


FLOWSENSE · THE ALL IN ONE PLATFORM FOR MONITORING YOUR SEWER NETWORK

See the whole network. Find the water that should not be in it.

FlowSense turns telemetry from any logger fleet into an operations picture your team can act on today. It turns the same data into the engineering evidence you need to justify capital spend. One continuous chain, in one system, so nobody re-keys a number between stages.



Network Overview. It opens on what needs attention: one row per site, beside a map that stays in view. Six figures sit across the top, with alarms to acknowledge kept apart from any still breaching. The Overview refreshes every sixty seconds. Screens throughout this brochure are the running platform, photographed on a demonstration network.

Monitor

The fleet, live. Alarms, blockage early warning, battery and comms health.

Analyse

Infiltration and inflow (I/I) in kilolitres, dry weather flow, hydraulics, data quality.

Engineer

Calibrated models, the I/I heat map, and the cost per kilolitre of removing it.

WHAT IT IS

Three products in one, and one record underneath them.

Each layer below is normally bought separately, from a different supplier, with a spreadsheet in between. In FlowSense they share one database, one set of units and one timeline. That is the whole idea. The number in the board paper and the reading in the chamber are the same number. You can trace one to the other without leaving the browser.

Monitor

Live fleet status, site detail, alarms, telemetry and battery health, rainfall context, and early warning of blockages weeks before a pipe surcharges.

REPLACES THE LOGGER VENDOR'S PORTAL

Analyse


Rainfall derived infiltration and inflow in kilolitres, dry weather flow profiled, depth and velocity hydraulics, catchment flow balance and data quality scores.

REPLACES SPREADSHEETS AND MANUAL I&I STUDIES

Engineer

Project workspaces for calibrated rainfall response (RTK) models, a dynamic wave network model, the I/I heat map, and the cost per kilolitre of infiltration removed.

REPLACES A MODELLING CONSULTANCY ENGAGEMENT



The network map. Every site sits on aerial imagery, ringed amber while an alarm waits to be acknowledged and red while one is breaching. The Layers panel narrows the map to what is wrong: alarms, hydraulic risk, comms failures, low battery, stale telemetry or anomalies. It also names the layers that are at zero. The Weather layer adds rain radar or satellite cloud. Select a site to open its workspace.

DELIVERY

Cloud hosted and browser based. Desktop, tablet and mobile, nothing to install.

INSTRUMENTATION

Vendor independent. Levels, flows, pump runs, rain, water quality, power: whatever a site measures.

DATA IN


Loggers collected by EDS, five ways in from your own SCADA, and reviewed CSV import.

DATA OUT


Branded PDF and CSV, scheduled SFTP, FTPS and MQTT feeds, SWMM and GeoJSON.

The failing logger is named before it goes quiet.

A monitoring programme fails quietly. A battery flattens, a modem stops answering, and nobody finds out until the quarter's report is already wrong. FlowSense monitors the fleet itself, and puts what needs attention first.




Fleet Health. Every logger on the account, with what needs attention first. Battery voltage is judged for its chemistry, and each battery gets a replace by date projected from its own discharge trend. Each logger also shows how many of the last thirty days it delivered. A dying logger becomes a planned visit rather than a hole in the record.



Alarms. Every rule that has tripped, unanswered first. Alarms to acknowledge are kept apart from any still breaching, so a cleared alarm reads as a task rather than a fault. There are rules for thresholds, for no data and for pump condition, and alerts go out by email and SMS.

Every site, judged against its own usual

The Sites list, each site's page and the sites you pin to the Overview compare the latest hour with that site's usual pattern. The pattern is learned from the site's last four weeks, for that day and time. They say plainly when a site is running above or below it. Departures under 10% are not called out. A site whose newest reading is more than a day old gets no verdict rather than a guess. Rain in the last day is shown beside the verdict, because weather is not a fault.




On a phone. What needs attention first, and a tab bar for the rooms used most.

