

Robert Bridge.

Selected projects from a decade of operational and analytical work, across financial services, infrastructure, and operations. Not a comprehensive list; the work I consider most representative of what I do.

DISCIPLINE	PERIOD	PROJECTS	EDITION
Operations & Analytics	2014-2026	11	May 2026



Robert Bridge
Project Portfolio
MMXXVI

Contents.

Eleven projects organised by employer, presented in the order they appear in this portfolio.

I	BNY Pershing	AUG 2020 – NOV 2025
01	Market Claims Workflow Redesign	p. 03
02	Division-Wide KPI Dashboard	p. 04
03	SharePoint Hub & Asset Service Hub	p. 05
04	Power Platform Upskilling Programme	p. 06
II	Big Atom	AUG 2019 – JUN 2020
05	End-to-End Operational Data Pipeline & Reporting Build	p. 08
III	McDonald's Canada	MAY 2017 – APR 2019
06	Cash Discrepancy Investigation Process	p. 10
07	Cash Management & Banking Procedures Overhaul	p. 11
IV	Remittance 360	NOV 2025 – JAN 2026
08	Transaction-Level Reconciliation Tool	p. 13
09	Monthly Fee Management Tool	p. 14
V	Amey	DEC 2014 – APR 2017
10	Fuel Efficiency & Theft Detection	p. 16
11	Contract P&L Reporting Suite	p. 17

BNY *Pershing*.

AUGUST 2020 – NOVEMBER 2025 · LIVERPOOL, UK

Five years across asset services operations. Beginning with a single team's claims workflow and expanding into divisional reporting, a SharePoint hub serving the wider division, and a structured upskilling programme that left Power Platform capability embedded in every team.

01	Market Claims Workflow Redesign	p. 03
02	Division-Wide KPI Dashboard	p. 04
03	SharePoint Hub & Asset Service Hub	p. 05
04	Power Platform Upskilling Programme	p. 06

01 — WORKFLOW REDESIGN · FIXED INCOME & EQUITIES

Market Claims *Workflow Redesign.*

Reconstructing a fragmented manual exception process around a structured reconciliation template and a four-team production line, moving the team from partial visibility to a fully reportable pipeline.

182 → 97 <small>AVG. CLAIM AGE, DAYS</small>	40 → 95% <small>RECONCILIATION RATE</small>	~2,000 <small>LIVE CLAIMS AT ANY TIME</small>	4 <small>SPECIALIST TEAMS (PRODUCTION LINE)</small>
--	---	---	---

The market claims team reconciled exceptions arising from entitlement to coupon interest payments on fixed income securities and dividend payments on equity securities. The core issue is one of entitlement: when a bond changes hands after record date but before pay date, the seller receives the coupon from the registrar but is not entitled to keep it; the buyer paid for it as part of the purchase price via accrued interest. Because BNY operated at institutional scale, BNY typically sat in the middle of a chain, claiming from the counterparty who sold to them whilst forwarding to the counterparty they sold to.

The workflow was entirely manual and fragmented. A daily macro-run Excel file pulled all exceptions and assigned them by stock ticker letter to individual analysts, each managing 300–400 exceptions. Analysts worked largest credits and oldest claims first, leaving a significant middle band unworked and ageing. Tracking depended on freehand comments, so if an analyst was absent, their reconciliations were uninterpretable to colleagues.

STAGE ONE: THE RECONCILIATION TEMPLATE

A structured Excel workbook was designed and made mandatory across the team, using tabs as stages of the claim. **Tab 1** accepted event and exception data once. **Tab 2** auto-cross-checked trade dates against record and pay dates, validated quantities, identified counterparty codes, and calculated the claimable amount. **Tab 3** generated a pre-populated email, with body, addresses and SSI details auto-populated from counterparty codes and client. **Tab 4** was a replica of the formal claim letter, exportable as PDF and uploadable directly to the digital signature system, removing print-sign-scan entirely.

STAGE TWO: THE PRODUCTION LINE

Each analyst owned their claims end-to-end, reconciling, chasing, negotiating, escalating and settling whilst handling inbound claims from counterparties. The constant context-switching made prioritisation impossible. A production line approach was implemented: Reconcile handled new exceptions and sent the first claim. Contact managed correspondence and negotiation. Payment handled fund movements and settlement. Escalate handled complex claims and disputes.

