order, and the forecast disappears.
- **Deleting a specific forecast** leaves its day **excluded**, so that nightly generation does not create it again.
- **Editing the schedule** only regenerates if the rule or the content of the template order changes. On regeneration, the generated orders that are **pending and untouched** from today onwards are replaced; those that have something of their own (attachments, comments, changed staff, responses, another date) or already have work done are **kept**, and the web warns of their codes so that the user can decide. Pending orders from past days are kept: they are overdue preventive maintenance.
- **Removing the schedule** keeps the template order; of the generated ones, the untouched pending ones are deleted and the rest remain as the asset's history.
- The record of a closed or cancelled template order with an active schedule offers **Remove schedule**.

6.5 Next visit: fixed or floating

Two ways of calculating when the next visit is due (UNE-EN 13306):

MODE	THE NEXT VISIT GOES...	WHAT FOR
Fixed (default)	To the previous theoretical date plus the frequency. A delay does not move the following ones	Statutory inspections, inspections, contracts with a calendar
Floating	To the actual end of the previous visit plus the frequency. A delay moves the next one	Wear and consumables: not bringing work forward

In floating mode the rule is only “every N days, weeks, months or years”, and while a visit is open the next order is not created, only forecasts.

6.6 Preventive compliance

The dashboard’s **Preventive compliance** indicator (chapter 20) and its list use the same criterion: orders generated by schedules dated today or earlier, excluding cancelled ones, measured against their **deadline**. The customer portal shows the planned preventive maintenance (chapter 22).

07

Incidents

AVAILABILITY	Maintenance module, with the incidents capability
SURFACES	Web · Mobile · Customer portal · Public API
WHAT IT IS FOR	It separates the <i>report</i> from the <i>work</i> . An incident is something someone has reported —a customer, a technician, an external system— that is not necessarily an order yet. It makes it possible to triage, measure response time and decide what becomes work and what does not.

7.1 Origins

Catalogue of input channels with multilingual name, icon, colour and order. There are eight system origins: **Technician app**, **Customer portal**, **Call / phone**, **Email / register**, **IoT sensor**, **Inspection / audit**, **Meter reading** (those the system opens because of an out-of-range reading, 7.15) and **Integration (API)** (7.16). The organisation can add its own. This is what makes it possible to answer “how many reports come in through the portal compared with the phone?”.

7.2 Reporter

An incident can come from:

- A **system user** (who is linked to it).
- A **customer portal user**, linked to their account: they can always see the incidents they reported (22.2).
- An **external party**: free-text name and contact (email or phone) so that they can be notified.

7.3 Content and classification

- Human-readable code generated automatically, with a sequence per organisation and year (format INC-2026-00056).
- Title, separate **public description** and **internal notes**: internal notes are visible neither to the reporter nor in the portal.
- **Incident type** (the organisation’s own catalogue, 7.14), **status** and **criticality level**.
- **Assigned reviewer**.
- **Customer**: the one specified or that of the assets. An incident **can remain “unclassified”** while it is open, because whoever reports it does not always know whose problem it is; but without a customer it cannot be resolved, closed or converted into an order.

- Responsible branch, resolved automatically.

STATUS	MEANING
Pending	Just reported, awaiting review
Approved / WO created	Being handled with a work order
Resolved directly	Resolved (also when the order handling it is completed)
Closed	Definitively closed
Rejected	Discarded; also the status of deleted incidents (7.11)

Resolution date. It is stamped on moving to resolved or closed, and it is the one that measures the SLA deadline. Returning to an open status clears it.

Editing according to status. Pending and approved incidents can be edited; a resolved one can only be reopened or closed by an administrator; closed, rejected and deleted ones cannot be edited.

7.4 Location

An incident is located on one or more **assets**, a **point** or an **area**:

- **Affected assets:** an incident can have **several**; the first is the main one.
- **GPS position** exactly as captured by the device.
- An **affected area** of its own —a pavement, a stretch of park— drawn on the map, of up to 50 km².
- Fallback text address.

On the map, the context menu offers “**New incident here**”, and an asset’s record lists its incidents. Those coming from the app also store whether the phone flagged its position as **simulated** (14.10).

7.5 SLA and deadlines

Each incident in an organisation with SLA receives its **response and resolution deadlines**, with the same criterion as orders: the incidents SLA for its criticality, from the customer, the organisation or the global one, counted in working hours. They are recalculated when the criticality, the customer or the branch changes while it is open; a resolved, closed or rejected incident **keeps its deadline**.

The status of each deadline is **visible** on the card and on the record:

STATUS	MEANING
Met	Responded to or resolved before the deadline
Missed	Responded to or resolved late
On time	Open and within the deadline
At risk	Open and with the risk percentage consumed (19.5)
Overdue	Open and past the deadline

The **response** is met when the incident leaves pending for the first time; the **resolution**, with the resolution date. The deadlines feed the dashboard and the customer portal.

7.6 Conversion to a work order

- **Generate an order from the incident:** creates the pending order, with type, priority (proposed from the criticality) and person in charge, and inherits the **customer**, **all the assets** or the **point** or **area** and the address. The incident moves to *Approved / WO created*.
- An incident is handled with a **single live order**: a second attempt is rejected, indicating which order is handling it, even if two people press at the same time.
- Only an **open** and classified incident can be handled with an order.
- **Link an existing order**, with a “*resolves*” or “*detects*” relationship and notes.

7.7 Files

Attachments with title, description and **thumbnail**, from the web, the mobile app and the portal. Each file is checked by its content, not by its extension:

- **Web:** images, PDF, modern Office documents without macros and text, up to 20 MB.
- **App and portal:** images.

Files are served after checking access, and the portal only downloads those of its customer. The record shows the photos in a **gallery**.

7.8 Display and statistics

- **Incident layer** on the map, with its own legend and filters (status, criticality, type, dates, customer, person in charge or staff of the order handling them), applied in the same way to the map, the framing and the list. Affected areas are drawn in a layer of their own.
- **List** of cards with the SLA status and an actions menu.
- **Statistics:** reported, resolved, rejected, pending, resolution rate, average time to first response and average resolution time.

7.9 Automatic alerts

Every incident created pending generates an alert, wherever it comes from (web, app, portal, out-of-range reading or API): of high or critical severity if its criticality is, medium otherwise. It is received by those who have incident alerts assigned in their branch, never by whoever created it, and it closes for everyone as soon as the incident stops being pending.

When an order ends without resolving the problem (7.12), the “**Unresolved incident**” alert is also raised, so that someone creates the next intervention.

7.10 Status history

Every status change is recorded with the previous status, the new one, the author, the date and notes, including those caused by the linked order, the orders that ended without resolving the problem and deletion with its reason.

7.11 Deletion

Deletion is a **soft deletion** with strict rules:

- It can only be done by a **supervisor** or an **administrator**.
- The **reason is mandatory** and is stored.
- **An incident with a current work order cannot be deleted.**
- The incident disappears from lists, indicators and the map, and its alerts are withdrawn.
- It can only be deleted from the web; the app removes it at its next synchronisation.

7.12 Status synchronised with the work order

With the answer “**Yes, it is resolved**” (5.24), the incident’s status follows that of the order, whether it is changed from the web or from the app:

THE ORDER MOVES TO...	THE INCIDENT MOVES TO...
Completed	Resolved
Closed	Closed, keeping the resolution time
In progress, reopened	Approved / WO created
Cancelled	Approved / WO created, unless another live order is still resolving it

With “**No**”, the incident remains open with its deadline running, accepts a new order and the unresolved incident alert is raised. If the customer answers “**not satisfied**” in the portal, the incident goes back to open and the review order takes over resolving it (5.21).

7.13 Offline incidents

The app can create incidents without coverage and upload them later. **A resend does not duplicate it.** On upload all its data is validated, and assets created without coverage in the same upload are processed first, so that the incident finds them.

The app **does not move the status backwards**: what the office has already approved, resolved or closed does not go back to pending. Changes to an incident that the office has also touched are resolved as **synchronisation conflicts** (23.15).

7.14 Incident types

Each organisation’s own catalogue, with name, description and asset flag. A type in use cannot be deleted: it is deactivated.

7.15 Incidents from out-of-range readings

When a meter reading falls outside the valid range of its measurement point, and the point has this enabled, the system opens a “**Possible faulty meter / Leak**” incident, of high criticality, with the photo of the reading if there is one. At most one per reading. Details in 17.9.

7.16 Creating incidents via the API

An external system (an ERP, a reporting platform) can create incidents through the public API, with the source system’s **reference**, protection against resends and keys that can be restricted to one customer. They go through the same process as web creation —same validations, alerts and deadlines— and come in with the *Integration (API)* origin. Subsequent changes can be notified by webhooks (chapter 25).

08

Maintenance types, tasks, trades and skills

AVAILABILITY	Core; maintenance types and tasks require the maintenance module; trades, skills and training depend on the competencies capability, and task materials on the warehouse capability
SURFACES	Web · Mobile (viewing, and confirmation of PPE on starting)
WHAT IT IS FOR	Defines what work exists, what you need to know to do it and what you need to bring . It is the layer that turns work assignment from “let’s see who is free” into “who is qualified, certified, available and nearby”, and work preparation from “remembering” into “what the task calls for”.

8.1 Maintenance types

The organisation’s catalogue (cloned from the vertical). Each type defines:

- Name and description.
- Whether it is **cyclical** (recurring preventive maintenance).
- Whether it **requires the customer’s signature**, which applies to customers whose signature policy is “according to the maintenance type” (5.9).
- **Estimated duration** of a work order of this type (chapter 12).
- Whether it is **active**.

Every organisation also has the system type “**Immediate corrective**” and the “Quick work report” form, which are used by quick reports from the app.

A type that has already been used in a work order cannot be deleted: it is deactivated.

8.2 Tasks

Within each maintenance type, the specific tasks are defined. Each task has:

- Name and description.
- **Its own estimated duration**, which is inherited from the type if left empty.
- Whether it is **recursive**: it is broken down across the asset’s components according to its hierarchy.
- Whether it is **active**.
- What is needed to carry it out: **materials, tools and PPE** (8.8).

In a work order, each task has its own **status** (pending, in progress, completed), and is completed with its forms (5.7).

8.3 Trades

Catalogue of specialities: electrician, plumber, gardener, aerial work platform operator... They are linked in two directions:

- **Task** → **required trades**: which speciality is needed to carry it out.
- **Person** → **trades**: which specialities each technician has.

When a task requires trades, having **one** of them is enough; having none **blocks** eligibility. A trade in use is not deleted: it is deactivated.

8.4 Skills and certifications

Catalogue with name, code, **category** and **type**, whether it is **official** (a recognised certification), whether it **expires** and its **default validity in months**, and the default **issuing authority**.

8.5 Staff skills

Each person accumulates their skills with:

- **Competence level**: basic, intermediate, advanced, expert.
- **Issue date** and **expiry date**.
- **Certificate number**, **issuing authority** and attached **document**.
- **Status** (active, expired, revoked, suspended) and **verification**: who verified it and when.

In the application this appears as the “**Training / Courses**” module, and each user sees their own in their profile.

- **Verification** serves as evidence: neither the person who entered the record nor **the person themselves** can verify it, and changing the level, the dates, the certificate or the document **voids the verification**.
- Removing a skill marks it as revoked; it is not deleted.

A skill whose **expiry date has passed does not count**, even if it is still marked as active: the eligibility engine, the planner, dispatch and work order validation all require it to be valid **on the day of the work**.

8.6 Skill requirements for work

Skills are required at two levels, with inheritance: **by maintenance type** (all its tasks inherit it) and **by task**, which can override what is inherited. Each requirement declares:

FIELD	VALUES
Requirement level	mandatory (blocks) · recommended · preferred

FIELD	VALUES
Minimum competence level	basic / intermediate / advanced / expert

Mandatory skills filter: anyone who does not hold them, valid and at the required level, is not suitable. **Recommended and preferred** skills never rule anyone out: they score in the eligibility engine (10.1) and in dispatch, and the planner tries to give each work order to someone who has them (11.10).

8.7 Expiry control

A view of **skills about to expire** with a configurable horizon, filterable by alert level, branch and category, with an aggregated summary, so that renewals can be anticipated before a technician stops being eligible for their usual work.

8.8 Task materials, tools and PPE

Each task in the catalogue can declare what is needed to carry it out:

WHAT	HOW IT IS DEFINED	WHAT HAPPENS IN THE WORK ORDER
Materials	Inventory item with its quantity, fixed or per asset , and notes	It is proposed when the work order is created; it is never added automatically (5.23)
Tools	Items flagged as a tool	Informative list on the record and in the app; they do not deduct stock
PPE	The organisation's catalogue of personal protective equipment	The app asks the technician to confirm it when starting the work order

PPE catalogue. Name, description and whether it is active, with a **basic list** (helmet, gloves...) that can be loaded in one click. PPE in use is not deleted: it is deactivated.

What a work order calls for is calculated from its current tasks: a pending work order picks up any PPE added to its tasks later.

PPE does not block work. The app works offline, so starting is not prevented: what each technician confirms is recorded, item by item, with the reason if any item is missing, and a “**Start without PPE**” safety alert is generated for the supervisor (chapter 24).

09

Staff, shifts and availability

AVAILABILITY	Core; shifts, public holidays, absences and exceptions depend on the working calendar capability
SURFACES	Web · Mobile (declaring on-call duty and overtime)
WHAT IT IS FOR	Models the real availability of the workforce. It is what stops the planner from assigning work on a public holiday, during a holiday period or outside the person's shift.

9.1 Staff record

- Name, phone, email, notes.
- **Internal or external** (subcontractor), with the company name.
- **Hourly cost**, with a history of rates by date (5.10).
- Linked **application user** (optional); a record without access can be given access later (1.9).
- Assigned **branch** and **default warehouse**.
- **Default shift** and **default vehicle**.
- Their own **time zone** (optional); if empty, it inherits that of their branch.
- **Home location**.
- **Where the working day starts and ends**: at their **branch** or at their **home**. This is used by fleet trips (chapter 13) and by the proximity in “Plan this WO” (5.4).

Restricted data. The rate, the notes, the home address and the reason for absences (sick leave is health data) are visible only to those who can edit staff; everyone else sees who is available, not how much they are paid or where they live.

Deactivation and erasure. Setting a person as inactive while they have unfinished work orders or materials in their personal warehouse warns of what is outstanding and allows confirmation. **Deleting them** is a **GDPR erasure**: they are anonymised along with their account and their location data, their work orders, costs and histories are kept, and their timesheet is frozen (chapter 25).

9.2 Shifts

Catalogue of shifts:

- Name, start time and end time.
- **Working days** of the shift.

- **Break** within the shift, which the planner respects as a pause.
- Shift linked to a **branch**.
- Whether it is **active**.

Shifts are assigned to people and to teams. A shift in use is not deleted: it is deactivated.

Shift times are **the technician's local time**. A shift that **crosses midnight** (for example, from 22:00 to 06:00) belongs to the day on which it starts and ends on the following day. This is how it is read by availability, the planner, dashboard capacity, GPS tracking and automatic end-of-day closure.

9.3 Public holidays

Public holiday calendar, with:

- Date, name and **scope** (national, regional, local or company).
- A specific **branch**, or global for the whole organisation.
- Whether it is **active**: a deactivated public holiday (work goes ahead that day) does not count.
- **Bulk entry** of a year's calendar. The upload is **matched by date** against what already exists: what remains is kept, what has changed name is renamed and what is new is added; what no longer appears is proposed for removal, but is not removed without confirmation.

Each public holiday applies only on its date: movable holidays (Easter) cannot be repeated correctly from one year to the next, so each year's calendar is imported or created. Within a single calendar, a day either is a public holiday or it is not: there cannot be two public holidays on the same date.

A branch's calendar shows **its own public holidays and those of the whole organisation**, which is how availability, capacity, SLA and the planner deduct them.