Alarms are edge triggered: you are told when the state changes, not once a cycle for as long as it lasts. A threshold rule deliberately ignores stale data, so a dead logger frozen on a breaching value cannot hold an alarm open. Catching a dead logger is the job of the no data rule, which is a separate rule on purpose.

Whatever a site measures, shown in the form that suits it.

A gravity sewer, a pump station, a trade waste discharge and an overflow chamber each open on one page. The page is laid out for what that site carries: levels, flows, pump runs, pH, conductivity, rain, power. Every chart zooms, pans and exports, and its window travels in the link you send a colleague.



The site page. The verdict against usual sits beside the site's name, with any alarm to acknowledge below it. Then comes a tile for each measurement, set against this time yesterday. Here they are flow, depth, velocity and temperature, plus the daily volume, with the hours of data it is based on, and the logger. A pump station's pumps share one status tile, showing running and starts. A tab the site has no data for is not shown.




A month on one time axis. Each measurement has its own pane, captioned by channel. The storm of 13 September lifts flow to 254 L/s. Depth and velocity rise with it, and the alarm it raised is flagged on the trace. Pump runs and alarm contacts draw as lanes on the same axis. Every trace breaks where the logger went quiet, rather than drawing across the gap.

The daily volume tile says how it was worked out and how many hours of the window it is based on. An engineer checks that first. Each local day's volume is its average flow multiplied by 86,400 seconds, the way a water authority's own spreadsheet does it. The tile links to the working.

Weeks of warning, read from the pipe's own overnight behaviour.

A gravity sewer drains almost empty in the small hours. An obstruction shows up in that overnight trough long before it shows up anywhere else, and long before anyone rings the council. Blockage Watch reads that trough. It runs on every gravity sewer you monitor, without being asked.



Blockage Watch. One summary row sits above the watchlist. Sites needing a decision come first, each with its score, mechanism and case. Clear sites are folded below, with the pump stations that the Pump Station Manager (PSM) watches. "We looked and it is fine" is also an answer. The selected site opens as a case file. Here its score has eased from 44 at the last run to 28 now, so the case is watched, not closed.

WHAT IT READS

Six signatures come from the level trace. They are overnight minimum elevation, its trend, amplitude collapse, departure from the site's own learned day and time profile, flat topping and step change. A site that also carries velocity gets two more. The strongest is **conveyance at matched depth**, which compares recent depth and velocity pairs with the baseline, bin by bin.

WHY THAT ONE MATTERS

Both a restriction and a busier catchment lift the overnight trough. A fall in velocity at the same depth measures the pipe, not what arrives through it, so it tells them apart. Only a restriction moves less water at a higher level, and only a restriction is worth a jetting crew.

Rain produces every one of those signatures more strongly, so the analysis runs on **dry days only**. A site whose elevation appears only after rain scores zero and says so. A fortnight too wet to judge is reported as insufficient dry weather, not as normal.

WHEN IT RUNS

Every gravity sewer, every twelve hours. It takes the overnight minimum between 01:00 and 05:00 and compares the latest fortnight with the rest of the last 75 days.

A DRY DAY


Less than 2 mm of rain on the day, and under 2 mm of the last three days' rain still draining. That leaves out a storm's recession.

BANDS

Watch from 35, warning from 55, critical from 75. A case opens at watch and you are told at warning.

A mechanism and an action, not a red dot on a map.

An early warning that only says "something is wrong at this site" moves the problem rather than solving it. Every verdict here names what the evidence says is happening and what to do about it.



One case file, read top down. It opens with the verdict on its score ladder, with the mechanism and the action inside it. The key numbers follow, then the case, then the evidence: every night behind the verdict, including those set aside and why. Here the overnight level is up but velocity rose with it. That means increased load, not an obstruction, and no jetting.