Handover was automated via email subject codes, with **CT** routed to Contact, **PT** to Payment and **ET** to Escalate. Server rules categorising replies automatically. A structured status column (“Pending tax documents”, “Awaiting counterparty confirmation”) made the pipeline reportable for the first time and fed directly into the divisional KPI dashboard.

OUTCOME: POST 60-LIVE

On new exceptions, the reconciliation rate rose to 95%, giving the team full sight of liabilities for the first time. Daily new exceptions ran at 10–15, manageable within the production line. Average claim age fell from 182 to 97 days. Legacy exceptions ran in parallel until cleared.

Division-Wide *KPI Dashboard*.

Replacing a constellation of legacy aged packs with a single Power BI dashboard, refreshed every 15 minutes, with row-level security across six tiers and metric-level data freshness indicators.

15 min <small>REFRESH CADENCE</small>	6 <small>RLS TIERS</small>	1 <small>CONSOLIDATED DIVISIONAL VIEW</small>	∞ → 1 <small>REPORTS → SINGLE SOURCE</small>
---	--------------------------------------	---	--

Divisional performance reporting was handled by an internal MI team producing aged packs in legacy `.xls` format. A separate report per metric, each stored in its own location, each updated daily or weekly through a combination of old macros and manually-maintained lookup tables. The lookup tables themselves were inconsistent: some records keyed by employee name, some by ID, some by initials, some by email. To review their own performance numbers, an analyst had to locate the latest version of each report, understand how each one worked, and hold the outputs in their head simultaneously.

REFERENCE LAYER FIRST

The first stage was mapping the existing MI reports, identifying what each measured, how it was keyed, and what its lineage was, before building a central reference layer that connected all the inconsistent identifiers. This reference layer became the foundation that made cross-report consistency possible for the first time.

MULTI-SOURCE MODEL, LIVE STATUS

The data model pulled from Excel documents, CSV exports, start-of-day source files, and live API connections. Every metric card carried a dynamic status: **“Pending update”** where data was at least a day stale, **“Updating”** during refresh, **“Updated at [timestamp]”** once a same-day refresh completed. Status reflected metric-level freshness, not just report-level refresh timing.

LAYERED DRILL-DOWN

A summary page presented high-level figures across all metrics, with volumes displayed as colour-coded pipelines, segmented by breaks, age, or value depending on what was operationally relevant. Every metric was clickable to a dedicated dashboard with full drill-down, and from there to the underlying data or the original MI report. Coverage spanned exception and break volumes, aged items, throughput, and pipeline status by stage.

ROW-LEVEL SECURITY

Six tiers, covering Individual Contributor, Team Leader, VP, Senior VP, Division Head and Director. Each user seeing only what was appropriate to their role and span of responsibility, with drill-down available at every tier.

Before the dashboard, performance data had been held at senior level only. Individual contributors were accountable for outcomes against numbers they could not see. The dashboard closed that gap entirely, and by hosting it inside the Asset Service Hub it became the natural starting point for every user's day: accessible in one place, visible at the right level, and refreshed continuously throughout the working day.

POWER BI

POWER QUERY

DAX

RLS · 6 TIERS

API + CSV + XLSX SOURCES

03 — SHAREPOINT · ASSET SERVICES DIVISION

SharePoint Hub & Asset Service Hub.

A single platform replacing fragmented network drives and personal notes, expanded across all five Asset Services teams under one divisional home, the natural starting point for every user's day.

<p>5</p> <p>TEAM SITES UNDER ONE HUB</p>	<p>10</p> <p>ONE-CLICK PLATFORMS IN HEADER</p>	<p>One source</p> <p>OF TRUTH, PROPAGATING AUTOMATICALLY</p>
---	---	---

The project addressed two problems. First, the division operated on a shared network drive, well-organised for compliance but reliant on the legacy Office file-sharing model, producing a chaotic daily cycle of overwritten changes, locked files and lost work. Second, each contributor maintained their own personal notes containing links to the platforms the role depended on; when a URL changed, the update had to cascade through every person's notes individually, and frequently did not.