9.4 Absences

Holidays, sick leave and other leave, with:

- Date and time range, and a **full day** or partial flag, in the **worker's time**.
- **Absence type** (holiday, sick leave, personal leave, training, maternity/paternity, unpaid, other) and **status** (confirmed, provisional, cancelled).
- **Person covering** the absence.
- **Overlap check** against confirmed absences.
- Cancellation without deleting the record. An absence that has already begun is not deleted, because it is part of the past calendar: it is cancelled.

Only **confirmed** absences reduce availability: in the planner, the eligibility engine, dispatch and the timesheet.

9.5 One-off calendar exceptions

For the specific day on which someone starts later, works a split shift or changes shift. Five types:

TYPE	EFFECT
Day off	The person does not work that day
Shift change	Another shift from the catalogue applies that day
On-call duty	Availability window with a start and end time
Extra working day	Works on a non-working day
Extended working day	Hours in addition to their shift

On-call duty or an extra working day can cross midnight, just like a night shift.

The technician can **declare from the app** an on-call period, an extra working day or an extended working day themselves. This exists because in a night-time emergency nobody creates it in time, and without it their GPS tracking and mileage would not be recorded. It takes effect immediately, is marked as **self-declared** and **pending review**, and a manager approves, edits or deletes it.

Every user with a staff record sees “**My calendar**” in their profile (1.11).

9.6 Cascade for resolving the effective schedule

For a person and a date, the schedule that actually applies is resolved in cascade:

1. **One-off exception** for that day, if there is one.
2. The person’s **default shift**, if that day is a working day in the shift.
3. With no shift, the **weekly schedule** of their branch or of the company (9.8).
4. Deducting the **public holiday** applicable to their branch.
5. Deducting **confirmed absences**.

Times are read in the person’s time zone. The same cascade is used by the planner, GPS tracking, automatic end-of-day closure, the timesheet, the eligibility engine and the creation assistants.

9.7 Availability query

An availability query for a person at a given moment and for a given duration, which returns whether they are available and the **reason**: deactivated, absence, another overlapping work order, day off, on-call duty or extra working day, public holiday, no shift but within company hours, shift deactivated, non-working day, outside working hours, on their break or within their shift. The work order form shows these reasons with their readable label.

9.8 Company working hours and working days without a shift

The **weekly working schedule** is configured for the company (2.8) and, if a branch needs it, as the **branch’s own schedule** (3.5). It has up to **two time slots per day** of the week (split day) and supports **night-time slots**; a day without slots is not a working day. It is read by availability, SLA (19.9), the work order’s working day, GPS tracking, the timesheet and the planner.

- A person **with no assigned shift** works within the slots of their branch's or the company's schedule, including in the planner.
- The planner adjusts each resource's shift window to its branch's schedule day by day and treats the gap between slots as a break. A night shift is not cut short by office hours.

9.9 Time zone

Shifts, absences and schedules are in **local time**. A person is read in their own zone, or in that of their branch, or in that of the organisation (2.13). In this way, dashboard capacity deducts each person's absences in their zone, the planner builds each resource's working day in theirs, and absences are edited in the worker's time.

9.10 Hours per technician

A screen and CSV export with the **hours worked per person and day** over a range of up to one year, filterable by person and branch. They come from the work orders' **timesheet** (5.19), with normal hours and overtime separated and anything pending review shown apart.

10

Eligibility and smart dispatch

AVAILABILITY

Core; planning a work order from its creation requires the planning module; dispatch to the nearest and insertion into the working day, the real-time field module

SURFACES

Web

WHAT IT IS FOR

Answers the real operational question — “who do I give this to, and for when?” — by combining qualification, valid certification, time availability and **real road distance**, not straight-line distance. Travel times are those of the road network, **without traffic**.

10.1 Eligibility engine

For a task, a date and time and a duration, the system evaluates the entire active workforce and returns for each person:

DATA	MEANING
Is eligible	Meets everything : available + valid mandatory skills at their minimum level + required trades
Is available	Free in that time slot (shift or schedule, public holiday, absence, other work orders)
Has the mandatory skills	All of them, active, not expired on the date of the work and at the minimum level required
Has the required trade	Any of those the task calls for
Recommended and preferred skills	How many of each group they meet
Same branch	Belongs to the work order’s reference branch
Blocking reason	Why they are not eligible: unavailability, missing mandatory skills, a trade they do not have, or a role with no field work
Score	The result of the assessment

Competency score: only eligible candidates are scored, with 30 points for each recommended skill they meet and 10 for each preferred one.

People with a **read-only** or **warehouse** role are shown as unavailable: they do not do field work.

10.2 Real proximity scoring

On top of that base, the system adds a **proximity score of 0 to 50 points**:

1. It takes the work order's location: its own point or the centre of its assets.
2. It calculates the **real road travel time** to each active branch.
3. It converts it into points: under 15 min, 50; under 30, 35; under 60, 20; under 90, 10; more, 0.

If the work order has no location or route calculation is unavailable, the results are returned **without** proximity rather than failing. "Plan this WO" (5.4) uses a variant that measures from **each unit's actual starting point**.

10.3 Team eligibility

There is an equivalent engine for **teams**, with the same criteria applied to the **union of the members present** in the time slot: a team is eligible if any member is available, the union of all of them covers the mandatory skills and one of them has the required trade. It also returns **why it is not eligible** when it is not: no member available, non-working day or time slot outside the team's shift, missing skills or trade.

10.4 Multiple tasks in one work order

When the work order has several tasks, a person is eligible only if they meet the requirements of **all** of them. Desirable skills are counted without inflating the score because they are shared between several tasks.

10.5 Search for the first time slot with a suitable resource

With no date set, the system **searches for the first date and time** at which there is at least one unit — team or technician — that is qualified and free, within the slots of the weekly schedule, with the work fitting entirely within one slot and never proposing lunchtime. It looks ahead at most **45 days**. If it finds nobody, it returns the evaluation of the first time slot, so that the interface can explain **why** there is nobody rather than showing a blank.

10.6 Dispatch to the nearest

Dispatch for emergencies: given a point on the map, it returns the **nearest technicians** (5 by default, up to 20) among those with **active tracking** and a position from the last two hours, with:

- Their **last position** and how long ago it was received.
- Their active **vehicle**.
- Whether they already have a work order **in progress**: they are still listed, marked as busy.
- Filter by **required skills**, valid and at their minimum level.
- **Maximum search radius** (20 km by default, up to 200), which is widened once if nobody is found.

It proposes only people who can go: **active** staff with **no confirmed absence**. A recent position is not enough, because the phone of someone on holiday may still be sending points.

Ordering uses **real road travel times**: the technician 2 km away on the other side of the river is not the nearest.

For a **specific work order**, dispatch additionally requires the mandatory training for its tasks at their minimum level, and shows for each candidate their **desirable training** (with stars, without ruling anyone out) and whether they **would be missing materials** (11.11). It is used by the “**Assign to the nearest...**” action, from the work order, the map or the dispatch board.

10.7 Assistant warnings

“Plan this WO” adds two warnings, which do not rule anyone out:

- **Cannot reach their next work order**: if finishing, travelling and leaving a 20-minute margin does not fit before their next work order within the following 12 hours.
- **The work order has nowhere to measure from**: with no located assets and no location of its own, it is planned on suitability alone, and this is stated explicitly.

10.8 Fitting a work order into a technician's working day

For a **pending** work order, each candidate is shown with **what it costs to fit it into their working day**: the system recalculates their route for the day without the new work order and with it, and the difference is the impact.

Candidates. Nearby technicians with a GPS position and, in addition, those whose **branch** is nearby, with or without GPS, with the mandatory training and without office roles.

Which day. First **today**, with what remains of their working day; if it does not fit, each of the **next three days**, until the first one in which it fits. The working day is the planner's: shift, schedule, public holidays, absences and breaks.

Their schedule. Pending work orders that only they hold can be moved, within their deadline; shared ones, paused ones and those without a location are not moved, and the work order in progress is respected. An **urgent** work order can **displace** lower-priority ones, which go back to “To dispatch”; a non-urgent one does not push out anything of higher priority.

What each candidate shows. Whether it is viable and on which day; what time they would arrive and whether they arrive **within the deadline**; how much **extra driving** it involves; which work orders on their schedule **change time** and which **no longer fit**; their desirable training; and whether the materials are available or they would have to pick them up (11.11).

Nothing is applied automatically. The supervisor chooses a candidate and, on confirming, the system first checks that none of the work orders involved has changed since the proposal; if any has, it must be recalculated. Each change is recorded in its work order's timeline.

11

Multi-day work order planner

AVAILABILITY	Planning module
SURFACES	Web (“Optimiser” view in the Planning tab)
WHAT IT IS FOR	Automatically distributes the backlog of pending work orders among the available resources over several days, minimising travel and respecting all the real constraints of the workforce. It replaces manual planning on a calendar.

11.1 What it solves

A vehicle routing problem with time windows over several days, using real road times and distances (without traffic). The scope is **local**: it is not designed for routes of hundreds of kilometres.

Each combination of **resource × available day** is a vehicle in the problem, so that a single optimisation distributes the work across **the whole horizon**, instead of solving day by day and chaining suboptimal decisions.

A **team is a single route unit**: its daily capabilities are the union of those of the members present that day.

Which work orders are included. Pending work orders not yet planned whose due date (5.25) is not later than the range. If no specific work orders are chosen, only those near the resources’ branches are included and, with a customer selected on the map, only theirs. If they are chosen, the selection prevails, and those that are not candidates are explained with their reason. A work order without a location remains unassigned with that reason.

11.2 Constraints it respects

CONSTRAINT	SOURCE
Shifts and exceptions, including night shifts	Effective schedule (9.6)
With no shift, the working day from the weekly schedule	Schedule of their branch or of the company (9.8)
Team’s default shift	Team record
Working days and the lunch gap	Weekly schedule of the resource’s branch
Public holidays	By the resource’s branch; in a team, each member’s
Confirmed absences and days off	Staff calendar

CONSTRAINT	SOURCE
Breaks within the shift	Shift definition
Schedule already assigned	Open work orders of those people within the range (11.9)
Valid mandatory skills at minimum level, on each day of the plan	Task or type requirements
Required trades	Trades of the task and of the person
Recommended and preferred training, as a preference	11.10
Materials, if the organisation requires it	11.11
Service time of each work order	Time budget (chapter 12)
Work order window	Planned start and due date
Travel times and distances	Road network
Start and end point of the working day	The technician's or team's branch

The validity of certifications is checked **day by day**: a certificate that expires halfway through the horizon is valid until that day and not after.

Night shifts. A shift that crosses midnight is planned in full, with its break inside it, and whatever falls after midnight is marked “(+1)”.

Teams. A day counts with the members present, excluding those absent and those with a public holiday in their branch; when applying, work is assigned only to those present.

Starting point. Each technician **leaves from their branch and returns to it**. If the branch has no location, they leave from a warehouse and the plan warns so that it can be completed.

Excluded days. The plan lists each resource-day that could not be used and why: no shift, non-working day, public holiday (with its name), absence, all members absent, the previous day's night shift or a full schedule.

11.3 Priority with due-date pressure

The priority used by the optimiser adds a **due-date pressure** to that of the work order:

POINTS	
Base	Low 10 · Medium 50 · High 80 · Urgent 100
Already overdue	+30
Due in 2 days or less	+20
Due within the planned range	+10

(with a maximum total of 100). An overdue medium-priority work order competes on equal terms with a non-urgent high-priority one: the work orders under pressure push those that are not off the route. An overdue work order is planned in the first available slot.

11.4 Balanced distribution

“**Balance workload**” option (enabled by default): it limits the number of new work orders per resource and day, so that the work is spread out instead of being concentrated on whoever has the cheapest route.

11.5 Operational limits

- **Maximum horizon: 14 days.**
- Up to 500 resource-days and **500 distinct locations** per plan.
- Minimum useful window of 1 hour after deductions for schedules, absences and breaks.
- Each resource is planned **in its own time zone**.
- At least one team or one technician must be selected.

11.6 Plan lifecycle: draft → applied

A plan **changes nothing** when it is generated. It is created as a **draft** and stores the range, the author, the complete input and output of the calculation (so that every decision can be explained) and a **summary**: routes per resource and day with the arrival and departure time for each work order, assigned work orders, unassigned work orders with their explanation (11.8), excluded days, warnings, and total distance and duration.

From there it can be:

- **Applied** — before anything is written, the system checks that the plan **is still valid**: if any work order has changed, if someone is no longer available at that time or if the planned time has already passed, **nothing is applied** and the changes are reported, so that another plan can be generated. If it is valid, the work orders receive their staff and their planned dates, remain pending and appear in the calendar, the Kanban and the per-technician view. Each resource’s existing schedule is not touched.
- **Discarded** — the plan is archived with no effect.

11.7 Resilience

If the route calculation or optimisation service goes down, the system stops trying for a few minutes and says so clearly, rather than leaving requests waiting. If road times cannot be obtained, planning goes ahead anyway with an **approximate estimate** and the plan carries a warning to review it before applying it.

11.8 Why a work order is left uncovered

The plan detail explains, under each unassigned work order, the reason and what to change to resolve it:

REASON	WHAT THE EXPLANATION SAYS
No qualified resource	Which skills or trades are missing, by name
Duration longer than the working day	Duration versus working day
Too far away	Kilometres and breakdown of outward + work + return
Incompatible dates	The work order's window
Full working days	How many it would fit into, and whether they have reached the balanced-distribution cap

Before the calculation, a work order can also be left out for having no location, because its fixed staff are not in the plan, or for lack of materials (11.11).

11.9 Schedule already assigned

Time that staff have already committed is not free. Open work orders of the selected technicians with a time slot within the range are included as **fixed stops**: they take up their work and travel time, only within their time slot, and **are not replanned**. The detail shows them in their place on the route as “already on their schedule”. A working day that is already full receives nothing new.

11.10 Recommended and preferred training

Desirable training does not rule anyone out, but the planner tries to give each work order to someone who has it. The result is only accepted if it **covers at least as much** as the plan without preferences; if not, the usual plan is kept and the plan warns of this. Each work order shows how many of its desirable skills the person doing it has.

11.11 Materials when planning

With the warehouse capability, the organisation chooses what the optimiser and dispatch do with the work orders' **outstanding materials**:

MODE	EFFECT
Ignore	Materials are not checked
Warn (default)	Each stop states whether the warehouse it would come from is short of anything, with quantities. It does not change the distribution
Require	A work order with materials only goes to someone who can take them

In **Require** mode, for each resource: if they carry it in their **van**, they go directly; if it is in their departure warehouse, they load it on leaving; if it is in another warehouse, they **stop there first to pick it up** (a pick-up stop); otherwise, the work order is left uncovered with the list of what is missing. When the plan is applied, the lines are switched to come from the chosen warehouse. Tools do not count.

12

Time budget (effective duration)

AVAILABILITY	Core
SURFACES	Web (work order creation, maintenance type and task records) · used by the planner, dispatch, capacity and the assistants
WHAT IT IS FOR	Ensures that the estimated duration of a work order is not a made-up number. It is the piece that turns the real execution history into an estimate that improves on its own with use.

12.1 Resolution hierarchy

A single source of truth for “how long is this going to take”, with this priority:

1. **Duration set manually on the work order:** if it is filled in, it always prevails.
2. **If the work order has tasks,** each task contributes its explicit duration or, if it has none, the **real historical median** for that task, provided there are at least 5 samples. If all of them resolve, they are added together; if any does not, the greater of that sum and the type’s duration is taken (**floor rule:** never estimate below what the type says).
3. **Without tasks,** the maintenance type’s duration: its **historical median** if there are at least 5 samples; otherwise, the configured one; and as a last resort, 30 minutes.