THE SIX VERDICTS

- **Silt or grease:** desilt the reach.
- **Sudden obstruction:** send CCTV, because a step change is debris or a partial collapse.
- **Downstream restriction:** inspect below the meter.
- **Increased load, not an obstruction:** the verdict that saves the money. The overnight level is up because more water is arriving, and there is nothing in the pipe to clear.
- **More flow overnight,** at a site with flow and no level: check the level before sending a crew. Flow alone cannot tell extra water from a restriction.
- **Mechanism unclear:** inspect before committing a crew.

A suspect channel goes to data quality rather than to a crew. A drifting instrument is not a blockage and should not cost a truck roll.

MEASURED ON YOUR NETWORK

Crossing the threshold opens a case with an assignee. The case closes on a recorded outcome. Detection performance is computed from those outcomes, so the hit rate you see is your own.

What it does not do

Blockage Watch is a statistical early warning read from each sensor's own history. It says the pipe reaching that sensor is behaving differently for a reason rain does not explain. It does not locate the obstruction, and it cannot see one downstream of the meter. A verdict needs at least four dry days in the latest fortnight and eight in the weeks before it. The velocity signatures need a working velocity channel. A pump station's wet well is left to PSM. We would rather you knew that now than discovered it later.

TESTED AGAINST YOUR OWN HISTORY


A one click **backtest** replays the detector against the overflows already in your register, using only the telemetry that existed before each one. It is the fairest test we can offer. You can run it on day one, against history you already hold.

Infiltration and inflow, quantified rather than described.

"This catchment is wet" is not a business case. A volume, a peak, a lag and a percentage over dry weather flow is. FlowSense separates the storm response from the sanitary flow underneath it. It reports what the difference is worth, per catchment and per storm.

Peak RDII 236.7 L/s	RDII volume 11,710.6 kL	Over dry weather flow 26.13 %	Storm lag 1.8 h	Dry weather flow, average 17.29 L/s	Night baseline 6.77 L/s
----------------------------------	--------------------------------------	--	------------------------------	--	--------------------------------------

What the storm cost, in numbers. The peak inflow, the volume in kilolitres and the percentage over dry weather flow. Beside them sit the lag from rain falling to the sewer answering, the dry weather flow itself and the night baseline.



The RDII hydrograph. Measured flow is blue, with the modelled dry weather baseline dashed green beneath it. The difference is the rainfall derived infiltration and inflow (RDII), in red. The area under the red trace is the volume that should never have entered the sewer. Every trace breaks at an outage rather than drawing across it.

THE BASELINE IS THE HARD PART

Dry weather flow is the daily sanitary pattern plus groundwater infiltration. We take it from qualifying dry days only and split it by day type, so weekend sanitary flow cannot masquerade as inflow. The groundwater part moves through the year, so it is tracked rather than held constant. A fixed baseline would quietly charge the difference to storm response.

ACROSS THE CATCHMENT

Reconciling upstream with downstream gives the mean and peak balance, the unaccounted volume and the average error. A catchment can then be closed rather than sampled. Alongside it come multi site comparison over a shared window and trade waste discharge analysis. A storm by storm split shows the wet weather uplift each catchment carries.

The number that ends the argument

The Flow Analysis report splits every catchment's flow three ways: sanitary flow, groundwater infiltration and rainfall derived inflow. The three parts sum exactly to the metered total. That turns "we think this catchment is leaking" into a kilolitre figure a treatment cost can be applied to. The capital plan later ranks its options against the same figure.

PATHWAY

Inflow dominant when the median storm lag is two hours or less, infiltration dominant at six hours or more, mixed between.

SEVERITY

An I/I score from 0 to 100, weighted across inflow, infiltration and the peak wet weather factor.

COST OF I/I


Priced from your own treatment and pumping cost per kilolitre. No default prices are assumed.

Check the instrument before you trust the number.

Every figure on the pages before this one rests on the record being sound. So the platform scores the record first and tells you which channel is the weak one.