SITE ARCHITECTURE: FOUR ZONES

The market claims site was structured across three pages. The homepage carried four distinct zones. Zone 1 was a row of ten large-format buttons spanning the full width: EASYWAY, TLM VIEW, NetX360+, NEXUS, EDI, Document Manager, SWIFT AMH via VDS, SWIFT EMX, NEXEN ERR, and DYNAMICS. Every system an analyst would open in a normal day. Zone 2 was the migrated team file structure within SharePoint's document library, mirroring the existing compliance folder organisation. Zone 3 was a Quick Links column for regular workflow tools (Counterparty Search, Markets Volumes, Ticksheet, IBAN Convertor) plus a secondary row for less-frequent platforms (My Euroclear, My Swift, My Clearstream). Zone 4 was internal escalation contacts with name, title and photo, alongside team member cards organised by office (Liverpool, Pune).

The right-hand sidebar carried world clocks for New York and Pune, a live countdown to the US Tax Deadline, the BNY switchboard number, and a team events feed. Two further pages, Recorded Sessions and External Contacts, covered training videos and counterparty escalations.

FROM ONE TEAM TO A DIVISION

The model was expanded to all five Asset Services teams. Each received its own site built to the same structure, with role-based access ensuring members could only access their own site, while managers had visibility across the teams they were responsible for. These individual sites were linked under a parent platform, the **Asset Service Hub**, which served as the divisional home page. The KPI dashboard (Project 02) was hosted directly within it.

DAILY WORKFLOW

Open the Asset Service Hub. Review performance data on the dashboard at the appropriate level. Navigate into the team site to access every tool needed to act on it. Migration to SharePoint enabled live co-authoring, autosave and version history; current-year files lived on SharePoint, prior years on the network drive for compliance.

Power Platform *Upskilling Programme.*

Building Power BI capability inside each operational team rather than centralising it, leaving every Asset Services team with someone who understood their BAU deeply enough to report on it usefully.

5	3+	600 hrs	0
TEAMS, ONE VOLUNTEER EACH	INDEPENDENT REPORTS PRODUCED	SAVED ANNUALLY, ONE WORKFLOW	STALLED PROJECTS ON DEPARTURE

The division-wide KPI dashboard was the first time Power Platform had been visible at scale across BNY Pershing, used daily by individual contributors, team leaders, VPs and senior management simultaneously. That visibility created demand. Each department wanted its own custom reporting, built closer to the detail of their BAU workflows than the divisional summary. Handling that volume centrally was not feasible; more fundamentally, my BAU knowledge was specific to Market Claims. I did not have the operational depth in the other teams' workflows to design reporting that would genuinely reflect how they worked.

BUILDING CAPABILITY IN PLACE

Rather than route requests through the continuous improvement team (a process that typically involved formal submission, a wait of several months, and the risk of being deprioritised), capability was built directly within each team. Volunteers were recruited from across the five Asset Services teams, one per team, and a structured upskilling programme was designed and delivered.

The first session was a group lesson on fundamentals: Power Query, covering what it does and how it transforms and models data; Power BI, covering how it sits on top of that model; how the two services interact. The second session was delivered one-to-one, a walkthrough of the divisional dashboard itself, showing not just what had been built but how: the data model, the relationships, the logic behind the measures.

REPLICATE, DON'T IMPROVE

The first task for each volunteer was deliberate in its constraint: replicate an existing report in their team, not improve it. Replication allowed them to practise the technical mechanics (connecting to data, shaping it in Power Query, constructing visuals) without the additional cognitive load of designing structure from scratch. It also produced something testable in parallel against the original.

OUTCOME

Every Asset Services team had a Power Platform-capable individual who understood their own BAU deeply enough to build reporting genuinely useful to it, something that could not have been achieved centrally.

One volunteer built a Power Query workflow that calculated daily dividend rounds on unit trust funds by shareholder, saving 600 hours annually. The monthly markets statistics and the team goal report were both automated. Beyond Asset Services, I automated the Start of Day report for the Reconciliations department and trained their team to build the End of Day report themselves, applying the same model of capability transfer rather than dependency. When I left, none of my Power Platform projects stalled.

Big *Atom*.

AUGUST 2019 – JUNE 2020 · CHESHIRE, UK

A tyre recycling start-up with a closed-loop circular proposition, operating its existing collection business while a new processing plant was designed. Hired to take ownership of operational data and accounts, and to close the visibility gap between the Cheshire site and London-based investors.

End-to-End *Operational Data Pipeline.*

From paper weighslips to live Power BI dashboards: a complete chain (weighbridge, GPS, schedule, CRM, accounting, debt collection) automated end to end, with no manual data entry remaining.