The creation form always states **where the figure comes from**.

12.2 What prevails over what

The history **prevails** over the configured estimate for the **type** — real data is worth more than the initial configuration — but it **never** prevails over the duration set on the work order or over a task’s explicit duration. A planner who knows that this particular work order is going to take four hours can say so, and the system does not correct them.

12.3 Who uses it

- The **planner**, as the service time for each stop, and **dispatch**, to measure the impact on a working day.
- **Staff capacity** (20.11).
- **Work order creation**, which fills in the duration automatically and proposes the **planned end**. The **deadline** is not derived from the duration.
- The planning assistants, which require the duration to **fit entirely within one slot** of the schedule.

- The **PDF report**, alongside the actual duration, and the type and task records, with their historical median and the number of samples.

The estimated duration is a **budget**, not the labour: each person's cost comes from their timesheet (5.19).

12.4 Where the history comes from

The median is calculated on completed work orders, using the first available source:

PRIORITY	SOURCE	CONDITIONS
1	Time in progress (sum of segments, 5.18)	Between 1 minute and 72 hours, with no inconsistent segments
2	Geofence : time inside the zone	Stays of between 1 minute and 12 hours
3	Actual dates	Only without segments; between 1 minute and 24 hours

Time in progress excludes pauses and nights: a work order started one day and finished the next does not contribute a full day of work. A task's median is fed only by work orders that have **that single task**. The history is updated every 15 minutes.

13

Vehicle fleet

AVAILABILITY

Fleet and costs module

SURFACES

Web · Mobile (vehicle check-in)

WHAT IT IS FOR

Attributing the real travel cost to each work order, separating work kilometres from outward and return kilometres, and knowing at all times who is driving what. Without fleet, a work order's cost covers materials and hours; with fleet, it also covers kilometres and fuel at real prices.

13.1 Vehicle inventory

Record with:

- **Registration number, type** (van, car, light lorry or motorcycle), make, model and year.
- **Fuel:** diesel, petrol, electric, hybrid, LPG or CNG.
- **Consumption** in l/100 km or kWh/100 km, mandatory according to the fuel.
- **Tank capacity and odometer**, which updates automatically with each trip.
- **Branch.**
- **Active/inactive status.** A vehicle is not deleted: it is deactivated, and it cannot be deactivated while it has assignments in progress.

A warehouse **van** (3.7) is linked to a vehicle: the materials travel with the vehicle, whoever is driving it.

13.2 Fuel prices

Two complementary sources:

a) The organisation's own prices, with a validity period, in its currency. They reflect the real price on the fleet card or under the supply contract. A new price closes the previous one, and a one-off price within a longer one splits it in two without losing it.

b) Automatic official index. Every **30 minutes** the system checks the fuel prices published by the **Spanish Ministry for the Ecological Transition** and stores the average price per fuel type. It applies only to organisations that work in **euros**; with another currency, or for electric vehicles, an own price is needed.

13.3 Technician ↔ vehicle assignment

The record of which person has which vehicle and since when, with a validity period, method, confidence level, status and review.

METHOD	CONFIDENCE
Check-in by scanning the vehicle's QR code (mobile)	0.99
Check-in by choosing the vehicle from a list (mobile)	0.95
Assignment by the supervisor from the web	1.00

Overlap rules:

- **One vehicle, one technician at a time**, so that the same trip is not counted twice. At check-in, whoever scans the vehicle takes it.
- **One assignment per technician at a time**: the new one cuts short the previous one, and a one-off assignment within a longer one splits it without losing it.

13.4 Review of doubtful assignments

A low-confidence assignment generates a **vehicle pending review** alert, with its own list and **confirm** or **correct** actions. It exists only in organisations with the fleet module. The **Alerts** tab in Vehicles brings together these alerts and the auto-pause ones (14.6).

13.5 Route legs and true cost

Each relevant journey is stored as a **route leg** with:

- Origin and destination (base, home, work order...) and their time.
- **Real road distance and duration**, with the **route trace**.
- **Vehicle** in use at that time.
- **Cost snapshot**: the fuel price that day, the vehicle's consumption and the **cost** of the leg. The cost at the time is stored: recalculating months later with today's price would distort the accounts.
- **Category**: work, travel from base, return to base or administrative.
- **Work order** to which it is charged, only for work legs.

How legs are created. Each time the geofence confirms a technician's arrival in a work order's zone (chapter 15):

- If they had already left another work order on the same working day, the **work** leg between the two is generated, charged to the destination work order.
- If it is the **first arrival of the working day**, the leg starts from their **base** — branch or home, according to their record — as **travel**, without being charged to any work order.
- The **return to base** for each completed working day is calculated automatically. If the working day ended with the technician inside a zone (phone switched off), no return is invented.

The working day is the shift's, not the calendar day: a night shift does not generate at midnight a return and a departure that nobody made. A base more than 150 km from the work order does not generate a leg, and neither do arrivals detected with **simulated positions** (14.10): they do not prove the journey.

Where the cost is shown.

- The fuel for **work** legs makes up the **vehicle** line of the work order's true cost (5.10), with a breakdown by vehicle and the journey on the map.
- **Travel and return** kilometres are not charged to any work order: the cost report shows them separately (20.10).

13.6 Viewing route legs: Journeys tab

The **Journeys** tab in Vehicles shows, for a vehicle and a period, the number of trips, the **kilometres split into work and travel**, the time behind the wheel and the fuel, the map with the traces and the table of trips. A person's trips can also be seen from their GPS tracking, and a work order's, from the work order. A leg can be created manually. The tab warns of technicians without a base with a location, whose first trip and return cannot be counted.

13.7 Fleet on the dashboard

The "**Fleet and tracking**" section of the dashboard (20.12), for periods of up to one year: trips, kilometres (and how many are not work), driving hours and fuel, by person, vehicle or month, with legs without a cost counted separately so that the total does not look complete when it is not.

14

Telemetry and GPS tracking

AVAILABILITY	Real-time field module
SURFACES	Mobile (sending) · Web (viewing, map and GPS Tracking screen)
WHAT IT IS FOR	Knowing where the field team is in order to dispatch emergencies, recording presence at the work site and calculating real journeys — with a design built around the legal obligations on workplace geolocation.

Seeing other people's position is a permission in its own right, separate from seeing staff: by default the supervisor and the administrator have it.

14.1 Tracking sessions: the legal anchor

Tracking **is not continuous**: it is organised in **sessions** with an explicit start and end. Each session stores the time and position of the start and end, the **closing reason** (the technician closes their working day, automatic end of working day, inactivity, withdrawal of consent or a time slot recovered after being without coverage), the **version of the consent** under which it was captured and the device.

This structure is what gives effect to compliance with article 90 of the LOPDGDD (geolocation in the workplace): there is a versioned and recorded consent, a bounded start and end, and not an indefinite collection.

14.2 Receiving positions with working-hours and consent validation

The app accumulates positions and sends them in batches. The system:

1. Sorts them by their capture time.
2. Discards one by one those that cannot be stored (too old, in the future, impossible), without rejecting the whole batch.
3. **Checks consent** against the organisation's record: if the technician has not accepted the current notice, or has withdrawn it, nothing is stored and the app knows whether it should request acceptance.
4. Checks for each position whether it falls **within that person's working hours**, with 15 minutes' grace, taking into account shifts, shift changes, days off, on-call duty, overtime and night shifts. Positions outside working hours are **rejected**.
5. **Discards resends**, so that a repeated batch does not duplicate anything.
6. Passes the new positions through the geofence (chapter 15).

Each accepted position stores: location, **accuracy**, speed, heading, **battery**, whether the device is moving, whether the position was **simulated** (14.10) and that person's **active vehicle** at that moment.

14.3 Retention

Positions are kept for **60 days** and are deleted automatically. Other location traces also expire: after 60 days the start and end positions of sessions, the trace of route legs and the position of work order events are cleared, while dates, kilometres, duration and cost are kept. The log of access to location data is kept for two years.

14.4 Queries

- **Current positions:** the last position of each technician from the **last 2 hours**, with how long ago it arrived, battery, speed, vehicle, whether it was simulated and the **work order they have in progress**, linked on the map. This is the “GPS positions” layer.
- **My route, my sessions and my export:** each technician views and downloads their own history from the app.
- **A person's route** for a date range, together with their fleet trips.
- **Export** of a person's location data in JSON or CSV.

When a route or an export reaches its limit, this is stated explicitly and the point from which to continue is indicated.

Every access to a person's location data is logged: who, about whom, what action, what range and how many points, in compliance with the principle of integrity and confidentiality (art. 5.1.f GDPR).

14.5 LOPDGDD rights

RIGHT	HOW IT IS HANDLED
Access (art. 15)	The technician views their route, their sessions and their consent history, and exports their location data themselves
Erase (art. 17)	Deletion of a person's location data, which anonymises what has operational and accounting value (legs, events) without revealing where they went. It is irreversible and is logged
Withdrawal of consent	The technician exercises it from the app: it stops capture immediately, without deleting what has already been collected

14.6 Automatic end-of-day closure

Every 15 minutes the system checks, in the time zone of each technician with open tracking, whether their working day has ended (with its margin). It never closes while the technician is within their working hours. If it has ended:

1. It **closes the tracking session**.
2. It **pauses** each work order in progress on which **nobody is still working**: if two technicians share a work order and one finishes at 15:00 and the other at 19:00, the work order remains in progress

until the second one finishes.

3. It issues a “**Work order paused at end of working day**” alert so that the resumption can be confirmed the following day.

It solves a real and very common problem: work orders that stay “in progress” for days because nobody remembered to close them. The automatic pause does not stop the SLA clock.

14.7 Off-shift activity and regularisation

Positions rejected for falling outside working hours **are not stored** — outside the working day there is no legal basis for storing them — but the system notes, per technician and day, **how many there were and between what times**, without coordinates.

The **Off-shift activity** tab in GPS Tracking lists those time slots, to detect on-call duty or overtime that nobody declared. The manager creates the **calendar exception** that covers it, and the time slot is regularised; the exception enables tracking from then on.

14.8 Geolocation notice and consent

The data controller is the organisation, so **the text of the notice is published by the organisation**:

- The **Geolocation notice** tab publishes versioned wordings. Earlier ones **are not deleted**, because the data collected under them needs their text to prove what was communicated.
- With no version published, **nothing is captured**.
- The technician **accepts** the notice from the app (version, date and device are recorded), can **withdraw** it and can view their history.

14.9 Schedule windows for the app

The app knows the technician’s own **schedule windows** for the coming days, using the same criteria as the server, so as not to waste battery or data on positions that are going to be rejected.

14.10 Simulated positions

The phone’s operating system flags positions that come from a fake location application. They are stored flagged (or, if the organisation prefers, discarded), also in photos, signatures, incidents, quick reports and form GPS fields, and the supervisor receives a “**Simulated positions**” alert stating how many and when, never where. An arrival detected with simulated positions **does not generate kilometres** (13.5). It detects the usual case; it is not proof: a tampered phone can hide it.

14.11 Capture with the screen off

The app can capture and send the position with the screen off. Those batches go through exactly the same checks for working hours, consent, duplicates and geofence, and a bad position never causes the rest of the batch to be lost.

14.12 GPS Tracking screen

The module's web screen, with three tabs: **Off-shift activity** (14.7), **Geofence review** (15.9) and **Geolocation notice** (14.8).

15

Geofence and field presence recording

AVAILABILITY	Real-time field module; the cost of legs also requires the fleet module
SURFACES	Mobile (detection and prompt to the technician) · Web (work order activity, review and configuration)
WHAT IT IS FOR	Ensures that the technician's arrival and departure are recorded with geolocated evidence, that the app asks them at the right moment whether they are starting or finishing the work, and that the journey is calculated with its cost. The system never changes a work order's status because of the geofence: status changes are always decided by the technician.

15.1 Arrival and departure detection

Each position is compared with the **zones of the open work orders assigned to that person**: their assets and, if it has one, its own location (5.14). The distance is measured to the geometry, not to its centre: inside a polygon it is zero.

The **arrival radius is fixed: 200 m**. Departure is counted at **300 m**: between the two radii nothing is confirmed, so that a technician on the edge does not generate looping entries and exits.

RULE	DEFAULT
A position with more error than this does not count	50 m
Arrival: one stationary position inside, or at least 2 consecutive positions inside over 5 minutes	2 positions / 5 min
Departure: at least 2 consecutive positions outside over 10 minutes	2 positions / 10 min

That is why **driving past without stopping is not arriving**. The crossing is dated at the **first position** of the run, which is when it was really crossed, and a resend does not generate crossings twice. Positions that arrive very late record their crossings, but do not generate a leg or an alert.

15.2 Geolocated work order events

Each relevant detection generates a **work order event** with its type, **what triggered it**, time, position, **confidence level** and context (asset reached, distance, accuracy). It is a record that complements the status history: the history says *what* changed, the event says *where and why*. The event's position expires after 60 days (14.3).

15.3 Starting on arrival: the technician decides

The system does not start, resume, pause or close any work order because of the geofence.

- **The app detects and asks.** On arriving at a pending work order it asks the technician whether to **start** it (or whether to **resume** a paused one); on leaving with the work order in progress, whether **they have finished** it. It does not offer this if they already have another work order in progress, nor for work orders planned for another day. The zones travel to the app so that it works offline.
- **The system records and alerts.** It records arrivals and departures (proof of presence, durations, legs) and, if the technician **leaves with the work order in progress**, alerts the supervisor with “**Technician out of zone with WO in progress**”, unless another technician on the work order is still inside. The alert closes automatically when they return or the work order is no longer in progress.

When the technician accepts the prompt, the status change is marked as “**Started on arrival**” or “**Resumed on arrival**”, with the phone’s position. Nobody closes a work order automatically: closure involves dates, forms, photos and a signature that only the technician can provide.

15.4 Automatic route leg

With the fleet module, each confirmed arrival generates the **route leg** to the work order, with real road distance and duration, trace and cost: from the last departure of the working day, as work charged to that work order, or from the **base** if it is the first arrival of the day, as travel (13.5). If the calculation fails, the positions are stored anyway.

15.5 Alerts that resolve themselves

When a work order changes status, the alerts whose cause has disappeared close automatically: the **paused at end of working day** alert when it is resumed or finished, and the **technician out of zone** alert when it is no longer in progress. If a closure from the app is undone (for lack of a signature, form or acceptance), the alerts return to their previous state.

15.6 Detection of stalled work orders

A daily task marks as “**needs review**”, without changing their status, work orders that have been paused for more days than configured (3 by default). The decision remains a human one.

15.7 Work orders in progress without positions

Without positions, a departure cannot be confirmed. Every 10 minutes the system looks for work orders in progress whose technician was inside the zone and has gone more than **45 minutes** (configurable) without sending positions — phone switched off, battery flat or app closed — and marks them for review without touching their status.

15.8 Organisation settings

In **Settings** → **Geofence**, each setting with its explanation: whether or not to alert on leaving with the work order in progress, maximum accuracy, positions and minutes to confirm arrival and departure, speed below which a technician is considered stationary, departure radius margin, minutes without

positions before marking a work order, and whether simulated positions are discarded. The arrival radius (200 m) is not configurable.

15.9 Review on the web

- **Work order Activity tab:** a timeline with the status history and its origin, arrivals and departures, review alerts, and planned and actual dates. With the tracking permission, the distance and accuracy of each crossing are also shown and “**View on map**” draws the surrounding positions; that access is logged.
- **Geofence review** in GPS Tracking: technicians out of zone with a work order in progress, work orders in progress without positions, phones with simulated positions, and possible unrecorded visits pending confirmation or dismissal.

16

Inventory, warehouses and purchasing

AVAILABILITY	Maintenance module, with the warehouse capability; the cost of an order also requires the costs capability
SURFACES	Web · Mobile (consumption on the order and “My van”)
WHAT IT IS FOR	Controlling material by warehouse—including each technician’s van—, knowing how much is already committed by open orders, valuing it correctly and knowing how much material has gone into each asset and each order.