Depth against velocity. A month of simultaneous pairs, coloured by flow. Silt, surcharge, backwater and a failing meter each leave a different shape. Here the readings hold to one rating curve. The few points at lower right, fast water at a low level, are what a reviewer should ask about. The correlation, the ranges and the sample count sit under the plot.



Site diagnostics. Three scores, each named for what it scores: operator health, telemetry and anomalies. Each channel's completeness and gaps are shown, with the I/I risk in words and checks for flatlines, impossible values and drift. Drift needs a week of readings, so a day's window says not assessed rather than guessing.

Data scrubbing and conditioning

Defect detection runs the full QARTOD test battery, not a single test. It covers gross range and negatives, climatology, Hampel spikes, rate of change, flatlines, cross channel residuals and attenuated signal. Each gap's repair is chosen by testing candidate methods against good data. Where a survey or a model needs a continuous series, depth, velocity and flow can be conditioned together.

A repair never quietly replaces a reading. **Filled gaps live only in the analysis and the export.** A correction to a recorded value is written back only by an explicit merge, which keeps the original on record. The reading the instrument produced is always there to go back to.

Manning

Colebrook-White

QARTOD flag vocabulary

Nash-Sutcliffe

Kling-Gupta

Theil-Sen


Hampel

IQR and MAD

ISO 8601 on every boundary

Severity painted onto the real network, with the confidence to read it by.

Anyone can draw a heat map. The value is in the method behind it. FlowSense offers six defensible methods. It attributes severity along the gravity network by catchment and flow direction, and tells you which method applies.



The I/I heat map. Severity is painted onto the sewer network over aerial imagery, with a legend giving the kilometres of pipe in each band. Key figures run across the top, and the findings are written in words beside the map. The inspection priority list ranks the pipes most likely to be sources, each with the confidence behind its score.

- | | | |
|--|--|---|
| <p>1 Network attribution
Observed severity per gauge, attributed along the gravity network</p> <p>4 Depth and velocity
Pipe condition read from the rating curve</p> | <p>2 RDII event analysis
Storm by storm response, measured rather than inferred</p> <p>5 Mass balance
Upstream against downstream, closing the catchment</p> | <p>3 Hydraulic capacity
Where the network is closest to running out of room</p> <p>6 RTK unit hydrograph
Counts only flow that rainfall explains, and only where the fit is good enough</p> |
|--|--|---|

Kriging, the default way of estimating between gauges, carries a **0 to 1 confidence overlay**. The map shows the estimate and how much of it to trust: pipes far from a gauge fade rather than pretending to a number. Flows in upstream pipes are calibrated estimates from the downstream meter, not measurements, and the interface labels them that way.

SIMULATED STORM

The same map driven by a catchment model's response to a design storm, with a time slider.

INTERPOLATION

Ordinary kriging fitted by REML by default, with inverse distance and radial basis surfaces to compare it against.

ASSET REGISTER AUDIT


Checks your register against itself and against the flow, and writes nothing back to your GIS.

EXPORTS

GeoJSON, CSV tables, a PNG or JPG figure, an A4 PDF assessment report, or all of it in one ZIP.

Every number carries the working that produced it.

Modelling is where a platform is usually asked to be taken on trust. This one is built the other way. Each figure is shown beside what it means, what it can and cannot be judged against, and how far it can be pushed.



A rainfall response (RTK) calibration, scored. One header holds the four tools, the project and scenario, and the workflow. A line under the fit score (the Nash-Sutcliffe efficiency, NSE) gives the **best score available on this record**. Observed inflow is metered flow less a modelled baseline, so it carries that baseline's noise, and no parameter set can beat it. Judging the fit against 1.0 instead would make an honest calibration look like a failure. Below the scores, the fit states what it had to work with: here two storms, where three are wanted.


