# The DemandFlow Platform

DEMANDFLOW®

Shape it to your business, automate the work between the records, connect it to everything you run, and let your AI assistants use it

PLATFORM OVERVIEW

Every operator ends up with a dozen systems that each hold part of the truth and a team whose job is to reconcile them. DemandFlow® is built the other way round: one platform, one set of records, and every module from fibre inventory to investment planning reading and writing the same facts. The cable in the street, the service sold over it, the change that touches it, the project that paid for it and the ticket about it are one connected record set.

That is only useful if the platform bends to your business rather than the other way round. So DemandFlow is definition-driven: the entities, fields, panels, rules and screens are definitions you can change, extend and create, and the automations, integration flows, APIs and AI assistant interface work on whatever you define.

## What it is

A definition-driven, multi-module operations platform for telecom, utility and infrastructure operators. Over a thousand entity definitions ship with it, covering networks, fibre, RAN, data centres, services, change, monitoring, projects, portfolios, investment plans, finance, procurement, people, recruitment, IT assets and service desk. Each is a definition you can tailor per tenant, and new entities can be added from a form or from a plain-English description. One write path applies your rules, fires your automations, keeps the history, delivers your webhooks and pushes live updates to every open screen. It runs serverless on AWS with each tenant's data in its own tables and buckets, hosted in the UK, EU or US region you choose, and can also be deployed into your own AWS, Google Cloud or Azure estate.

## What you get

- **Entities you can shape.** Change any field, panel, lookup or rule of any shipped entity for your tenant, hide what you do not use, and add fields, panels and whole entities without code.
- **New entities from a description.** Describe what you need to track and the AI drafts the entity, its fields and its relationships for you to review and publish.
- **Automation on every record.** Triggers on create, update, delete and approval, with filters, that update records, create related ones, send email, start approvals, run flows or create CRM records, plus scheduled runs.
- **Integration flows without middleware.** Webhook, schedule, manual and inbound-email triggers stepping through external calls, DemandFlow reads and writes, mapping, branching, loops, Slack and email, with secrets held in a vault.
- **A full REST API.** Every record and every entity, with the same rules, automations, history and webhooks as the application, authenticated by personal access tokens.
- **Built for AI assistants.** An MCP server that lets Claude and other assistants read, search, analyse and draft across the platform, audited call by call.
- **Serverless scale.** Functions that scale with demand, on-demand database capacity, a per-tenant table design and a global edge network, with no servers for you to size.
- **Your data, your region, your cloud.** One table per tenant, hosted in the UK, EU or US, or the same platform deployed into your own AWS, Google Cloud or Azure account.

## Flexibility: entities, fields and screens

Every entity in DemandFlow is a definition: its fields and their types, the panels that group them, the lookups to other entities, the grid columns, the toolbar actions, the write rules and the relationships. More than thirty field types cover text, numbers, dates, choices, multi-select, tags, user and multi-user pickers, lookups to any other entity, record references, files and images, notes, links, gauges, grids and nested object arrays. Panels can be plain forms or any of over a hundred and fifty specialised panel types, from a gantt, kanban, map, calendar or timeline to a splice diagram, rack elevation, tower layout, cash flow grid or RACI matrix.

A definition editor in the application lets an administrator change any of this for their tenant. Rename a field, change its type, add a lookup, hide a panel, add a validation rule, reorder a form, add a new panel of related records, or add a brand-new field to a shipped entity. The tenant's version of the definition replaces the shipped one, so changes survive every release.

**New entities.** Create a new entity from a form, or describe it in plain English ("a wayleave register with landowner, parcel, annual fee, renewal date and documents") and the AI drafts the entity, its fields, lookups and panels for review. New entities get list screens, record editors, search, history, automations, import, export, API and assistant access the moment they are published, and they can be placed in the menu under your own groups.

**Screens and dashboards.** Grids, forms, kanbans, gantts, calendars and charts all come from one component library, so every screen behaves the same way and every grid filters, groups, sorts and exports to Excel, CSV and PDF. A dashboard designer builds KPI, chart and grid dashboards from any entity or dataset, with templates to start from. Workbooks add SQL, charts, KPI tiles and plain-English questions over the same records, with governed edits back.

**Rules.** Write rules on an entity validate and refuse a save that breaks them, number sequences issue references per company, and approvals route by organisation and spending authority.