16.1 Categories and items

- **Hierarchical categories** with colour and icon.
- **Items** with internal code, name, description, manufacturer and reference, unit of measure, **cost price** (the average price, to 4 decimal places), **last purchase price**, sale price, tax, images, **preferred supplier**, whether it is a **tool**—tools never move stock— and which **asset type** it corresponds to.

The **catalogue shows all items**, including those with no stock, with their total stock, the **committed** quantity and the **free** quantity (16.2). An item used by an open order cannot be deleted; where deletion is allowed, its stock is adjusted to zero with its own movement and the item goes to a recycle bin that can still be consulted.

16.2 Stock by warehouse and committed stock

Stock **does not live on the item**, but in each warehouse: quantity, **minimum stock specific to that warehouse**, location within the warehouse (aisle, shelf) and date of the last stock count. This way, the same reference can have a minimum of 50 in the central warehouse and a minimum of 2 in each van.

Committed and free. The **committed** quantity of an item in a warehouse is what open orders still have to take out of it; **free** is stock minus committed. It is shown in the catalogue, on the item record and in each warehouse's stock, and you can check which orders are committing it.

There are three types of warehouse (3.7): **physical**, **mobile** (the van, linked to a technician or to a vehicle) and **virtual**. When the van is linked to a vehicle, the material goes **with the vehicle**: whoever is driving it that day takes it out.

A warehouse can only be deactivated when it has no stock, and an inactive warehouse does not accept receipts or transfers.

16.3 Stock movements: a single path

Every change in stock levels goes through **one single path**, which updates the quantity and records the movement in the same step: stock cannot change without leaving a trace, whether the change comes from the web, the app or a scheduled process. Movements cannot be edited or deleted, and two simultaneous operations do not overwrite each other.

Each movement records type, quantity, **resulting stock**, warehouse, related warehouse (in transfers), **unit cost**, supplier, work order, notes and author.

MOVEMENT	WHEN IT OCCURS
Receipt	Purchase or receipt of material
Consumption	Material consumed on a work order
Return	Material returned from a work order
Issue	Direct issue, without a work order
Adjustment	Inventory correction
Transfer	Movement between warehouses of the same organisation

16.4 Valuation: weighted average price

When stock is received with a purchase price, the system **recalculates the item's weighted average price**, to 4 decimal places. The cost of the receipt stores the real purchase price.

Consumption on an order is valued at the average price in force at the time of consumption, which is **fixed on the movement**: the cost of an order that has already been closed does not change even if the average price moves later. A return comes back at the same cost at which it went out.

16.5 Operations

- **Add stock** with purchase price and supplier.
- **Adjust stock** to a specific quantity (physical count), with a mandatory reason.
- **Direct issue** with a reason: internal consumption, loss, breakage, expiry or other.
- **Transfer** between warehouses.
- **Consumption and return** from the work order (16.6).
- **Movement history** of an item, filterable by warehouse.
- **Minimum stock** of an item in each warehouse.

Operations warn on screen if the warehouse falls **below minimum or into negative stock**.

16.6 Work order material

Material is **deducted when it is consumed, not when it is planned**, and only the usage that is reported moves stock:

MOMENT	WHAT HAPPENS TO STOCK
The office plans material	It does not move. The line fixes its warehouse and its quantity becomes committed
The technician reports what was used	Exactly what was used is consumed; if the figure is corrected downwards, the difference is returned
The order is completed without anyone reporting a line	Nothing moves and the line stops committing stock
The order is cancelled or a line is removed	What was consumed and not used is returned

Each line moves only the difference from what it has already taken out: resending the same data from the app does not deduct twice.

Which warehouse it comes from. On the web, each line has a “**Taken from**” selector: **Automatic** or a specific warehouse with the free quantity it holds. In automatic mode the system decides, in this order: the van of the vehicle the technician is driving that day, their van, the warehouse on their record, the branch's warehouse and the main branch's warehouse. A manual choice is kept. The technician does not choose a warehouse: whatever they add in the field comes out of their own.

Negative stock allowed. A consumption is never rejected for lack of stock: the warehouse may go into negative stock and an alert is raised (16.7).

Warning when planning. When an order with material is saved, the system warns, per warehouse, of anything that exceeds the stock or leaves it below minimum, and opens a “**Material missing for a WO**” alert for the person in charge of the warehouse, which is updated when the order is saved again and closed when the order is completed or cancelled.

Preventive maintenance, planner and dispatch. Preventive orders are created with the warehouse of each line already fixed and their material committed (6.3). Depending on the organisation's material setting, the optimiser and dispatch warn of what is missing or only give the order to someone who can carry it (11.11).

In the app. The technician reports what was used per line and can add unplanned material. Each line shows which warehouse it comes from and how much has been consumed, and “**My van**” shows, in view-only mode, the stock in their warehouse with its minimum, the committed quantity and the tools.

16.7 Stock alerts

Stock alerts are opened, updated and closed **automatically with each movement**, one per **material and warehouse**:

- **Low stock:** the quantity in that warehouse has fallen below its minimum.
- **Out of stock:** the quantity has gone negative.

The alert names the material and the warehouse and states what is left; if it is replenished, **it closes on its own**. Added to these is “**Material missing for a WO**” (16.6). They are received by the administrator, the supervisor and the warehouse role, limited to the warehouse's branch.

Query of **items below minimum** by warehouse, sorted by largest shortfall, with the preferred supplier and the last purchase price, exportable to CSV.

16.8 Suppliers

Supplier record with active/inactive status and **item catalogue per supplier** with its price. It allows a supplier to be marked as preferred for an item and records who each receipt came from.

16.9 Inventory reports

- **Inventory value**, by category and by warehouse, with a ranking of items by value.
- **Consumption by asset** over a date range: how much material has gone into each asset, exportable to CSV. It is the basis of the “expensive assets” analysis.
- **Cost of a work order**: materials, labour and vehicle (5.10).

16.10 What is not available yet

So as not to overpromise, the following are outside this phase:

- **Reservations** of units and **material in transit**: planning commits stock (16.2), but does not set units aside.
- **Purchase orders** with formal receipt and **transfer or purchase requests**: replenishment is recorded as a stock receipt with its price and its supplier.
- Control by **serial number** or **lot**.
- **Import** of the item catalogue, count sheet with bulk adjustment, material kits and QR codes per item.

17

Meter reading

AVAILABILITY

Meter reading module, with its dynamic readings and reading routes options (routes require working with teams)

SURFACES

Web · Mobile (fully offline) · Public API

WHAT IT IS FOR

Field meter reading campaigns (water and energy meters, running hours, pressures) with period-based planning, work without coverage and detection of out-of-range readings with an automatic incident.

17.1 Measurement points

They are defined at two levels:

a) By asset type — the template: name, unit, minimum and maximum valid values, **range basis** (17.9), reading frequency, tolerance days and the “**Open an incident if the reading is out of range**” checkbox. On top of the template, **conditional rules** are defined: “if the property *calibre* has the value *DN80*, then the maximum valid value is different and so is the frequency”, with priority between rules, which are applied to existing assets. Measurement configuration is a permission in its own right: a technician who takes readings does not change it.

b) By specific asset — the real instance, which inherits from the template and can be **overridden** individually. It also stores the **next scheduled reading date**, which moves forward on its own with each valid reading.

Deactivating a template deactivates its points (except customised ones), and a template with readings cannot be deleted: it is deactivated.

17.2 Recalculable thresholds

The validity thresholds of a point can be **recalculated statistically** from the real history, instead of being set by hand. The point stores which method was used to calculate them, with how many samples and when. A daily task recalculates them in bulk; for cumulative meters, on the consumption series.

17.3 Static reading routes

Fixed, ordered sequences of assets that are always read in the same order: the meter reader's round. Each route has a name, a **customer**, an assigned **team** and a **schedule** that generates its batches.

Visit order. The form shows the numbered, draggable list and the line of the round on the map. “**Optimise order**” calculates the shortest round:

- It groups the assets by **street segments** and covers each segment in a row, so that in a street with five manholes no routes are calculated between them.
- It orders the segments **on foot** or **by vehicle** (respecting one-way streets), with real distances and times. On foot, you choose whether to cover **both pavements at once** or one side out and the other side back. The starting point can be the first asset, a branch or a point, with or without return.
- The result states how many segments it found, and the total distance and time.

Editing. A change to the route can also be applied to the batch in progress, as long as it has no readings.

17.4 Reading batches (campaigns)

The batch is the unit of work. It contains name and type, origin (a route or a selection by asset type and sector), **assigned team**, customer and branch, **period** and frequency, and the **status** of the batch and of each asset and point, to know exactly what is outstanding.

A batch is **Pending** when generated, **In progress** as soon as its first reading arrives and **Closed** when it is closed; a closed batch cannot be reopened.

The **batch summary** counts pending, completed, in progress and with incident, and lists the readings that need attention: invalid, **out of range** (with the incident they opened), with notes or with photos.

Operations: creation, due date assignment, batch **calendar**, **closure** and deletion. An asset added after the batch has been generated joins the open batches that apply to it. **Closing** reports the points that were left unread, and from that moment the batch accepts no more readings, not even those arriving from the app without a connection. Only a batch that nobody has worked on can be deleted.

17.5 Automatic batch scheduling

Rules that generate batches on their own:

- Measurable asset type or source route, and customer.
- **Frequency** (days, weeks, months or years) and **days in advance**. Routes also allow a **one-off** run.
- **Assigned team**, with the option of a **different team per sector**.
- Start and end dates.
- Manual run ("Run now") in addition to the automatic one.

Periods are counted from the start date, so that a monthly schedule on the 31st does not get stuck on the 28th. Only assets with at least one active measurement point are included, and when the schedule is created the system warns of sectors with no team and of assets with no sector that would be left out.

17.6 Offline work per batch

Each batch has an **offline package** with everything needed to work on it without coverage: assets, measurement points and history. The app downloads it in one go and works entirely locally.

- It is prepared as soon as the batch is created, and **regenerated** before download if there have been changes.
- It is built separately: a download never picks up a half-finished package.

- When the batch is **closed**, the package disappears.

Combined with the **offline mapping** of assets and the base map (chapter 21), the technician can do a full day of readings without a connection.

17.7 Readings

Each reading records value, date, **measurement point**, asset, batch, **person who took it**, **photograph** (mandatory if the type requires it), notes, whether it is valid and whether it is **out of range** (17.9). Invalid readings do not count as the last reading and do not move the next scheduled date forward. Readings are also available through the public API (chapter 25).

17.9 Out-of-range readings and automatic incident

When each reading is received, the system compares it with the valid range of its point. Each point declares its **range basis**:

- **Consumption** (cumulative meter): the reading **minus the previous valid one** is compared. This is the default for new points.
- **Reading** (instantaneous measurement): the value is compared as it is.

The range already takes into account the template, the applicable rule, the recalculated threshold or the manual setting. The web shows “Consumption outside the valid range” or “Reading outside the valid range”.

The incident. If the reading is out of range and the point has “**Open an incident if the reading is out of range**” ticked, a “**Possible faulty meter / Leak**” incident is opened, with **high** criticality and origin *Meter reading*, on the asset of the reading, with the value, the previous reading, the range and the photo, and with the new incident alert. At most one per reading. If the reading is corrected later, the incident **does not close on its own**: a note is added to it and a person decides.

Lists. The lists (20.6) offer “**Out-of-range readings**” and “**Overdue readings**” (points whose scheduled date plus tolerance has already passed).

18

Dynamic forms

AVAILABILITY	Maintenance module, with the forms capability
SURFACES	Web (design and viewing) · Mobile (completion)
WHAT IT IS FOR	Letting each organisation capture in the field exactly the data its process requires —an inspection checklist, a commissioning record, a condition sheet — without any programming, and ensuring that what is required cannot be skipped when the work is closed.

18.1 Templates

Template with name, description, **version** and active status, on two levels: **global** (belonging to the vertical) and **organisation-specific** (cloned and editable). A template can always be deactivated: it is no longer offered and whatever has already been filled in is kept. A template in use cannot be deleted.

18.2 Sections

Each template is organised into sections with name, description and order, **collapsed by default** if desired, and **repeatable**: a detail section that is filled in N times (for example, one row per luminaire inspected), with a minimum number of rows. A section with responses cannot be deleted: it is **retired**, it no longer appears in new forms and whatever has already been filled in is kept.

18.3 Fields

Fourteen field types:

Yes/no · Text · Integer · Decimal number · Single choice · Multiple choice · Radio buttons · Checkboxes · Date · Time · Photograph · Signature · GPS position · Information text

Each field defines:

- **Multilingual** label and help text.
- **Options** (for the choice types).
- **Mandatory flag** and **validation rules** (minimum, maximum, pattern, “must be ticked”).
- **Dependency rules**: show, hide, make mandatory or optional, enable or disable the field depending on the answers to others.
- **Responsive layout**: columns it occupies on mobile, tablet and desktop, so that templates do not have to be duplicated per device.

What has already been answered is protected. A field with responses cannot be deleted (it is retired), does not change type and does not lose an option that has already been answered; labels can be changed.

There is a **visual layout editor** and a **preview** of the form as the technician will see it.

18.4 Form assignment

A form is linked to the work through assignments that combine **maintenance type**, **task** and **asset type**. Each assignment declares whether it is **mandatory**, its **order** and its **scope**: once per order or once per asset.

When an order is created or edited, the system attaches to it the forms chosen **plus the mandatory ones** that apply to it, and each one stores the task it belongs to and whether it is mandatory **on that order**: an open order is not blocked by a new rule that its technician has not seen.

Editing the order keeps the responses: new mandatory forms are added and only deselected forms that have no responses are removed.

18.5 Responses

- **Response batch:** groups several responses on an order under a name (“Correct batch”, “Fault group”), to process in one go assets that are all in the same condition instead of repeating the same form N times.
- **Response:** linked to the order, the form, the asset and the batch, with who filled it in, when, its status (draft or completed) and the **GPS location**.
- **Values:** stored with their type, with a row index for repeatable sections.

Typed values make it possible to analyse responses, not just display them. Responses are shown grouped by batch and by section, in the designer's order, with photos and signatures; the same data feeds the order's PDF report.

18.6 Tasks completed by their forms

A task on the order becomes **completed** when all the forms linked to it have a completed response, and reverts to pending if any of them stops being completed. Tasks without forms keep their status.

18.7 Completion from the mobile app

- Responses are validated record by record when they are uploaded, and authorship comes from the session, not from the content sent.
- An order that has been **validated, closed or signed** no longer accepts changes to its responses.
- Photo and signature fields only accept images.
- GPS fields store whether the position was simulated (14.10).
- Synchronisation sends the app the replaced responses and the retired fields so that it removes them.

18.8 Mandatory forms and order closure

A mandatory form **blocks closure**: the order cannot be completed, validated or closed while any is missing, and the system says which ones. Cancelling is not blocked, and quick reports are exempt.

What counts as completed: a completed response if the form is per order, or one for each asset it applies to if it is per asset. A response marked as completed **with empty mandatory fields does not count**, taking into account dependency rules and checking row by row in repeatable sections.

This is checked on the web, in bulk actions and in the app; if a closure from the app does not meet it, the status change is undone and the rest of the submission is saved.

19

SLA

AVAILABILITY	SLA module; additionally, it can be enabled customer by customer
SURFACES	Web · Mobile (cause of the pause) · Customer portal (deadlines and compliance)
WHAT IT IS FOR	Committing to deadlines in writing and measuring them automatically. It is the feature that turns a maintenance contract into something both parties can verify.