A single score over the whole series can look respectable while the model misses every storm peak, because dry stretches dominate it. So every calibration and validation scores **each storm separately**, against the bands consulting practice verifies to. Simulated peak must be within 15% below to 25% above observed, and event volume within 10% below to 20% above. Peak timing is reported against a one hour tolerance. An event passes on peak and volume together, and a run needs at least three assessable storms before it may be called verified.

THE WORKSTATION AND ITS TOOLS, ORGANISED BY PROJECT AND SCENARIO

I/I Heat Map Severity painted onto the real network	RTK Calibration SWMM compatible unit hydrograph per catchment	RTK Validation Independent storms, per storm scorecard	Hydraulics Dynamic wave model, pumps, gates, control rules
Scattergraph Silt, surcharge, backwater, meter fault	Pump Station Volumetric testing and performance decay	Overflow Forecast Design or forecast storm to exceedance	Capital Plan Net present value, benefit cost ratio, cost per kL of I/I removed
Compliance Licence limits and an overflow register	Catchment Model Multi site simulation with routing	DWF Dry weather calibration and daily pattern profiling	Events Rainfall and response event detection
Data Scrubbing Defect detection with a method per defect	QA / QC Structured review and sign off	Compare Scenario against scenario	Export Excel, GeoJSON, CSV, or EPA SWMM 5

Station flow, and pump condition, from run times you already have.

Most pump stations are never going to get a flow meter. They do not need one. The wet well fills as sewage arrives and empties as the pumps move it. So its level and the pump run record show what arrived at every timestamp, what was pumped away, and the condition of the pumps.



Pump Station Manager. One station fills the full width, with its figures first: sewage arriving and pumped out, and their balance over whole pump cycles. Below them the wet well is drawn to scale, with its set-points and the kilolitres each cycle moves. A briefing says what is worth watching, and each pump's delivery against its rating shows on a dial.

PUMP DELIVERY, MEASURED NOT RATED

Each pump's delivery is the median of its measured days over the period you choose, thirty days by default. It is set against the pump's rating, or against its own first measured month where no rating is held. A falling figure points to an impeller wearing or a rising main starting to block, weeks before the station struggles.

Pumps sharing a rising main are never judged against their ratings added up, since together they deliver less.

Pump condition, and the time each pump has left

On dry days each pump's output is held against its own reference. A slow decline, a sudden drop of 10% or more over three days, and a stopped pump each raise a finding and an alarm. After six months, remaining useful life is given as a range to a stated threshold (ISO 13381-1). Causes are ranked, never diagnosed. FlowSense reads hydraulic output, not bearings or seals.

ENERGY, AND THE ARITHMETIC BEHIND IT

Where a station has a power tag, compare kilowatt hours per megalitre between stations and over time. More energy for the same water means lost condition, reported with the hours and the draw behind it.


Cycles the record cannot support are set aside and counted, never averaged in. With the wet well and overflow configured against a level sensor, FlowSense derives overflows from the station's own behaviour.

The report your board reads, generated from the same record.

Seven templates, each written for a different reader. **Flow Totals** suits operations and billing, and **Executive Summary** suits management. **Flow Analysis** covers inflow and infiltration, and **Engineering Diagnostic** serves technical teams. **Site Investigation** is a field ready pack, **Network Performance** is for asset planners, and **Pump Station Performance** covers every station PSM runs.



A generated report. It opens with headline figures, the network's daily volume stacked by site, and telemetry health for every site. The cost of infiltration and inflow, the data availability record, key risk indicators and alarms follow. It exports as a branded PDF or a CSV bundle.



The builder. Template, sections, sites and period, previewed before anything is exported.

BUILT FOR YOUR AUDIENCE

You can select several sites, set any date range and add a comparison period, such as this quarter against last. Sections can be included or excluded, so one template serves several readers. Your logo and report title appear throughout.