<p>0</p> <p>MANUAL ENTRIES BY GO-LIVE</p>	<p>5</p> <p>SYSTEMS JOINED ON TIMESTAMP + KEY</p>	<p>Intraday</p> <p>OPERATIONAL VISIBILITY</p>	<p>£0</p> <p>CAPEX ON HARDWARE REPLACEMENT</p>
--	--	--	---

Senior management were venture capitalists based in London with virtually no visibility into day-to-day operations or finances at the Cheshire site. There was no reporting infrastructure. Drivers attended scheduled collections, returned to site, used the weighbridge, which printed a paper weighslip recording the registration (typed manually by the driver) and the two weights. Slips accumulated and were entered the following day by the bookkeeper into a blank Excel workbook. From there, invoices were raised manually in QuickBooks. The process existed entirely for billing.

STEP 01: DIGITISING THE WEIGHBRIDGE

Replacing the physical weighbridge would have required significant capex. Instead I made a business case for **OpenWeigh**: software sitting between the sensor and the printer, intercepting the output and converting it to an exportable CSV. A dedicated always-on PC was approved, and scheduled CSV exports meant weighbridge data was available intraday with no manual intervention.

STEP 02: CLEANING REGISTRATIONS VIA GPS

Drivers typed their own registration at the weighbridge, so the data contained consistent typos, with the same vehicles appearing under multiple identifiers. Tracking system exports recorded GPS-verified registrations and locations at every timestamp. By matching tracking to weighbridge on timestamp and location, accurate registrations replaced the manually-typed ones.

STEP 03: JOINING SCHEDULE, CRM, ACCOUNTS

Weighbridge data, now carrying correct registrations, was joined to schedule data, which held the job, post-code and assigned vehicle. Customer numbers were pulled from the CRM using job postcode as the join key. The CRM held each customer's rate per tyre. The dataset was now complete at transaction level: **vehicle, customer, location, tonnage, value**. This connected to QuickBooks, generating weekly receivable invoices automatically; QuickBooks in turn connected to the Barclays business account for automated reconciliation and live debtor visibility.

STEP 04: AUTOMATED DEBT COLLECTION

The previous chase process was email, letter and phone. **ChaserHQ** connected to QuickBooks and the CRM and automatically sent templated chasers based on age of debt, escalating through a sequence until balance cleared. For customers operating by post (common among traditional car garages), ChaserHQ generated mail-merge letters for physical dispatch.

REPORTING LAYER

A **Collections Summary** presented daily / weekly / monthly tonnage against 12-week averages, per-vehicle performance, and a geographic bubble map of the North West & North Wales operational area, colour-coded by day of week. A **Tonnes Inbound Summary** tracked collected and delivered tonnage against monthly targets, top customers, and daily net weight against vehicle capacity.

McDonald's *Canada.*

MAY 2017 – APRIL 2019 • YUKON, CANADA

Assistant Manager for a franchisee with a flagship store and a satellite Walmart location. Recruited from Liverpool and sponsored for the role. Two parallel cash-control projects addressing individual cashier accountability and shift-level deposit discrepancies that had been accumulating untracked.

06	Cash Discrepancy Investigation Process	p. 10
07	Cash Management & Banking Procedures Overhaul	p. 11

Cash Discrepancy *Investigation Process.*

Building cashier-level accountability into a shared-till environment by reconstructing nine weeks of envelopes into a transaction log, and using absolute-difference methodology to expose offsetting errors.

<p>17 days</p> <p>RETROSPECTIVE DATA ENTRY</p>	<p>1%</p> <p>MCDONALD'S CANADA STANDARD</p>	<p>7.47 → .67</p> <p>BEST INDIVIDUAL % IMPROVEMENT</p>	<p>12</p> <p>SHIFT MANAGERS & TEAM LEADERS</p>
---	--	---	---

Tills were not assigned to individual cashiers but shared across the shift. When a cash-up was performed, the resulting discrepancy was attributed to the shift as a whole. If one cashier was \$50 short and another was \$50 over, the shift reconciled to zero and nothing was flagged. Individual patterns of error, whether from poor counting, float mismanagement, or deliberate short-changing, were invisible. McDonald's Canada required cash discrepancies be no greater than 1% of total cash tendered; without transaction-level records, the standard could not be meaningfully applied.