19.1 Criticality levels

LEVEL	MEANING
Low	Minimal impact, does not affect operations
Medium	Partially affects operations
High	Significant impact
Critical	Total stoppage or severe operational risk

19.2 Global SLA and organisation SLA

For each **criticality × type** combination a **maximum response time** and a **maximum resolution time** are defined. The type distinguishes **incidents** from **work orders**, because the commitment to attend to a report and the commitment to carry out a job are not the same.

Deadlines exist on two levels: the platform's **global SLA** and the **organisation SLA**, which inherits each value from the global one until it is customised, with the option to **restore the global value**.

19.3 Customer-specific SLA

Each customer can have **its own SLA**, with specific times by criticality and type that replace those of the organisation. The order of precedence is **customer → organisation → global**. The customer record is also where you enable whether it **accepts justified pauses** (19.8).

19.4 Application

Work orders. Each order receives its **response and resolution deadlines**, and stores which level they come from:

- Criticality comes from the order's **priority**, according to a configurable mapping (by default, Urgent → Critical, High → High, Medium → Medium, Low → Low).
- **When the clock starts:** for a corrective order, from its creation; for a **preventive** order, from its **scheduled date**, so that it does not reach the scheduled day with half its deadline used up. **Moving an order in the calendar does not move its deadline.**
- Deadlines are counted **in working hours** (19.9). An order created outside working hours starts counting at the next opening time.
- The resolution deadline is **extended by justified pauses** (19.8).
- On open orders they are recalculated when the SLA, the schedule, public holidays or the time zone change. Closed orders keep the deadline they were closed with.

Incidents. They receive their deadlines in the same way, based on their criticality and always from when they were created (7.5). A resolved, closed or rejected incident keeps its deadline.

19.5 “SLA at risk”

An open order becomes “**SLA at risk**” when it has used up a **percentage of its resolution deadline** —75 % by default, configurable—, counted in the same working time. It is proportional to the deadline: a 4-hour order raises the warning earlier than a 5-day one. Incidents use the same percentage.

It is used by the dashboard indicator, the pulse of the day, the label on the cards and the **SLA at risk with no staff** alert (19.10).

19.6 Configuration

The whole SLA is configured in a **single section**, with help on each concept and a label that distinguishes what *changes the deadlines* or *changes the alerts* from what *only changes colours*:

1. **Deadlines by criticality**, for incidents and for orders, with their origin (inherited or customised).
2. **Which deadline each order priority uses.**
3. **“SLA at risk” warning**, with examples calculated on the configured deadlines.
4. **How compliance is rated on the dashboard:** the targets that decide the colour of the indicators, without moving any deadline.
5. **Justified pauses:** what they are for and the **name the organisation sees** for each cause.

19.7 Visibility and control

- **Cards and records.** Orders and incidents show the status of each deadline —met, missed, on time, at risk or overdue— with its due date and, on orders, the justified pauses and the minutes deducted.
- **Dashboard:** on-time resolution and response for orders and incidents, orders at risk and with the SLA missed, and the **resolution and MTTR by criticality** table with mean, median and 90th percentile (20.3).
- **“Where the time goes until closure”:** a chart that splits the time of resolved incidents between waiting until work starts, repair, pauses by cause, other waits and administrative closure.
- **Customer portal:** each incident and each order carries the status of its deadlines, with the pauses that stop the clock, their cause and their duration; the portal dashboard gives the **service level** for

the period and its monthly trend (22.5).

19.8 Justified pauses

The SLA clock does not stop just because an order is paused: an order can also be paused because of the end of the working day or a break. What decides is the **cause** of the pause. There are four fixed causes, because what matters for the contract is who is responsible:

CAUSE	EXAMPLES	DOES IT STOP THE CLOCK?
Attributable to the customer	No access, customer absent, change of scope	Yes
Third party or supply	Manufacturer's spare part, permit from another authority, utility company	Yes
Force majeure	Weather, disaster, order from the authorities	Yes
Own	Break, another order, end of the working day	No

Each organisation can **rename** each cause as it sees it, but never add causes.

- **Only if the customer accepts it** under the contract.
- **Resolution only:** the response deadline does not stop.
- **In working hours:** the stopped period is deducted in working minutes.
- **Who chooses it:** whoever pauses, on the web or in the app. The cause of a recorded pause can be corrected, and this is kept in the audit trail.
- **Control through transparency,** without approval: the record shows the periods, the customer sees in their portal all the pauses that stop the clock with their cause and duration, and the export and the API include them.

Justified pauses only apply to work orders.

19.9 Working hours used in the calculation

The deadline advances through the **slots of the weekly schedule** (9.8): split shifts without counting lunch, **night** slots, a **branch-specific** schedule and excluding public holidays. Slots are counted in the branch's time zone, so that the night when the clocks change lasts as long as it actually lasts.

19.10 SLA alerts

The system maintains, minute by minute, two alerts on open orders:

- **“SLA at risk with no staff”:** the order is at risk, is not yet overdue and has nobody assigned. It closes on its own when staff are assigned, when the order is finished or when it becomes overdue.
- **“SLA overdue”:** the order has exceeded its resolution deadline and is still open. It closes when the order is finished or if the deadline moves.

Both are included in the default escalation rule and can be sent by email (chapter 24).

20

Dashboard and analytics

AVAILABILITY

Core, always active; the scope depends on the dashboard's level of detail; costs require the costs capability, and fleet its own module

SURFACES

Web · Customer portal (reduced version, chapter 22)

WHAT IT IS FOR

To answer maintenance management questions with data, using the indicators the sector recognises (MTBF, MTTR, availability, preventive compliance, backlog, deadline compliance) calculated on the real history of statuses and work periods, not on fields declared by hand; and, from the home page, to direct the day's work on the map.

20.1 Main KPIs

KPI	DEFINITION
Active WOs	Orders that are pending, in progress or paused
Open incidents	Pending or approved reports, not yet resolved
Overdue WOs	Unfinished orders past their due date
On-time resolution	% of orders finished in the period before their resolution deadline or, without SLA, their due date
Total assets	Size of the managed asset base
Items below minimum	References below the minimum in some warehouse
Completed WOs	Work closed in the period
Average resolution time	Average days between the real start and end
Pending WOs	Backlog not yet executed
Backlog in weeks	Pending backlog divided by the real weekly closure rate: "how many weeks of accumulated work do I have?"
Preventive compliance	% of preventive maintenance due in the period that was completed before its due date
Availability	% of time the asset base has been operational, deducting downtime due to breakdowns
SLA at risk	Active orders that have already used up the configured percentage of their deadline. Snapshot of now

KPI	DEFINITION
SLA missed	Active orders that have already passed their deadline. Snapshot of now
On-time response	Of the closed orders with a response deadline, % that were started on time
First-time fix	% of closed orders done in a single work period and never reopened
Reopenings	% of closed orders that returned to an active status
Wrench time	% of the interval between the real start and end during which the orders were in progress. The rest is pauses and waits: parts, access, travel, nights
Incidents: on-time response	% of incidents in the period that left pending status before their deadline
Incidents: on-time resolution	% of incidents in the period resolved before their deadline

SLA indicators only appear in organisations with SLA. Dates and days are counted in local time.

Comparison with the previous period. With a date range, flow metrics are compared with the previous period of the same length and with the same filters.

Targets and traffic lights. Fifteen indicators are coloured green, amber or red against two thresholds configured by the organisation, each with a help text explaining what it measures.

Basic level. MTBF, backlog in weeks or first-time fix are excellent indicators, and they are also what makes the maintenance manager of a small company conclude that the product is in a different league. With the dashboard at **basic level**, **seven figures that anyone understands without explanation** are shown —active, pending, overdue and completed orders, total assets, open incidents and items below minimum—, with the distributions, the monthly trend and the comparison with the previous period.

20.2 Technical analysis

- **Planned versus unplanned:** the ratio between preventive and corrective maintenance, the maturity indicator of a maintenance department.
- **Preventive compliance.**
- **Pause analysis.**
- **Failure count** in the period.
- **Distributions** of orders by status, priority and type, and of incidents by criticality and status.

20.3 Maintenance indicators

- **MTBF by asset type** — mean time between failures, on the real intervals between incidents on the same asset.
- **Resolution and MTTR by criticality**, kept separate as the UNE-EN 13306 standard does:
- **Resolution:** from the creation to the closure of the incident, in **working hours**, with mean, median (P50) and P90, against the SLA deadline and with the % **resolved on time**.

- **Technical MTTR:** the effective working time of the orders that resolve the incident, without pauses or waits, with mean, P50 and P90.
- **Where the time goes until closure** — stacked chart by criticality: until work starts, repair, pauses by cause (external or own), other waits and administrative closure.
- **Availability by breakdown and by maintenance**, kept separate.
- **Problem assets (*bad actors*):** ranking of assets by number of incidents. It is the 80/20 analysis that identifies where spending is concentrated.

20.4 Request analysis

Incidents reported, resolved, rejected and pending, resolution rate, **average time to first response** and average resolution time.

20.5 Team performance

- **Order status by technician.**
- **Workload distribution by team.**
- **Productivity by user:** assigned versus completed.

20.6 Trend and lists

- **Monthly trend** of orders created and completed, and of incidents created and closed.
- **Latest orders, latest incidents and overdue orders** with days of delay.
- Lists of orders and incidents with their own **filters**, exportable to **CSV** and **PDF**.

List index. The “Lists” button on the top bar gives access to each user's favourites and to an index grouped by category (operations, compliance, assets, costs, warehouse and fleet), with a search box:

LIST	CONTENT
Work orders · Incidents	With all their filters
WOs outside SLA or at risk	With SLA contracted
Preventive compliance	Orders by outcome and % done on time
Overdue readings · Out-of-range readings	With readings contracted (17.9)
Stock below minimum	With warehouse
Costs by customer	With costs and customers (20.15)

Each list only appears to those who have the permission and the module. Each figure on the dashboard opens the list of what it counts, with the same filter.

From the number to the map. “View on the map” applies the list's filter to the map and the operations panel: what is seen on the map is exactly what is in the list.

Lists are served with a cap of 500 rows; when it is reached, the screen **warns before exporting** and suggests refining the filters: a truncated file that is believed to be complete is worse than a warning.

20.7 Operational reports

- **Inventory value** by category and warehouse.
- **Material consumption by asset** over a date range, net of returns.
- **Real cost of a work order** on its record: labour per person, materials and vehicle (5.10).

20.8 Filters and views

The whole dashboard can be filtered by **date range**, **customer**, **branch**, **sector** and **class of work**: everything, only work on assets or only work without an asset. Filters apply to all the figures and to the lists opened from them.

20.9 Geographic dimension and work without an asset

By area. Breakdown by branch and by sector of active, overdue and completed orders and of open incidents, counting each order in its responsible sector (5.22).

Work with and without an asset. Percentage of orders and incidents without an asset and the areas with most work without an asset: a view of work on surfaces that are not in the inventory, which points to where it would be worth carrying out an inventory.

Heat map. The density of orders or incidents over the last 90 days is drawn as a map layer (21.1).

20.10 Real costs

Cost **incurred** in the period —not budgeted—, which can be grouped by customer, task type, branch, sector or month. Each amount is dated when it occurred:

COMPONENT	CALCULATION
Labour	Each person's timesheet (5.19) at their rate in force on the day worked . Hours pending review are reported separately and do not count until they are approved
Materials	Consumption minus returns, at the cost of the movement
Vehicle	Fuel for the route legs assigned to the order

If there is no rate for a person, or no cost for a route leg, the amount is marked **incomplete**: the total does not pretend to be complete when it is not. The amount spent on **rework due to “non-compliant”** results and on **non-billable** work is reported separately, as are the **travel** kilometres between the base and the orders. Amounts are in the **organisation's currency**.

20.11 Staff capacity

Per person and period (up to 93 days):

- **Shift hours**, deducting public holidays and absences.

- **Assigned hours:** estimated duration of the planned orders. **Load** = assigned / shift.
- **Hours worked**, from the timesheet (with the real-time field module). **Utilisation** = worked / shift, with a configurable target.
- **Unplanned:** pending work that does not yet occupy anyone.

20.12 Fleet and tracking

With the fleet module or the real-time field module, for periods of up to one year:

- **Route legs:** number, kilometres, driving hours and fuel, by person, by vehicle or by month.
- **Tracking alerts**, by person: detected visits not recorded, orders in progress without positions and days with out-of-shift activity not yet regularised.

20.13 Scheduled reports by email

The dashboard summary is sent by email, as a **PDF**, periodically and with the filters the dashboard had when it was scheduled:

PARAMETER	OPTIONS
Frequency	Daily, weekly (with day) or monthly (day 1 to 28)
Time and time zone	Sending time in the chosen zone
Report period	The previous period, period to date, last 7 or 30 days
Recipients	From 1 to 20 addresses
Amounts	Optional, only for those who can see them

It is sent in the organisation's language, currency and date format. If any recipient would not be able to see those figures in the application, this must be explicitly confirmed. The “**Sends**” button brings together the schedules, their history and “**Send now**”.

20.14 Operational home page

The home page is the full-screen map with the operations panel on the right (21.1): the day's command post –what needs attention, where it is and who is doing it– without leaving the map.

20.14.1 Pulse of the day

A strip of counters above the map summarises the orders that need attention today:

COUNTER	WHAT IT COUNTS
SLA missed · SLA at risk	Active orders past their deadline or close to using it up
Overdue · Due today	By due date, in each order's time zone
Unassigned	Pending with no staff

COUNTER	WHAT IT COUNTS
In progress · Paused	Current status
Review	Paused orders flagged for review (15.6)
To validate	Completed orders pending validation (“To close” in the supervised cycle)

Clicking a counter filters both the map and the list with the same criterion used to count, so the number and what is shown always match.

20.14.2 Synchronised map and list

- The list and the map **highlight each other**.
- “**Visible only**” limits the list to the area of the map on screen.
- The list loads 50 at a time, with search, sorting and counts by status.
- Each order has a “...” menu with the valid status changes.

20.14.3 Bulk actions

Multiple selection mode —cards, clicking on the map or a **rectangle drawn on the map**— to assign, change the priority, reschedule or cancel up to 200 orders (5.15).

20.14.4 Dispatch

On the per-technician view of the Planning tab:

- “**To dispatch**” queue: pending orders, those with no staff first.
- **Dragging** an order to a technician's row and day assigns and plans it; orders already planned are moved in the same way.
- “**Assign to the nearest**”: the closest technicians by driving time, even without GPS tracking, filtered by their mandatory training (10.6).
- **Impact on the working day**: for a pending order, what it means to fit it into each candidate's working day —today or in the next three days—, with arrival time, extra driving and orders that move or no longer fit. Nothing is applied without confirmation (10.8).

20.14.5 Map context menu

Right-clicking offers, depending on what is under the pointer: on an **order**, view the record and assign to the nearest; on an **asset**, view its record and create an order; and at **any point**, new order here, new incident here, nearest technicians and copy coordinates.

The fleet layer colours vehicles with an order in progress and draws the line from the technician to the order.

20.14.6 Up to date without reloading

The page updates itself every 30 seconds while it is visible: the alerts bell, the pulse of the day and the order and incident views receive only what has changed, and the list flags new or updated orders. When the tab is in the background, no queries are made.

20.14.7 Bounded views

The Kanban, the calendar, the dispatch queue and the optimiser load only what is visible or what is searched for: no view downloads the full backlog.

All these views respect the technician's scope (1.13).

20.15 Cost reports

The administration area has a **Reports** block with two screens that expand on the dashboard's real costs using the same criteria (20.10):

WO costs. What it has cost to carry out the orders in a period —staff, materials and fuel—, with filters and groupings, and **order by order** from highest to lowest cost, with hours, kilometres, technicians without a rate, route legs without a cost, hours to review and whether it is billable or a rework. CSV export in the user's language.

Costs by customer. What it has cost to serve each customer, **by month, quarter or year**, with charts by customer and by concept and comparison with the previous period. It is cost, not margin: there is no revenue.