SCHEDULED FEEDS OUT

Data is delivered automatically into your systems, over SFTP, FTPS or MQTT, as CSV files or JSON messages. Each feed has its own channel selection, a connection test, a payload preview and a run history. A feed marks its place only after a fully successful send, so a failed delivery cannot silently skip a window.

The round trip

Your SCADA feeds FlowSense, and FlowSense feeds your SCADA, historian, business intelligence tools or data pipelines. Very few products in this category accept the inbound half. Once a tag is mapped, its readings are ordinary channel data, with no separate view to learn. A read API, webhooks, and model runs exported as Excel, GeoJSON or an EPA SWMM 5 input file complete the loop.

REPORTS

PDF, CSV or both, on demand or scheduled daily, weekly or monthly by email, with your logo.

API

Sites, channels, readings and alarms by key. A key can be held to chosen sites, and never sees more than its owner.

WEBHOOKS

Alarm raised and cleared, forecast risk and blockage case events, signed with HMAC SHA-256 and retried.

Every way in, one validated record.

However a reading travelled, the same rules apply to it once it arrives. That is what makes a figure comparable across a network assembled from several sources over several years.

Automatic collection	Readings arrive on their own from your instrument fleet, on a schedule set per site. EDS operates and monitors that collection as part of the subscription, and checks every day that the platform holds everything the source does. Instruments from several manufacturers are supported side by side. New site types are added without a platform change, so the network you buy next still lands in the same record.
Your own SCADA	Your SCADA can send readings to us, or we can collect them. There are five ways in. HTTP push goes to our endpoint, and a file drop lands in an SFTP folder we host. We subscribe to your MQTT broker, or read your system by API polling with an API key or OAuth2 client credentials. Or we collect files by file pickup from your SFTP or FTPS server. Guided setup sends your IT team a form that needs no login. Nothing is installed at your end, and with a push or a file drop you open nothing inbound at all. Unmapped tags land in a review queue. They are counted and discarded until an operator maps them, so nothing a SCADA sends can create a channel by itself. A token used to send data can only write. It cannot read a single stored value.
Reviewed CSV import	Imports are staged and reviewed before they are committed, against a chart of old readings and new. A file in the wrong time zone is caught before it lands. You can replace a range or the whole history, or fill only what is missing. The time zone is chosen at upload, day and month order is detected or set by hand, and imports can be undone for 30 days.
Manual sites	History that never had a logger. You create a 2019 flow survey, or a meter since removed, by hand, place it on the map and give it channels. Every heat map, model and volume total then reads it exactly as it reads a live site. Nothing polls it, so it never counts as an outage.
Rainfall	Every geocoded site gets modelled rainfall for its own grid cell, from the ECMWF IFS forecast model via Open-Meteo, filed in the hour it fell. That way "did it rain there" is always answerable. Your own gauges, uploaded or reporting on your telemetry, take precedence, one source per window, so measured and modelled rain are never mixed. No rain is invented for an hour nobody reported.
Integrity rules	Field timestamps are kept, never the time of collection. A record without a trustworthy time is skipped rather than guessed. Synthetic "as of now" readings and future dated rows are rejected. Duplicates are dropped on channel and timestamp, so any file, message or batch can be re-sent safely. Writes are all or nothing, and every collection is logged.
Outage recovery	For loggers EDS collects, the Data Recovery tool compares the platform with the source, reaching back as far as 60 days. It fills the windows where a site went quiet. An outage no longer leaves a hole that surfaces only when a report is already wrong.
Time	Stored in UTC and shown on each reader's own clock, so a team across two states reads one record. Alarms use the site's own time zone, and daily totals are cut on the site's local day.

HTTP PUSH

Up to 10,000 readings a request, every row reported back.

MQTT

JSON or Sparkplug B, with optional TLS and the broker's certificate checked.

API POLLING

JSON or CSV over HTTPS, every 1 to 1,440 minutes.

FILES

CSV collected every 15 minutes by default, or dropped into our folder.