THE WORKBOOK: SIX SHEETS

A cash tracking workbook was built from scratch and data entered retroactively from physical envelopes. The **Entry** sheet provided a clean daily input form, associating each cashier's count with the supervising manager's initials. The **Data** sheet computed cash over/short, and derived an **absolute difference**, stripping the sign to give a direction-neutral measure of inaccuracy. This was deliberate: a net of zero could mask a \$5 short on one till and \$5 over on another, producing a false impression of accuracy.

The **Scorecards** sheet produced automated rankings updated in real time. For each cashier: total absolute discrepancy, number of tills counted, average per till, total tendered, and discrepancy as % of total tender, the metric used to apply the standard. The **Cash Policy** column tracked the date each cashier had signed the new policy, flowing in from a tracking sheet, a prerequisite for any future disciplinary action. A **Trend** sheet aggregated weekly absolute discrepancies, surfacing spikes for investigation.

HIGHLIGHT

The worst performing cashier at 7.47% discrepancy dropped to 0.67%.

The first report was distributed transparently after 17 days of collection, including my own figures, ranked among the managers. Only one manager met the 1% standard at 0.45%; one was at 14.58%, traceable to a single till with a \$164.43 float discrepancy undetected at shift start. No formal actions followed from the first report. The intent was to establish the data record, communicate that performance was now visible and attributed, and let the data drive behaviour before moving to formal process.

Across nine weeks, the effect was measurable. Of cashiers present in sufficient volume in both early (W5-9) and later (W10-13) periods, the majority improved. Several high-variance cashiers reduced their rates substantially: one from 7.47% to 0.67%, another from 4.01% to 0.22%, a third from 3.90% to under 1%.

Cash Management & *Banking Overhaul.*

A structured cash-control audit, two store-floor process documents, a two-store deposit tracker with RAG thresholds, and per-day investigation workbooks generating plain-English narratives for shift managers.

<h2>36</h2> <p>CONTROLS AUDITED, 5 CATEGORIES</p>	<h2>67%</h2> <p>INITIAL AUDIT SCORE</p>	<h2>-\$11,768</h2> <p>CONSOLIDATED 2018 SHORTFALL</p>	<h2>RAG</h2> <p>WEEKLY THRESHOLDS, BOTH STORES</p>
---	---	---	--

Whilst Project 06 addressed cashier accountability, a parallel problem existed at deposit level. Discrepancies there reflected errors introduced during counting and deposit preparation by shift managers, separate from cashier-level shortfalls. These had been accumulating untracked. January 2018 alone totalled **-\$782.04** against expected, annualising to a projected **\$9,207.89** loss.

CASH CONTROL AUDIT: 10 FEB 2018

Thirty-six controls assessed across five categories with full, partial or zero credit. Overall **67%**. Weakest areas: CCTV coverage **57%** (main safe, drop safe, and several front counter POS not adequately covered) and cash audit envelope completion **50%**. Cash controls **63%**, with POS codes not changed weekly, skims not made four times daily, panic alarms not worn or tested. Safe controls **83%**, with combination not changed when managers left, USD held in safe rather than deposited.

TWO PROCESS DOCUMENTS

Displayed in-store. The cashier document, “Your Till, Your Responsibility”, communicated four behavioural expectations: **count your float, use only your till, lock your till, check your change**. The manager document set out four mandatory deposit steps: **till count** (count every till, consolidate), **float count** (verify all unused floats at \$150, balance against backup), **backup count** (verify backup at \$2,250, sign for the count), **deposit count** (complete envelope, check against drawer, complete cheque book, countersign and seal).

DEPOSIT TRACKER & INVESTIGATION WORKBOOK

Daily deposits across Main Store and Walmart with auto-calculated variance. A weekly Email sheet pulled each day's variance, computed projected versus actual shortfalls, and applied **RAG thresholds**: green below 40% of projected, amber 40-70%, red above 70%. Where flagged, an investigation workbook performed a till-by-till reconciliation split by Open / Mid / Close / Overnight shift, with a named manager for each. A Conclusion Data sheet auto-generated plain-English narrative sentences, formatted into a structured investigation report sent directly to the shift managers involved.