21

Mapping and GIS infrastructure

AVAILABILITY

Core

SURFACES

Web · Mobile · Customer portal

WHAT IT IS FOR

To sustain a serious mapping operation —hundreds of thousands of features, editing on the map, external layers, use without coverage— without depending on third-party map services or their quotas for the essentials.

21.1 The map as the main interface

The web application is a full-screen map with:

- **Layer panel** with grouping by families, zonings and sectors, bulk selection and direct access to the list and editing of each layer.
- **Branches layer**: the bases from which technicians and teams set out.
- **Work order and incident layers**, with their point and, close up, the **area** of those located on a surface. A technician only sees their own.
- **Heat map** of orders or incidents over the last 90 days (configurable up to 730), recalculated when the map is moved.
- **Operations panel** with five tabs: Work orders, Incidents, Readings, Planning and Assets (20.14).
- **Customer filter**, which limits everything shown (3.1).
- **Address search** and **global search** (25.7), which finds orders, incidents, assets by code or attributes, and customers.
- **Information tool** to query any feature with one click, and **context menu** (20.14.5).
- **Advanced filters** for orders and incidents, which the map and the list apply in the same way.
- **Thematic maps** of orders and incidents (by status, priority...) and the “**Asset status**” thematic map (4.11), with its **legend**.
- **Alerts that lead to the map**: clicking an alert opens the feature on the map.
- **Integrated Street View**.
- **Geometry editing** on the map, to draw and modify assets, zones and sectors, and the point or area of an order or incident. Moving an asset leaves a trace in its position history (4.14).

21.2 Map layers

Assets are served to the map by area and zoom level, always from the organisation itself, excluding permanently decommissioned assets and the layer's hidden fields. There are dedicated layers for work orders and for incidents, and dedicated layers for the mobile app and for the customer portal, each with its own data scope. All of them require a session or a temporary credential.

21.3 Map cache and versioning

Each layer decides whether it is served from **cache**, so that an area already queried is not recalculated. When an asset is edited, **only the affected area is refreshed**, not the whole layer at every zoom level. Each layer has a **version** that changes with its assets, its definition or its styles, and open maps apply it without reloading.

21.4 External WMS layers

Integration of third-party WMS services (cadastre, urban planning, official mapping):

- **Automatic discovery** of the layers offered by the service, only towards public Internet servers.
- Configuration of address, layer, version, reference system, transparency, opacity, style, format, legend and default visibility.
- If the service does not offer the map's projection but does offer another known one (geographic, ETRS89, UTM...), the system **reprojects** the image.
- WMS layers also reach the app.

21.5 PNOA/IGN orthophoto

Generation of an **offline package** with the PNOA orthophoto published by the IGN, limited to the extent of the organisation's assets, at zoom levels 15 to 18, with a **prior estimate** of the size, status query and download. A failed generation does not destroy the previous one. It gives the field technician the real orthophoto of their area, offline.

21.6 Own base map

In-house generation of the **base map from OpenStreetMap** for any region in the world, with a catalogue of regions and a recommendation of the appropriate one. The base map does not depend on an external provider or on a key with a quota.

21.7 Assets for offline use

Generation of **map packages** with the organisation's assets and their attributes, which the app downloads and uses without a connection. They are published in one go—a download never sees one half-finished—and regenerated every night. The app can also request **only the assets modified** since the last generation, and the list of those it must remove.

21.8 Own route calculation

Route calculation on the road network, without depending on external services:

- Point-to-point **route** with distance, duration and geometry (fleet route legs, chapter 13).
- **Time and distance matrices** between many origins and destinations, the basis of staff proximity, dispatch, the planner and the ordering of reading routes.

It is used only by the platform's internal services, and if it stops responding the system says so immediately instead of keeping users waiting; the planner keeps working with an approximate estimate. It is designed for the local scale of a maintenance operation, not for continental journeys.

21.9 Map performance

The map loads and responds quickly with hundreds of thousands of assets: the extent of the asset base is precalculated, start-up requests everything it draws in parallel, layers and their versions are requested in one go, static resources are cached in the browser, responses are compressed, editing only refreshes the affected area and the attribute search has a time limit that warns instead of hanging.

21.10 Layers fed from outside

As well as being drawn or loaded from a file (4.4), a layer can be maintained from an external system:

- **External data source (REST API):** the layer declares the remote connection, the geometry field and its reference system, the identifying fields, the date field for **incremental synchronisation** and the **run times**, in the layer's local time. Each synchronisation recalculates the codes, refreshes the affected map areas and records the movements in the position history.
- **WMS** as a view-only background layer (21.4).

22

Customer portal

AVAILABILITY

Included in the customer management module, with a limit on customers and on portal users per customer; each customer also has its own portal switch

SURFACES

Standalone web, adapted to mobile

WHAT IT IS FOR

To give the end customer a transparent window onto the service they receive, to turn their reports into structured incidents without going through the phone and to let them take part in the close-out of the work.

22.1 Portal users

An identity completely separate from internal users:

- Email, own password, name, **job title**, phone, avatar and **language**, which users can change themselves.
- They belong to a specific **customer**.
- **Portal role** and, optionally, the customer **zones** to which they are restricted (22.2).
- Activation by email with an expiring link and **activation resend**.
- Single session, record of last access and own password recovery. **Deactivating** a user ends their open sessions.

Portal users are created from the customer record or, within their own company, by its **customer administrator** (22.8).

22.2 Portal roles and scope

Each user has a role whose permissions apply across the whole portal, with everything denied by default:

ROLE	WHAT THEY CAN DO
Basic requester	Reports incidents and sees only their own and the orders that address them. No assets on the map and no dashboard
Technical supervisor	Sees the assets, all incidents and work of their company, the dashboard and the reports. Reports incidents
Manager / approver	The same and, in addition, gives acceptance on completed work (22.4)
Administrator	The same and, in addition, manages the portal users of their company (22.8)

Zones. A user with assigned zones only sees the assets, incidents and orders that fall within one of them, and always the incidents they reported themselves. Without zones, they see their whole customer. This applies to lists, records, map, dashboard, preventive maintenance and email alerts.

22.3 What the customer sees

Always within the scope of their role and their zones:

- **Work orders** that affect them, active ones and history, with filters for **awaiting acceptance** and **reviews**. The record includes dates, location, forms, photos and the acceptance status.
- **Service deadlines:** the status of the deadlines of each order and incident and, for orders, the **pauses that stop the clock** with their cause and duration (19.8).
- **PDF report** of their completed orders, in a customer version: without internal notes and with the signatory's ID number masked.
- Their own **incidents**, open ones and history, with their location and **the orders that address them**.
- **Upcoming maintenance:** preventive orders and the forecasts for the coming days (60 by default, up to 180).
- **Assets** within their scope: record, photos and associated incidents.
- **Their own map** limited to their assets, with its styles, layers and filters.

They never see the internal notes of an incident or the contact details of whoever reported it.

22.4 What the customer can do

- **Create incidents**, with attachments and type, placing them at a **point** on the map, on the **affected assets** or by drawing an **affected area**. They come in as pending, linked to the user, and generate the new incident alert.
- Check the status and progress of their work.
- **Give acceptance** on completed work (manager or administrator): **“Satisfied”** or **“Not satisfied”**, with a mandatory comment in the second case. If the order is reopened and completed again, acceptance is requested again.
- Each customer decides whether acceptance is **required**: if so, the order cannot be closed without it. Otherwise, it is informative.
- It can set a number of **days for tacit acceptance**: with no response within that period, the work is deemed accepted.
- **“Not satisfied” opens a review order**, urgent, with no technician and not billable, linked to the original, and alerts the company so that it can choose who carries it out (5.21).

22.5 Portal dashboard

Dashboard with a summary and charts: incidents by type and by criticality, work by status, monthly trend, latest incidents and latest completed work. Indicators: open and critical incidents, work in progress and completed in the month, average resolution time and incidents outside SLA. In addition:

- **Service level**, with SLA: response and resolution on time for incidents and work, and their monthly trend.

- **Preventive:** compliance and inspections scheduled for the next 30 days.

22.6 Email alerts

The portal alerts its users by email, without them having to log in:

- **Incident received, in progress, resolved or not processed, and work completed,** with the PDF report attached and, if the customer has it, the acceptance request.
- Before sending, the system checks that the alert **is still true**: if a closure has been undone, the email does not go out.
- Each user chooses in “**My alerts**” what they receive, and each email is in their language with a direct link to the incident or the order.

22.7 Access security

- **Isolation:** the portal session only works for the portal, and all queries are restricted to its organisation, its customer and the scope of its role and its zones.
- **Company access:** on every access the system checks that the customer is still active and has the portal switched on; deactivating a customer or switching off its portal **closes its sessions**.
- **Files:** only those within their scope can be downloaded.
- **Login attempts:** temporary, escalating lockouts after several failures, without revealing whether the email exists.
- **Passwords:** the same policy as internal users (1.10).
- Logins, user creation, changes, acceptance responses and actions by the customer administrator are recorded in the **audit log**.

22.8 Customer administrator

A user with the **Administrator** role manages the portal users of their own company without going through the maintenance company: invite, edit, resend the invitation and delete. Only users of their customer, never another administrator, only with the zones they themselves have and within the plan's maximum number of users per customer.

22.9 Portal on mobile

The web portal is **adapted to small screens**: the map at the top and the panel in a collapsible sheet, which collapses when working on the map and opens when viewing a detail.

23

Mobile app

AVAILABILITY	Core; the fleet and real-time field functions depend on their modules
SURFACES	Mobile (app for iOS and Android)
WHAT IT IS FOR	To let the technician carry out the full working day from their phone, including in areas without coverage, and to have what they do return to the system without manual transcription.

23.1 Session and start-up

- **Its own login**, independent of the web one: logging in on mobile does not close the browser session, and a new login in the app invalidates the previous one.
- A session **unused for 60 days expires**, and users can end it from their web profile, for example after losing their phone. “**Log out**” also revokes it on the server.
- Start-up delivers, already resolved, what the app needs to decide which screens to build: permissions, modules, capabilities, vocabulary, **weekly working hours** and language.
- The app responds in the language of the user or of the phone.
- A user **without a staff record cannot log in** (1.3).
- After several failed attempts, the app knows how many seconds the lockout lasts, and the phone from which the user has already logged in can keep logging in even if a third party has locked the account.

23.2 Assigned work

- **My assignments**: overdue ones, today's, those for the next 3 days and those without a date, plus those completed in the last 3 days.
- **Order detail** with everything needed to carry it out, including its time zone, its actual dates, the time currently running, whether it requires a signature, its own location and its **due date**.
- **Detail of reading batches** and of incidents.

A technician can only reach **the orders assigned to them**. The status of an order only changes through **synchronisation** (23.11), with the same rules as the web: a valid transition in the organisation's workflow, the same permissions per status change (5.1), signature, mandatory forms, customer acceptance, propagation to the incident and stock consumption.

23.3 Materials, tools and PPE

- **Item catalogue**, including items with no stock, with the total stock level as a guide.
- **“My van”**: the technician's warehouse —that of the vehicle they are driving today, their van, their warehouse or their branch's— with how much there is of each item, the minimum and what is **committed** by open orders.
- Each material line on the order states **which warehouse it comes from** and how much has been consumed.
- **Add material and report what was used**, also offline.

What is used consumes stock. Stock goes down when the technician reports what was used, not at planning time. The app reports the total used on the line, and the system only moves the difference: a resend does not deduct twice and lowering the quantity returns stock. The technician does not choose a warehouse. Once the order is validated, closed or signed, the material can no longer change.

Task requirements. The app receives the **PPE** and **tools** required by the tasks of each order, and uploads the PPE the technician confirms on starting, with the reason if any is missing (5.23).

23.4 Field forms

Completion of the order's forms with a layout adapted to mobile, saving as a draft and marking as completed. Photo, signature and GPS fields capture directly from the device.

- Tasks are completed with their forms (18.6).
- Completing, validating or closing an order with an incomplete mandatory form is not allowed; responses arriving in the same synchronisation count.
- Withdrawn fields and sections disappear from the app, without losing what has already been answered.

23.5 Photographic evidence and signature

- Order **attachments**, with thumbnail, GPS coordinates, capture date and a simulated-location flag.
- Customer **signature** on the order, or a **signature exception** when they cannot sign, with an evidence photo if the reason requires one.
- The app knows, already resolved, whether the order **requires a signature**, and it cannot be completed without one (5.9).
- Every signature is **sealed** on receipt with a snapshot of what was signed —including the form responses—, so the responses must arrive with the signature or before it (25.16).

23.6 Assets in the field

- **Asset record** with its properties, **photos** and **documentation**.
- **Creation of new assets** from the field, which come in **pending validation**.
- **Catalogue of asset types**, their properties and their **statuses**.
- **Layer tree** (including WMS layers) and **styles** per layer, to draw the map the same way as on the web.

- **Geometries of the assets on the technician's orders**, and their own location, to highlight them.
- **Assets modified** since the last download, to update the local copy without downloading it in full.

Decommissioned assets, assets from disabled layers and assets of deactivated customers do not arrive as live: the app receives the list of those it must remove.

23.7 Change requests and asset editing

The technician can propose a **correction to the geometry or data** of an asset instead of modifying it directly: the proposal awaits approval by a manager (25.4), and pending ones are drawn faded on the map.

Anyone with permission can also create, modify and decommission assets offline, with the same rules as on the web; every change of position is recorded in the asset's history.

It is a data quality mechanism: the person who best sees the error is the one standing in front of the asset, but the inventory is not touched without control.

23.8 Incidents and quick report

- Creation of **incidents** from the field with type, criticality, location and photos.
- **Open incidents** on the map, with their affected area.
- The app **does not delete incidents** or move their status backwards.
- **Completing the order resolves its incident**, unless the technician answers that it is not resolved (5.24).
- **Quick report**: record unplanned work already done without going through the office, also offline. It is created as completed, with the GPS position as the order's location, and the system assigns its **final code**.

23.9 Meter readings

Download of the complete **reading batch**, taking readings with a photograph and synchronising back. A closed batch does not accept further readings. The system decides whether a reading is **out of range** and, if the point requires it, opens the incident, and the app receives what was decided (17.9).

23.10 Offline mode

This is the app's differentiating feature. Three packages that are downloaded with coverage and used without it:

PACKAGE	CONTENT
Work data	Orders, assets, catalogues, forms and their assignments. Prepared every night and on demand, and only published complete and verified
Reading batch package	Everything needed for a reading campaign (17.6)
Offline mapping	Assets and PNOA orthophoto of the area (21.5, 21.7)

The download is **full or incremental**: with the marker from the previous synchronisation it only brings what has changed, with a small overlap so as not to miss anything. It also brings the working hours, the pause causes, the task requirements and “My van”, and the list of what the app must **remove**: reassigned, cancelled or deleted orders, deleted incidents, superseded responses, withdrawn fields, deleted items and withdrawn assets.

23.11 Upload synchronisation

A single **synchronisation upload** channel receives everything done offline: status changes with their actual dates and the cause of pauses, quick reports, new incidents, materials consumed, PPE confirmations, **per-technician timers** for the timesheet, form responses, attachments, signatures, asset photos and documents, new and modified assets, readings and decommissionings.

Record by record. Each record is validated and saved separately: what is valid goes in and what is not comes back with its reason, without losing the rest of the working day. Each error carries a **stable code** with which the app decides whether to retry, resend as a creation or withdraw the record and show it to the technician: the parent element is missing, the order is no longer theirs or has been cancelled, a permission is missing, the status change is not possible, the order is already closed or signed, the signature, mandatory forms or customer acceptance are missing, the reading batch has been closed, the plan does not allow more elements, among others.

Permissions per record. Each record is checked against the user's permissions: the status change with the web rule, and material, assets and readings with their own.