Offline means **stopped delivering**, judged against each site's own upload cadence. The limit is twice its usual gap between deliveries, never less than a day and never more than a week. A meter that uploads twice a day and one that uploads every fifteen minutes are each held to their own rhythm. A single fleet wide threshold cannot do that.

Your team, your structure, managed by you.

WHO SEES WHAT

Access is role based, with an on or off switch per user for each page. The pages are the Network Overview, alarms, Blockage Watch, analytics, engineering, reports, data scrubbing, the heat map, modelling, flow surveys, data export, PSM and Trade Waste. You decide exactly what each of your people sees, down to individual sites and channels.

WITHOUT RAISING A TICKET

Account owners create users, reset passwords, suspend accounts and set access themselves. Named sub-teams mirror your own organisation, with inherited access. Individual Admin sections can be delegated one at a time, per person. EDS can require two factor authentication by role on any account.

ASSURANCE

Sign-ins, exports and administrative actions are recorded in an audit trail kept by EDS. Access is scoped per account and per user. Passwords are hashed, traffic is TLS protected, and third party credentials are stored encrypted. Tokens for sending data are write only. Data assurance keeps a monthly record of how much of each site's due data arrived.

Built to the published standards, and tested against them

In the product the list of standards sits on one page, generated from the running software. Each standard lists the calculations that actually cite it, and a standard nothing implements cannot appear on it. Conformance is tested against the code, not asserted once. Any of it can be walked through line by line against the running software.

STANDARD

WHAT IT GOVERNS HERE

CIWEM UDG Code of Practice for the Hydraulic Modelling of Sewer Systems, 2017	The per storm acceptance bands every calibration and validation is scored against, and dry weather flow as a daily pattern split by day type
WEF MOP FD-6 and WPCF MOP No. 36	Dry weather flow as the daily sanitary pattern plus groundwater infiltration, from qualifying dry days only
WSA 02 Sewerage Code of Australia	Self cleansing velocity, level of service checks, and Colebrook-White friction per pipe
AR&R 2019, Book 2	Areal reduction factors, and AR&R temporal patterns run exactly as supplied
AS/NZS 3500.3	The Rational Method for ungauged sub-catchments, and the minimum velocity band on scattergraphs
QARTOD	The quality control tests and flag vocabulary applied to every sample
EPA SWMM, including EPA/600/R-15/162A	Dynamic wave routing, and the three component RTK method for rainfall derived infiltration and inflow
NSW Liquid Trade Waste Management Guidelines, 2021	Trade waste charges: the annual fee, usage charge, excess mass charge with its escalating BOD5 rate, and the pH charge
ISO 13381-1	The terms pump condition uses: remaining useful life, counted down to a stated threshold and given with its uncertainty

FlowSense is built to the standards listed and tested against them. The instruments EDS supplies carry certification in their own right. The Detectronic LIDoTT range is MCERTS certified for flow metering, and rated ATEX and IECEx Zone 0 for hazardous atmospheres. A certified meter keeps its certification when its readings arrive in FlowSense, which stores the field time and the raw value the meter produced.

WHAT IT COSTS

One subscription. Every capability. Priced by the site.

Published here, because a price you can only discover through a sales cycle is itself a cost. There are no editions, no modules and no per user fees. Every capability in this brochure is in the one subscription. The only number that moves is how many sites you monitor.

FIRST 20 SITES \$30 per site per month	SITES 21 AND BEYOND \$20 per site per month
--	---

AUD excluding GST, billed annually on twelve month terms. Bands are cumulative, so a fifty site network pays the first rate on twenty sites and the second on thirty. Growing the network never re-prices the sites you already have. A site is a monitored location with at least one active channel. Archived sites cost nothing. Cancel at renewal, and your data exports with you in full. The only pass through is SMS alerting at carrier cost.