2018 END-OF-YEAR RESULT

Consolidated deposit discrepancy **-\$11,768.16** for the year (Main Store **-\$9,093.79**, Walmart **-\$2,674.37**), equivalent to a full-time crew member at \$14/hr for 5.5 months, or a part-time crew member for 11 months. Concentrated effort between March and July brought Main Store deposits consistently below the projected loss baseline, the first sustained period of improvement since tracking began.

Remittance 360.

NOVEMBER 2025 – JANUARY 2026 • BLK GROUP

An electronic money institution operating between clients and payment service providers. Two solo builds in under three months: a transaction-level reconciliation tool replacing day-total comparison, and a three-layer fee management system separating client intent from execution from reporting.

08	Transaction-Level Reconciliation Tool	p. 13
09	Monthly Fee Management Tool	p. 14

Transaction-Level *Reconciliation Tool*.

Replacing day-total reconciliation with transaction-level matching across statement, ledger and PSP records, with a stateless classification layer feeding an immutable, append-only state machine.

<h2>3</h2> <p>ARCHITECTURAL LAYERS, NO OVERLAP</p>	<h2>5</h2> <p>STATE TABLES OF RECORD</p>	<h2>09:15</h2> <p>DAILY COMPLETION TIME</p>	<h2>< 3 mo</h2> <p>SOLO BUILD, END TO END</p>
--	--	---	--

The existing reconciliation approach compared aggregate daily totals (statement versus ledger at day level), then worked backwards through transactions to identify the discrepancy. This is the inverse of industry practice. Rather than surfacing exceptions directly, the process began with a known difference and required manual investigation to locate it within potentially hundreds of transactions. Reconciliation typically completed around 1pm.

LAYER 01: PRESENTATION

Six tabs reflecting transaction *state* rather than data source: **Reconciliation**, **Exceptions**, **Pending**, **Matched**, **Off Platform**, **Audit**. A dashboard reconciles cash balances using ECB FX rates, with EUR as control currency, off-platform balances isolated and netted separately, and an operational summary driven entirely by table counts. The dashboard reflects state; it does not drive logic.

LAYER 02: CLASSIFICATION (POWER QUERY)

On every refresh, classifies each item deterministically into AutoMatch, AutoMatchInternal, or AutoException. Stateless and idempotent, with no knowledge of yesterday, manual resolutions, or pending items. Rebuilds from scratch on every run. The matched and exception views it produces are disposable; they exist only to feed the layer below.

LAYER 03: STATE MACHINE (VBA, SYSTEM OF RECORD)

Five tables hold the truth: **Matched_History**, **Exceptions_Outstanding**, **Pending_Outstanding**, **OffPlatform_History**, **Audit_Log**. Power Query refreshes first; VBA reads its output and uses Audit_Log as a global deduplication ledger. **Any ItemID that has ever appeared in the Audit_Log will never be auto-inserted again**, regardless of refresh count. Nothing duplicates across runs.

Exceptions follow an explicit lifecycle with user-initiated transitions: AutoException → Pending / OffPlatform / ManualMatch, with Recall paths back. Every transition physically moves the row, writes exactly one audit entry, and never mutates history silently. Manual matching enforces strict rules: selection required, currency must match, net must equal zero using ledger-reversal logic. Recall from Matched_History operates at group level: all rows sharing a ManualMatchID recalled together. The original audit record is preserved.

AUDIT VIEWER

Forensic access to any ItemID, surfacing current location across all state tables, full chronological history of every transition, and the complete match group for any item that was manually matched. The audit log is **append-only and immutable**.

Monthly Fee *Management Tool*.

Replacing a matrix spreadsheet that destroyed audit trails with a three-layer model: a deduplicated client master, a per-period transactional ledger, and a read-only dashboard derived from operational truth.

<h2>3</h2> <p>LAYERS, CLEAN SEPARATION</p>	<h2>1 row</h2> <p>PER CLIENT PER DUE DATE</p>	<h2>Auto</h2> <p>NEW-CLIENT ONBOARDING</p>
--	---	--

As an EMI, Remittance 360's client accounts did not carry overdraft facilities, so standing orders could not be used to collect monthly fees automatically. Fees were taken manually on or after the 15th. The legacy process used a matrix-style spreadsheet, with clients on Y, months on X, fee in each cell. When collected, the cell was zeroed.