Status is validated as on the web. If the technician went offline from pending to in progress and to completed, the system checks that this is reached through valid transitions. A simultaneous cancellation from the office does not end up as an order that is both cancelled and completed.

Resends without duplicates. An incident, a quick report, a material line or a PPE confirmation that is resent is not duplicated.

After saving, in the same upload, the timesheet and the tasks are recalculated, and for each order that has been closed the **signature**, the **mandatory forms** and the **customer acceptance** are checked; if anything is missing, that closure is undone together with its dates and its alerts, and the error goes to the order. Then the status is propagated to the incidents, stock is moved, the map of the uploaded assets is refreshed and the supervisor is alerted to simulated positions and detected conflicts.

23.12 Fleet and tracking from mobile

- **Vehicle check-in** at the start of the working day, by scanning its QR code or choosing it from a list.
- **Geolocation notice**: the app shows the current wording, records acceptance and allows it to be withdrawn (14.8).
- **Position capture** in batches, also with the screen off, only within the working day and with consent (chapter 14).
- **Geofence**: the app detects arrival and departure and **asks the technician** whether to start or finish (chapter 15).
- Viewing of **one's own route** and **one's own sessions**, and **export** of one's own location data.

23.13 Dates

Every date the app sends is an instant with its time zone, taken when the event occurs, not at synchronisation. All dates the app receives arrive with a time zone, and each order brings its own so that they are shown in the local time of the work site, not that of the phone.

23.14 Errors and traceability

Every error has a stable code, which the app uses to decide, and a text in the user's language. Each response carries a request identifier, which makes it possible to locate a specific error from what the user sees. A long synchronisation does not delay other users.

23.15 Synchronisation conflicts

When the office and the mobile change the same thing while the technician is offline, the upload does not simply apply “last writer wins” without leaving a trace.

Revision per record. Work orders, incidents, assets, form values and material lines carry a **revision** number. The app returns it on upload together with the value each field had when it was downloaded.

Field-by-field merge. If the record has changed in the meantime, what only the mobile changed is saved, what only the office changed is kept, and what both changed is a **conflict**, won by the **field owner**: by default the field prevails for actual dates, the incident description, the asset's properties and position, forms and material used; the office prevails for the title and internal notes of the incident and for the asset's status, customer, sector, parent and code. The administrator can change the owner of each group of fields. When the office wins, the record is sent back down to the mobile with its value.

Review in the office. Administration → **Synchronisation conflicts**: a list with filters by technician, record type, status and dates, restricted to each supervisor's branch. For each conflict it is possible to **restore the office value**, **apply the mobile value** or **accept it as is**. Decisions are recorded in the audit log.

Alert and retention. Each technician with pending conflicts generates **one alert**, which closes by itself only when none remain. Reviewed conflicts are deleted after 180 days and pending ones after 365.

23.16 Functions that depend on the app version

The server maintains compatibility with earlier versions of the app: everything new is optional. Some functions only take full effect with an app version that incorporates them: conflict merging (23.15), the question “Is the problem resolved?” (with earlier versions, completing resolves the incident), the pause cause, “My van” and PPE confirmation, the weekly working hours and the due date in the app, and the fingerprint of responses on signing.

24

System alerts

AVAILABILITY

Core; alerts from modules not subscribed to are neither generated nor shown

SURFACES

Web (alert centre in the top bar) · Email and webhook (escalation)

WHAT IT IS FOR

To make sure that what matters reaches the person responsible without them having to look for it, that it moves up a level if nobody deals with it and that it disappears by itself only when it stops mattering.

24.1 Alert catalogue

ALERT	WHEN IT IS GENERATED	FOR WHOM
High or critical criticality incident	Creation of a pending incident, wherever it comes from	Whoever receives incident alerts
New incident	Creation of a low or medium criticality incident	Whoever receives incident alerts
Unresolved incident	An order finishes without resolving the problem (7.12)	Whoever receives incident alerts
New work order	Creation of an order with assignees, including by preventive maintenance	Each assignee
Urgent work order	Creation of an urgent priority order	Each assignee
Staff without the required qualification	Assignees who do not meet the task requirements, or preventive maintenance without assets or without staff	Each assignee
Customer not satisfied	“Not satisfied” response in the portal (5.21)	General
SLA at risk without staff	Order at risk with nobody assigned (19.10)	Whoever receives SLA alerts
SLA breached	Order that exceeds its resolution deadline (19.10)	Whoever receives SLA alerts
Stock below minimum	Each movement, by material and warehouse	Whoever receives stock alerts
Stock-out	Each movement, if it goes negative	Whoever receives stock alerts

ALERT	WHEN IT IS GENERATED	FOR WHOM
Material missing for a WO	When saving an order with material the warehouse does not cover	Whoever receives stock alerts
Start without PPE	A technician starts an order without some required PPE	Whoever receives safety alerts
Order paused at end of working day	Automatic end of working day (14.6)	Whoever receives tracking alerts
Technician outside the area of a WO in progress	Geofence (15.3)	Whoever receives tracking alerts
Simulated positions	Mobile with a fake location (14.10)	Whoever receives tracking alerts
Vehicle pending review	Low-confidence vehicle assignment	Whoever receives fleet alerts
Synchronisation conflicts	A technician's mobile has overwritten changes made by the office (23.15)	Whoever receives synchronisation alerts
New asset · geometry modified · data modified · asset deleted	Requests from the mobile app	Whoever receives asset alerts

24.2 Alert structure

Each alert has a **severity** (critical, high, medium, low) and a **status**; a title and message, which are shown **in the language of whoever is viewing them**; the **related element** it navigates to when clicked; the **branch** it affects; who has seen it, **who is handling it** and its **escalation** level (24.7).

24.3 Alert centre

In the top bar: pending counter, list sorted by severity, direct navigation to the element, mark one or all as seen, and **full history**.

Seen is per user. A general alert only stops showing for the person who marks it; the others keep seeing it until someone **resolves** it or it resolves itself.

“I’ll handle it”. Whoever receives an alert can take charge of it: the alert does not close—it remains in the others' bell, showing who is handling it— but it **stops escalating**. The bell also shows the escalation level and where the alert has been.

24.4 Who receives each alert

An alert addressed to a person is only shown to them. A **general** one is shown to anyone who meets all of the following:

1. **The organisation has the module that produces it.**
2. **They have that family of alerts assigned** in their role: incidents, GPS tracking, stock, fleet, assets, safety, SLA or synchronisation (1.7). By default, the administrator and the supervisor have all of them and the warehouse role has stock.

3. **It belongs to their branch**, if the organisation works with branches. Administrators and users without a branch see them all.
4. **They did not trigger it themselves**: whoever creates an incident does not receive the alert for what they have just done.

Anyone who has received an alert **through escalation** always sees it.

24.5 Alerts that close by themselves

- **Incidents**: when the incident stops being pending or is deleted.
- **Stock**: a single alert per material and warehouse, which is updated with each movement and closes on restocking.
- **Missing material**: updated when the order is saved and closed when it is completed or cancelled.
- **Work order**: the pause at the end of the working day closes on resuming, and the technician outside the area, when the order stops being in progress; if a closure from the app is undone, they reopen.
- **SLA**: when the order has staff, finishes or changes its deadline.
- **Customer not satisfied**: when staff are assigned to the review.
- **Synchronisation conflicts**: when they have all been reviewed.

24.6 On-screen alerts

Some operations also return **alerts in the response itself**: creating an order or a preventive schedule (stock and qualification), inventory movements (below minimum or stock-out) and the planner (excluded resources and uncovered orders).

24.7 Alert escalation

An alert that nobody deals with **moves up a level** according to rules the administrator manages in **Administration** → **Alert escalation**:

- A **rule** applies to certain **alert types**, can be active or not and can **use shifts**.
- Each rule has up to 5 **levels**. Each level states after how many **minutes unattended** it escalates, **to whom** —the branch manager or deputy, administrators, a role, a specific person or whoever is **on call**— and **through which channel**: the bell, **email** or a **webhook** (25.13).
- With “**use shifts**”, anyone not on shift is skipped and the alert goes to the on-call person. If nobody is left, it goes to the administrators: an escalation is never lost.
- Escalation **stops** when someone clicks “I’ll handle it” or when the alert closes by itself.

Each organisation starts with the “**Critical alerts**” rule: SLA at risk and breached, critical incidents and urgent orders → after 15 minutes, to the branch manager; 30 minutes later, to the administrators; in the bell and by email. Alerts are not sent by SMS or as notifications on the phone.

25

Public API, integrations and data governance

AVAILABILITY	Core; the public API is enabled by creating API keys and webhooks by registering destinations
SURFACES	Public API · Outgoing webhooks · Web (administration)
WHAT IT IS FOR	To let Maptainer's data live in the customer's information system —their ERP, their corporate GIS, their management dashboard— and to let the customer's data enter Maptainer without manual exports; and to make what the system stores demonstrable, portable and kept only for as long as it should be.

25.1 Public API

An independently documented API, with its own OpenAPI specification and interactive documentation accessible without a key, which includes the webhook catalogue. There is an integration guide, a reference client and, for integration with an ERP, the “Integration with an ERP” manual with its mapping sheet for the most common maintenance ERPs.

It is a **read API with a single write operation**: the creation of incidents from another system (25.14).

Authentication by API key, which the organisation generates from the application.

Scopes. Each key declares what it can access: assets, readings, work orders, incidents (read and create), inventory, staff, files and full export. Granting a scope is not enough: it is also validated against the **contracted modules**. A key with access to readings on an organisation without the readings module receives an explicit error, not an empty list. A key can also be **restricted to a customer**, so as to create only that customer's incidents.

Quota of 120 requests per minute per key (60 incident creations and 10 exports), which the response communicates on every call.

What can be read:

QUERY	WHAT IT RETURNS
Catalogue	What that key can read, with its declared limits
Data	Any of 60 resources , paginated by cursor: assets, work orders and their full record, incidents, readings, photographs, documents, signatures, position history and the catalogues and master data (customers, branches, sectors, zones, users, statuses, work types, forms, staff, inventory)

QUERY	WHAT IT RETURNS
Changes	Incremental synchronisation: what has changed since a date, entity by entity, distinguishing creations, modifications and deletions
Deletions	Deleted records, so that a replica can remove them (kept for 90 days)
Files	A specific file from the record that contains it, with its fingerprint to verify integrity

Reading paginates by **cursor**, not by page number, so that going through millions of assets does not get slower with every page. Resource and field names form a **versioned public contract** that does not expose the internal structure. Personal data of staff and users is not returned when they are read only as a supporting catalogue.

25.2 API key management

Administrator self-service, in *Administration* → *API and data copies*:

- **Creation** with name, scopes (only those of contracted modules), optional **expiry** and, if desired, a **customer**.
- **List** with the prefix, scopes, status, expiry and last use; activation, deactivation and deletion.
- **Rotation**: creates a new key with the same scopes and keeps the previous one live for a few days, so that it can be changed in the integration without interrupting the service.
- The full key is only shown when it is created or rotated, and every change is recorded in the audit log.

25.3 Remote connections (data input)

Connections to the customer's **external REST APIs**, with up to three addresses with their credentials. Combined with a layer's external data source (21.10), they allow Maptainer to **import and keep up to date** an inventory that lives in another system:

- **Scheduled synchronisation** at each layer's times, without overlapping a manual synchronisation.
- **Incremental**: with a previous successful synchronisation, only the changes since then are requested, with a small overlap; a record marked as deleted decommissions its asset.
- **Caps** on size and time.
- **Effects**: codes are recalculated, the map of the affected areas is refreshed and every change of position is recorded in the history.

The credentials of these connections belong to third-party systems and are treated as such: **they are not shown again** once saved, they are stored **encrypted**, only the administrator manages them and the connection only goes out to **public Internet destinations**, without following redirects. A connection in use cannot be deleted.

25.4 Change requests

A data governance flow for corrections proposed from the field (23.7):

- The request stores the asset, who requested it and when, the **type of change** (creation, modification or deletion) and what is proposed: geometry, properties, photo or status.
- List and query with **display on the map** and comparison between the current and the proposed value.
- **Approving** or **rejecting** is a permission in its own right, and nobody resolves their own requests except the administrator.
- **The decision is kept**, with who reviewed it, when and their **notes**, so that the technician can be told why it was rejected. Each decision is recorded in the audit log.
- Approving a change of position records it in the asset's history with the requester as author (25.17).

25.5 Signature policies

A catalogue of policies —signature always mandatory, never mandatory, or depending on the maintenance type—, assigned per customer. The policy is **enforced on the server**: an order that requires it cannot be completed without a signature or exception, whether the change comes from the web, from bulk actions or from the app (5.9). Each signature is sealed (25.16).

25.6 File uploads

Each upload checks what it receives and has its own cap:

UPLOAD	FORMATS	CAP
Logos	Image	5 MB
Inventory images	Image	10 MB
Asset photos	Image	20 MB
Incident attachments	Image, PDF, Office without macros, CSV, text	20 MB
Order documentation	Image, PDF, Office without macros, CSV, text	50 MB
Asset documents	The above plus DWG/DXF and ZIP	50 MB
Geospatial and CSV imports	Those of each importer	150 MB

Content is validated internally, not by extension: images that open as images, PDFs without executable actions, modern Office documents without macros, controls or embedded objects, and rejection of old formats with macros and of double extensions. The name on disk is set by the system. Every night, uploads that nobody uses any more are removed.

25.7 Global search

A search box accessible from the header that finds work orders and incidents (by code, title and description), assets (by code and by the value of their visible attributes) and customers, and navigates directly to the result. It does not return deleted items, and whoever only sees their own orders only finds their own.

25.8 File access security

- **Every** file and all mapping require a session or API key **and** membership of the organisation. A file from another organisation responds as non-existent.
- Uploaded content is never interpreted in the browser: anything that is not an image or a PDF is downloaded.
- Order documentation supports **expiring links** (10 minutes by default), to share a specific document.
- The full data copy is downloaded with the administrator's session, so that the audit log records who took it.
- Underneath all of this, the database restricts each query to its organisation (2.14).

25.9 Scheduled tasks

TASK	FREQUENCY	WHAT IT DOES
End of working day	Every 15 min	Closes tracking and pauses in-progress orders (14.6)
Inactive tracking	Every 10 min	Closes tracking with no positions in 8 hours
Stalled orders	Daily	Marks for review those paused for too long (15.6)
In-progress orders without positions	Every 10 min	Marks them for review (15.7)
Return to base	Hourly	Closes each working day with its return trip (13.5)
Retention of personal data	Daily	Removes expired positions and trails (25.15)
Fuel prices	Every 30 min	Downloads the official index from the Spanish Ministry (13.2)
Recurring orders	Daily	Creates the preventive orders and forecasts (chapter 6)
Reading batches	Daily	Creates the batches according to their rules (17.5)
Timesheet	Daily	Recalculates labour for the last 35 days (5.19)
SLA pauses	Every 5 min	Maintains the calculation of justified pauses (19.8)
SLA alerts and escalation	Every minute	Opens, closes and escalates alerts (19.10 and 24.7)
Tacit acceptance	Hourly	Deems accepted any order not responded to in time (22.4)
Portal alerts	Continuous	Sends the portal emails (22.6)
Offline data and mapping	Daily	Regenerates the packages for mobile (23.10 and 21.7)
Reading thresholds	Daily	Recalculates the statistical thresholds (17.2)
REST layers	Every 5 min	Runs each layer with an external source at its scheduled times (25.3)

TASK	FREQUENCY	WHAT IT DOES
Scheduled reports	Every 5 min	Sends the reports by email (20.13)
Webhooks	Every minute	Sends and retries deliveries (25.13)
Signature timestamping	Hourly	Builds and timestamps the daily batch (25.16)
Data copies	Continuous	Generates and purges full copies (25.10)
Audit	Daily	Purges entries older than the retention period (1.12)

25.10 Full export and data portability

What it is for: to let the organisation take **all** its information whenever it wants, without having to ask for it or depend on anyone. It answers “what if one day I want to leave?”, covers the right to data portability under Art. 20 GDPR and feeds a data warehouse or a corporate GIS without manual exports.