## Automation

An automation watches an entity for a create, update, delete or approval and runs when its filter matches: a status change, a date passing, a value crossing a threshold, a field left empty. Its actions update the record or a related one, calculate a count of child records, create a new record with fixed or copied values, send an email to an address, a user or the person in a field, start an approval workflow with its own emails and links, run an integration flow with the before and after state, or create a company, contact or deal in HubSpot. Several actions chain in one automation, and a scheduler runs the time-based ones every five minutes, every fifteen minutes or daily.

Modules ship with their own automations ready to switch on: warranty and renewal reminders, overdue maintenance, retirement notices, ticket lifecycle emails, change approvals, status change history, review-due knowledge articles, support agreement expiries and many more. They are editable and versioned, so an upgrade keeps the copy you customised. Provisioning automations set up whole modules in one step. Every run is recorded, and a failed step is visible rather than silent.

## Integration flows and webhooks

Integration flows are built in the application, step by step, without middleware. A flow starts from an inbound webhook, a schedule, a manual run, an inbound email or an automation. Its steps fetch from or post and patch to any external system, read, create, update, look up and list DemandFlow records, map fields between shapes, branch on conditions, loop over lists, and post to Slack or send email. Secrets such as API keys are held in a vault

**Outbound webhooks.** Subscribe an endpoint to an entity's changes and receive signed, timestamped payloads with a deduplication header, timeout and retry, carrying the changed fields or the full record. Payloads come in DemandFlow's own format or as TM Forum 639 resource and TM Forum 638 service events for telecom integration.

under their own rotating key and referenced by name, so they never appear in a flow definition or a log. Every run keeps its steps, inputs and outputs for inspection.

Inbound webhooks give each flow its own address per tenant, with none, header or bearer authentication, so an external system can hand DemandFlow an event and have the flow turn it into records, tickets, changes or notifications.

**Ready-made connections.** Microsoft 365 for mail, calendar, files, users and Intune devices; Google calendar; Microsoft Entra single sign-on; HubSpot; Dynamics 365 with in-app import and inbound events; Xero; IP Fabric and NetBox for network discovery; AWS, Azure and Google Cloud network reads; Power IQ, EcoStruxure IT and EkkoSense telemetry; Slack and Teams; Stripe for product sites; Idibu, Indeed and Monster for recruitment; and a Power BI feed.

**Imports.** Generic CSV import with column mapping, a dry run and an error report; a bulk endpoint; and module importers for ledgers, purchase orders, patents, recruitment, fibre, data centre records, device libraries, site plans, RAN OSS exports and MS Project.

## REST API

Everything in DemandFlow is reachable through a versioned REST API. Get, create, update and delete any record by id; stream a whole entity as newline-delimited JSON; run several queries in one call; and read public files. Any entity you add is on the API the moment you publish it. Requests authenticate with personal access tokens that users create and revoke themselves.

A write through the API is a write through the platform: the same write rules validate it, the same automations fire, the same history is kept, the same webhooks are delivered and the same live updates reach open screens. An integration can never bypass the governance your administrators set.

Beside the REST API sit the TM Forum read APIs for resources and services, the OData feed that Power BI and Excel read directly, the SIEM ingest endpoint for events, the IT agent and network collector endpoints, the careers site and support portal endpoints, and the e-signature and customer presentation portals. Each uses the same tenant model, so a token or key only ever reaches its own tenant's data.

## MCP: built for AI assistants

DemandFlow runs a Model Context Protocol server, so Claude and other MCP-capable assistants can work with your operation directly. Core tools list entities and read their definitions, fetch records by id or reference, batch read, list, search and run multi-queries, walk a record's cascade and parent, and create and update records. RAN tools return a site or a cell as one bundle. Estate and service graph tools answer capacity questions, find where equipment fits, trace a path, compute blast radius and dependants, run a failure what-if, allocate addresses, compare discovery snapshots and draft plan scenarios and changes. Portfolio tools give summaries and headroom, create demands and scenarios, compare them, assess plan impact and draft status reports.

An assistant can therefore answer "which customers lose service if this duct is cut", "where can six 8 kW racks go in COL1", "what is the headroom in Packet Core for 2027" or "draft the status report for Northern Trunk Main" from the live records, and leave a draft for a person to approve. Drafting tools create drafts only, there is no delete tool, every call is authenticated by a personal token or OAuth 2.0 and audited, and the assistant only ever sees the tenant its token belongs to. The same interface powers the Ask cell in Workbooks and the AI features inside the modules.