USERS	ADDED MID TERM	AT RENEWAL
Unlimited, with no per user fee and no editions.	A new site is charged at the rate of the band it falls in, for the days left in the term.	The rate card applies for the current term and may change at renewal, on 60 days written notice. Either side may end the subscription at renewal on 30 days written notice.

WHAT THAT MEANS FOR A REAL NETWORK	SITES	PER YEAR	PER MONTH	BLENDED, PER SITE PER MONTH
Small shire	5	\$1,800	\$150	\$30
Regional council	20	\$7,200	\$600	\$30
Mid size council	50	\$14,400	\$1,200	\$24
Large regional network	90	\$24,000	\$2,000	\$22
Bulk utility	300	\$74,400	\$6,200	\$21

The bands are set deliberately against procurement reality. For most Queensland councils, a fifty site network sits under the \$15,000 small contractual arrangement threshold. That means a single quote and a delegated signature rather than a procurement round.

Getting started

STEP	WHAT IT COVERS	PRICE
Free trial	90 days, the whole platform, running on your own uploaded flow history. Upload last winter's CSV exports and see your infiltration ranking the same afternoon. No hardware required.	\$0
Assisted onboarding	GIS network import, site and channel configuration, alarm setup, commissioning any SCADA connection you want feeding the platform, and training.	\$2,500 once waived at 50 sites and above
Guided calibration, optional	The calibration capability is included in every subscription and you can run it yourself. This optional service has EDS do the work with you. It is delivered in the platform, where you can re-run it, not as a PDF.	\$7,500 per catchment

Onboarding and guided calibration are one-off services. The subscription, and SMS at carrier cost, are the only recurring charges.

Where the boundaries sit.

Stated here rather than discovered later, because an engineering buyer will ask. The answers are good ones.

GIS network	We integrate with your GIS. Connecting a council's own network layer is a scoped integration task at onboarding, not an out of the box switch.
Network model	The dynamic wave solver is one dimensional: gravity network, pumps, gates and control rules. It does not do 2D overland flood mapping, water quality or sediment transport. A run exports as a SWMM input file, so work started here stays portable where those layers are needed. Preparing and running a large network takes minutes, not seconds.
Capacity screening	The screening table is steady uniform flow from the recorded pipe geometry, and is only as good as the GIS attributes behind it. Dynamic behaviour is the network model's job.
Model verification	The scorecard reports how the model performed against the storms in that window, to a published standard. It is not a promise about the next storm. A window with fewer than three assessable storms is reported as insufficient events rather than passed.
Overflow forecast	As good as the rainfall forecast fed into it. It is a risk ranking tool, not a guarantee.

EDS HAS MEASURED WATER AND WASTEWATER FOR





Logos are of organisations EDS has supplied, monitored or worked for since 1991; they are the trademarks of their owners and do not indicate use of FlowSense.

NEXT STEP

See it running against a real network, not a slide.

We will demonstrate FlowSense live, including the blockage backtest replayed against a real overflow register. We will show you the working behind every number on the screen, so bring your hardest question about how a figure was calculated. To test it on your own data first, send us the history you already have. We will stand up a 90 day trial on it.



Every drop, measured.

Environmental Data Services

Australian owned and operated since 1991
13/20-22 Ellerslie Road, Meadowbrook QLD 4131
Offices in Brisbane, Richmond, Boronia and Adelaide

1300 721 683

eds@e-d-s.com.au
www.e-d-s.com.au
edsflowsense.au

EDS FlowSense platform brochure, reference EDS-FS-BRO-01, revision 2.6, issued 12 October 2026 in place of revision 2.5 of 12 October 2026. Capabilities described are those in the released platform at the date of issue. Every screen reproduced in this document was photographed from the running platform on a demonstration network whose sites, telemetry and pipe network are synthetic; no client's data appears in it. Prices are AUD excluding GST and are the founding rates current at the date of issue. Third party product and vendor names are the trademarks of their respective owners. Environmental Data Services, ABN 31 729 813 151.