25.10.1 What it includes

The API's 60 resources, always restricted to what the key can read:

DOMAIN	CONTENT
Assets	Assets, types, fields, statuses, photos, documents and position history
Records	Work orders with their full record: tasks, staff, work segments, materials, assets, status history, comments, documents, attachments, signatures with their seal and forms
Incidents	Incidents, attachments, history, assets and linked orders
Readings	Measurement points, readings and batches
Forms	Templates, sections, fields, submissions and responses
Organisation	Customers, branches, sectors, zones and users
Catalogues	Statuses, priorities, work and task types, incident types and statuses, criticalities and origins
Staff	Staff, teams, shifts, working hours, trades and skills
Inventory	Categories, items, warehouses, stock levels, movements and suppliers

Explicitly requesting a resource without permission returns an error: **the request is not silently trimmed**, because a trimmed dump that the integrator believes to be complete is worse than an error.

25.10.2 The flow

The copy is requested, its progress is followed and, when it finishes, it is downloaded through a **short-lived signed link**. It is possible to request only what has been modified since a date, and to include or exclude deleted items, files and spreadsheets, as well as to choose the language of the texts. There is **one copy in progress per organisation** at a time.

“Full backup” from the panel. The administrator requests it in *Administration* → *API and data copies* without needing a key, **with no options that could leave it incomplete**: everything contracted, with files, with deleted items and with spreadsheets, in the chosen language. The request and the download are recorded in the audit log.

25.10.3 The package format

A self-describing ZIP: a package **descriptor**, one data file per resource (NDJSON, one record per line, which can be processed without loading it entirely into memory), deletions in incremental copies, the files with their index and their SHA-256 digest, and a **“Spreadsheets”** folder with the same content designed to be opened in Excel without knowing anything about the API—one CSV per entity, with headers in the chosen language, names already resolved, the assets of each layer with their fields as columns and each photo or document linked to its file— plus a README.

The descriptor records the format version, the moment of the copy (which is the starting point for the next synchronisation), the currency of the amounts, the filters applied, the number of records of each resource and **the files that could not be included**, which are listed instead of aborting the copy.

The entire dump is read from **a single snapshot**: a record that changes while it is being generated does not come out half-written or with inconsistent relationships.

25.10.4 Data conventions

- **Geometries**: always GeoJSON in longitude and latitude (WGS 84).
- **Dates**: ISO 8601. **Identifiers**: text. **Decimals**: numeric.
- **Multilingual texts**: in the requested language.
- **Names**: those of the public contract, not the internal ones.

25.10.5 The limits, declared

Each resource declares whether it supports incremental copying and whether it can report its deletions, and the descriptor warns of what applies to each copy. Publishing these limits is deliberate: it prevents an integrator from building a replica with gaps while believing it to be up to date.

25.10.6 Scale

A copy of 2.2 million records from the largest organisation is generated in a few minutes, and **memory consumption does not grow with volume**: data is written directly into the package, without loading anything in full.

25.10.7 Execution and life cycle

The copy is generated outside the application and starts immediately. An interrupted copy is deemed failed after a while, so that another can be requested, and the package is kept for **3 days** from when it finishes.

25.10.8 Audit

Each copy leaves a permanent record: who requested it, from where, which resources, with which filters, how many records of each, the size and the fingerprint of the package. A full copy includes personal data, so the trail is not optional.

25.11 Errors and request traceability

All errors follow a common structure: a **stable code** with which the client system decides what to do, a text in the user's language and the data that specify it. No error returns internal details. Each response carries a **request identifier**, which makes it possible to locate a specific error from what the user sees.

25.12 Audit log

Sensitive actions are recorded in a permanent log —logins, users and staff, permissions, personal data, operations and data governance, integrations and organisations— that cannot break the operation it records, is not lost if that operation is undone, stores no secrets and cannot be modified or deleted from the application. Retention is configurable (one year by default, never less than 90 days). The administrator consults it in *Administration* → *Audit* (1.12).

25.13 Outgoing webhooks

What it is for: to let another system —an ERP, middleware, a management dashboard— learn of changes to work orders and incidents **as soon as they happen**, without polling every so often. It is the natural return path for incident creation via the API: the ERP reports the fault and Maptainer notifies it of the order, its closure and its outcome.

25.13.1 Events

EVENT	WHEN
Order created	Creation of a work order
Order: status change	Any status change
Order completed · validated · closed · cancelled	Transition to each of those statuses
Incident created	Creation of an incident
Incident: status change	Any status change
Incident approved · resolved · closed · rejected	Transition to each of those statuses
Alert escalated	An alert moves up a level with the webhook channel (24.7)

A status change produces both the generic event and the specific one: anyone who only wants closures subscribes only to “order closed”.

25.13.2 How they are generated

Events are recorded **at the same moment and in the same operation as the change**: no matter who makes it —the web, the app, a scheduled task or the API— the event goes out, and if the change is undone, so is the event. Each event also states whether the change was made by an API key, so that the receiver can recognise the echo of its own creations.

25.13.3 The content

Each event carries its identifier (the same in all retries), type, version, instant and organisation; the entity with its human-readable code, the previous and the new status, the customer and the external reference; who made the change; the **full record** with the same contract as the API; and related items: for an order, its incidents with their external reference; for an incident, its orders.

25.13.4 Delivery signature and security

- **Signature according to the Standard Webhooks specification**, verifiable by the receiver with a secret that is shown only once and that can be **rotated without interruption**: during a grace period both the previous and the new secret are valid.
- **Only https and public destinations**, without following redirects and with bounded timeouts.

25.13.5 Delivery and retries

- Deliveries to each destination are made in order, and a slow receiver does not hold up the others.
- A successful response closes the delivery; anything else is retried for **about 42 hours**, with increasing intervals.
- Delivery is **at least once**: the receiver discards duplicates by the event identifier.
- A destination **without a successful delivery in 3 days** is deactivated automatically and the administrators are notified by email.
- Events are kept for 30 days.

25.13.6 Configuration

Up to **10 destinations** per organisation, with name, address, events (only those of contracted modules) and, if desired, a **single customer**. From the same screen it is possible to **test** the destination, view the **delivery history** with its response and **resend** a specific delivery.

25.14 Incident creation via the API

What it is for: to let an ERP or a control system report “there is a problem” and have Maptainer take care of the rest with its own rules: review, order, planning, execution and closure.

- **Catalogue** of incident types, criticalities and statuses, to build the equivalence table with the ERP.
- **Creation** through the same process as the web —customer, branch, SLA, alerts and code—, in pending status and with the origin *Integration (API)*.
- **Everything by code**: the type, the criticality and the assets (up to 50) are indicated by their code, and an unknown or ambiguous code returns an explicit error. Without assets, a point and an address can be given.
- **External reference** (the ERP notification number), unique per organisation: a second creation with the same reference returns the incident that already exists.
- **Protection against resends**: repeating the same request within 24 hours returns the same response without creating another incident.
- **Search by external reference** and query of an incident with its status.
- A key restricted to a customer only creates and sees that customer's incidents.

To find out when it is resolved, the integrator subscribes to the incident and order webhooks, which carry the external reference.

25.15 Data subject rights and retention of personal data

What it is for: to handle the rights of access, portability and erasure (Arts. 15, 17 and 20 GDPR) without breaking the maintenance history or its figures, and to have location and access trails expire by themselves.

25.15.1 Export of a person's data

A ZIP with all their data in JSON, one CSV per section, a README and the files they uploaded and their signatures. It can be requested for a user, a member of staff or a portal user, and each person can download their own from their profile. Each download is recorded.

25.15.2 Erasure by anonymisation

Erasing a user, a member of staff or a portal user **anonymises** rather than deletes, because deleting would break the history: orders would be left without an author and without a cost. What identifies the person disappears —name, email, phone, notes, home address, contact details in the incidents they reported, ID number in their signatures— and their location is deleted; orders, assignments, histories and the proof of the work done for the customer are kept.

Erasure does not change any cost figure: before erasing, their timesheet is fixed, and the audit log stores the economic fingerprint before and after as proof that nothing changed.

25.15.3 Retention by period

What has its own purpose is kept; the “where” expires:

DATA	PERIOD	WHAT IS DONE
GPS positions	60 days	Deleted
Start and end positions of each tracking session	60 days	Cleared; the session and its dates remain
Path of each route segment	60 days	Cleared; kilometres, duration and cost remain
Position of order events and status changes	60 days	Cleared; the event and its time remain
Log of access to location data	2 years	Deleted
Unused app sessions	90 days	Deleted

Portal users of customers deactivated more than 12 months ago are anonymised automatically.

25.15.4 Acceptance of legal texts

When activating an account, internal or portal, the privacy policy and the legal notice must be accepted, and each acceptance is recorded with document, version, date, origin and browser, with no possibility of modification. The application does not use third-party services that receive the visit.

25.16 Integrity seal and timestamping of signatures

What it is for: to make the customer's signature **proof of what they signed**, not just an image.

The seal. On receiving each signature or signature exception, the system stores a **snapshot of what was signed** —the order, the signatory, the position, the time, the signature image and **all the form responses** of the order— together with its **SHA-256 digest** and the server time.

- **The signatory is pseudonymised** in the seal, so that erasing the person (25.15) does not invalidate the seal.
- **Immutable:** the seal cannot be modified.
- **Verification:** the order record shows the seal status of each signature —**intact**, **modified** (and what changed) or **unsealed**— and allows what was signed to be downloaded.

Third-party timestamping. Once a day, the digests of all the previous day's seals are combined into a single one, which is sent to a **timestamping authority** (RFC 3161), which signs “this digest existed at this time”. Only the digest reaches the authority, never any data. From the record it is possible to download a **proof** with what was signed, the inclusion proof, the authority's response and a program that verifies it **outside Maptainer**.

The seal also travels in the order PDF, in the public API and in the full copy. It is not a qualified electronic signature: what it provides is integrity and a certain date for what was signed.

25.17 Asset position history

Every movement of an asset leaves a record of the previous and the new position, **when**, **who** and **from where** (web, mobile, approved request or REST synchronisation), whatever the path (4.14). It can be consulted in the asset record, through the public API and in the full copy, with its own spreadsheet.

A

Modules, permissions and limits

A.1 How access is governed

Four distinct, non-interchangeable mechanisms:

MECHANISM	ANSWERS	WHO DECIDES
Contracted module	What have you subscribed to?	The organisation, with the sales team
Organisational capability	How does your company work?	The organisation's administrator, at no cost (2.3)
Level of detail	How much rigour do you want?	The organisation's administrator, at no cost (2.4)
Permission	What can this person do?	The organisation's administrator (1.7)

For the two in the middle, a **starter profile** (2.5) sets them in one go. Beneath all of them, the database restricts each operation to its organisation (2.14), and assignment-based scope limits the technician to their own orders (1.13).

A.2 Complete matrix

Surfaces: **W** = web · **M** = mobile · **P** = customer portal · **A** = public API.

#	FUNCTIONAL MODULE	AVAILABILITY	SURFACES
1	Authentication, users, roles and permissions	Core	W · M · P · A
1	Audit log	Core	W
2	Organisations and verticals	Core	W
2	Settings, capabilities, levels and profiles	Core, at no cost	W
2	Languages (Spanish and English)	Core	W · M · P · A
2	Help and user manuals	Core	W
3	Customers and customer zones	Customer management	W · P
3	Sectors and zonings	Core, with its capability	W · M
3	Branches	Core, with its capability	W

#	FUNCTIONAL MODULE	AVAILABILITY	SURFACES
3	Teams	Core, with its capability	W
3	Warehouses	Maintenance, with the warehouse capability	W · M
4	Assets and GIS layers	Core	W · M · P · A
5	Work orders	Maintenance	W · M · P
6	Recurring scheduling	Maintenance, with the preventive capability	W
7	Incidents	Maintenance, with its capability	W · M · P · A
8	Maintenance types and tasks	Maintenance	W · M
8	Trades and skills	Core, with the competencies capability	W
9	Staff	Core	W
9	Shifts, absences, public holidays and exceptions	Core, with the working calendar capability	W · M
10	Planning of an order	Planning	W
10	Dispatch to the nearest	Real-time field	W
11	Multi-day planner	Planning	W
12	Time budget	Core	W
13	Fleet: vehicles, fuel, assignments and segments	Fleet and costs	W · M
14	GPS tracking	Real-time field	W · M
15	Geofence and attendance record	Real-time field (+ fleet for cost)	W · M
16	Inventory, movements and suppliers	Maintenance, with the warehouse capability	W · M
17	Meter reading	Meter reading	W · M · A
18	Dynamic forms	Maintenance, with its capability	W · M
19	SLA	SLA (+ per customer)	W · P
20	Dashboard and reports	Core	W · P
20	Order and per-customer costs	Core, with the costs capability	W
21	Mapping, external layers and offline	Core	W · M · P
22	Customer portal	Customer management	W · P

#	FUNCTIONAL MODULE	AVAILABILITY	SURFACES
23	Mobile app	Core	M
24	Alerts and escalation	Core	W
25	Public API, API keys and full copy	Core	W · A
25	Outgoing webhooks	Core	W · A
25	Remote connections	Core	W
25	Change requests	Core	W · M
25	Synchronisation conflicts	Core	W
25	Personal data (export and erasure)	Core	W

A.3 Default permissions of the system roles

Roles belong to **each organisation**: whatever its administrator ticks or unticks does not affect any other.

ROLE	DEFAULT PERMISSIONS
Superadministrator	All, in all organisations
Administrator	All, within their organisation
Supervisor	Full work orders, including closure , quick reports and incidents · full assets, with configuration and approval of changes · inventory with movements · full readings, with their configuration · reports with export, costs and tracking · viewing of staff, customers, sectors, branches and SLA · all alerts
Technician	View and edit their own orders, without closing them · quick reports and incidents · view, create and edit assets · view inventory and record movements · readings. No dashboard or reports
Warehouse	View orders and assets · full inventory, with warehouse control · reports · stock alerts
Read-only	Viewing of orders, assets, inventory, readings, staff, customers, sectors and branches · reports with export, without costs or tracking

They are the starting point: once the role has been created, its permissions are the organisation's decision, and per-user exceptions take precedence over those of the role.

A.4 Subscription limits

LIMIT	WHAT IT CAPS
Users	Internal users of the organisation, except erased ones
Customers	External customers, including deactivated ones; the organisation's own customer does not count

They are checked on **creation**, through every channel: web, app, synchronisation, bulk uploads and synchronisation with an external API. A bulk upload that does not fit is rejected in full before inserting anything, stating how many elements fit: a half-finished import is worse than none. From 80 % of the quota onwards a warning is shown, and the administration area shows the consumption of each limit.

The limit caps growth, **it does not confiscate what already exists**: an organisation above its cap continues to operate everything it has; it simply can no longer add more.

A.5 Subscription validity

The organisation has a status (active or inactive) and a **validity date**, inclusive. These are checked on access and in every open session, on the web, the app, the portal and the API: deactivating an organisation also takes effect for anyone who already had a session open. The message names the organisation and, if it has expired, the date.

Thirty days before expiry, the application shows a notice with the days remaining. When a piece of data is doubtful, the system lets the user in: locking out a customer who has paid is much worse than letting someone in for one day too many.



CONTACT

Something missing that you need to read?

If there is a question the documentation does not answer, it is useful for us to hear it.

EMAIL

info@maptainer.com

PHONE

[\(+34\) 960 659 055](tel:+34960659055)

WEB

maptainer.com