## Scale, residency and security

The application is delivered from a global edge network and the back end is serverless: hundreds of functions that scale with demand behind a managed API, on-demand database capacity, a dedicated search cluster and a per-tenant table design with no shared business data. There are no servers to size and no maintenance windows for upgrades. Demo tenants run tens of thousands of plant records, thousands of devices, a hundred thousand ports and millions of log events on the same design.

Each tenant's records live in their own database table, with separate log and sensing tables and their own private, public and archive buckets. The tenant is resolved from the signed-in user, never from the request. Tenants are hosted in the UK, EU or US region of your choice, and for customers who need their own cloud the same platform deploys into your own AWS account or as a container into Google Cloud or Azure, with one database per tenant.

Sign-in uses short-lived signed tokens and rotating refresh tokens in secure cookies, two-factor authentication enforced per tenant, Microsoft Entra single sign-on with group-to-role mapping, new-location fingerprinting and one-time reset links. Every create, update and delete keeps a before-and-after history with the actor, taken from the database's own change stream so nothing bypasses it. Secrets sit in a vault keyed per tenant, data is encrypted at rest and in transit, a web application firewall fronts the application, and rate limits protect public endpoints. A supplier assurance pack is available for procurement and vendor risk reviews.

## Summary of specifications

| Item | Detail |
|---|---|
| **Definitions** | Over 1,000 shipped entities; per-tenant definitions replace shipped ones and survive releases; in-app definition editor; new entities from a form or an AI-drafted description; custom menu groups |
| **Fields and panels** | 30+ field types including lookups, references, multi-user, tags, files, images, notes, gauges, grids and object arrays; 150+ panel types including gantt, kanban, map, calendar, timeline and module-specific diagrams |
| **Rules** | Write rules that validate and refuse; number sequences per company; approvals by organisation and spending authority; record versioning with conflict refusal |
| **Automations** | Triggers create, update, delete, approval with filters; actions update property or count, update related, create or update object, email address, user or field, approval workflow, trigger flow, HubSpot create; scheduler 5 min, 15 min, daily; shipped module automations, editable and versioned |
| **Integration flows** | Triggers webhook, schedule, manual, inbound email, automation; steps fetch, post, patch, DemandFlow create, update, lookup, list, map, branch, loop, Slack, email; secrets by name from a vault; run history |
| **Webhooks** | Outbound in DemandFlow, TMF639, TMF638 formats, signed and timestamped, dedupe, timeout and retry; inbound per flow with none, header or bearer |
| **REST** | Versioned API: objects get, post, patch, delete; entity streams; multi-query; public files; PAT bearer; same rules, automations, history and webhooks as the app; TM Forum read APIs; OData feed for Power BI |
| **MCP** | Core tools (entities, definitions, records, batch, list, search, cascade, parent, multi-query, create, update); RAN site and cell; estate graph (capacity, fit finder, path trace, blast radius, dependants, what-if, allocate, compare snapshots, draft scenario, draft change); portfolio (summary, headroom, demand, scenario, compare, plan impact, status report draft); PAT or OAuth 2.0; audited; drafts only |

| Item | Detail |
| --- | --- |
| **Connections** | Microsoft 365, Google calendar, Entra SSO, HubSpot, Dynamics 365, Xero, IP Fabric, NetBox, AWS, Azure, Google Cloud, Power IQ, EcoStruxure IT, EkkoSense, Slack, Teams, Stripe, Idibu, Indeed, Monster, Power BI |
| **Imports and exports** | CSV with mapping, dry run and error report; bulk endpoint; module importers; Excel, CSV and PDF on every grid; server-side PDF rendering |
| **Scale and hosting** | Serverless functions behind a managed API; on-demand database capacity; one table per tenant; dedicated search cluster; global edge delivery; tenants in UK, EU or US regions; deployment into your own AWS account or as a container on Google Cloud Run and Azure Container Apps |
| **Security** | Signed tokens, rotating refresh tokens, per-tenant 2FA, Entra SSO, new-location fingerprinting, one-time reset links; before-and-after history from the change stream; vault secrets keyed per tenant; encryption at rest and in transit; WAF; rate limits; assurance pack on request |

**www.demandflow.com**
Book a demo: www.demandflow.com/

DemandFlow® is a registered trademark. Registered in England No. 09950859. 71-75 Shelton St, London, WC2H 9JQ. All rights reserved.