Three failures: **(i)** zeroing destroyed the audit record entirely, with no record of when, by whom, or why a fee was delayed; **(ii)** a single shared comments column at the end of the months covered every period for each client, forcing comments to be either absent or compressed; **(iii)** new clients were added manually from memory, with omissions inevitable, drift over time.

LAYER 01: CLEANCLIENTDATA

The operational master list. Power Query pulls from all business entities within BLK Group, cleans and standardises, and deduplicates clients who have moved between entities, so each client appears exactly once regardless of history. Output: one row per active client with fee rules, fee status, relationship manager, and entity. Sole input to fee generation. New clients appear automatically on next refresh.

LAYER 02: MONTHLYFEEDATA

The transactional ledger. **One row per client per due date.** Populated exclusively by the AppendMonthlyFees macro; users cannot write directly. Each new month's entries generate automatically with due date calculated to the 15th business day. Fee Amount Paid and Date Settled are user-entered; Remaining is derived. Each row carries its own Comments field: a per-period record rather than a single shared note. Clients with no monthly fee generate a row marked accordingly, ensuring the ledger is complete.

LAYER 03: OUTSTANDING FEES DASHBOARD

A read-only reporting view built on a pivot of MonthlyFeeData. Sliceable by VIP entity, client status, relationship manager, fee status, and balance state. **Grand totals reconcile to remaining balance, not original fee amounts**, meaning the dashboard reflects what is actually owed, not what was originally billed. Format was designed to match the export managers were already familiar with, minimising the transition.

ARCHITECTURAL PRINCIPLE

CleanClientData defines who should be billed and at what rate. MonthlyFeeData records what actually happened, month by month, with a full audit trail per row. The dashboard surfaces operational truth derived from that record.

Amey.

DECEMBER 2014 – APRIL 2017 · BUSINESS EFFICIENCY

Civil infrastructure services across utilities, defence and government contracts. Placed within the Business Efficiency function on contracts that were underperforming financially, finding savings without affecting productivity, through workflow redesign, waste-cost elimination, and reporting that surfaced data which actually drove decisions.

10	Fuel Efficiency & Theft Detection	p. 16
11	Contract P&L Reporting Suite	p. 17

10 — TELEMATICS × FUEL CARD RECONCILIATION

Fuel Efficiency & *Theft Detection.*

A tank-capacity reconciliation methodology turning telematics consumption against fuel-card purchases into an arithmetic test for fraud, formalised into a fleet-level reporting tool covering 47 contracts.

-10% FUEL BILL, NO MILEAGE DROP	1 → 47 CONTRACT GROWTH	9 ROLLING WEEKS PER ANALYSIS	Targeted CCTV REQUESTED PER OUTLIER
--	----------------------------------	---	--

Fuel was among the highest cost lines on any contract involving field gangs. On the United Utilities contract, a deliberate programme reduced fuel cost across three fronts: renegotiating fuel card deals, promoting efficient driving to reduce idling and harsh handling, and deterring fuel card fraud. The fraud pattern was specific: employees entrusted with fuel cards were filling jerry cans or personal vehicles at the pump in addition to their work vehicles. Without requesting CCTV for every transaction across the entire fleet (neither practical nor proportionate), the only viable detection route was through the data.

THE TELEMATICS INSIGHT

Unlike MPG-based estimates derived from mileage, the tracking system used a sensor in the fuel injector valve to report **actual fuel consumption**. The telematics reading was a direct measure of what went into the engine, not a calculation. Combined with fuel card transaction data showing what was purchased at the pump, a mathematical discrepancy could be calculated at vehicle level.

METHODOLOGY: TANK CAPACITY AS THRESHOLD

Fleet records held the precise tank size for each vehicle type: a Sprinter van, for example, carried a 90-litre tank. **Starting from an assumed empty tank**, each fuel card purchase was added and each telematics consumption reading subtracted, producing a running tank balance across a rolling nine-week window. If the balance exceeded the vehicle's physical tank capacity, it was arithmetically impossible for all the purchased fuel to have entered that vehicle. The starting assumption of zero was conservative (the tank would never truly be empty), making any breach unambiguous regardless of starting conditions.

Random vehicles were selected for analysis. Outliers had CCTV requested for those specific transactions. Disciplinary proceedings followed where footage confirmed the fraud. The investigation workbook structured one tab per source: telematics journey data, fuel card statement, week lookup, vehicle tank and MPG reference; with a running per-row tank balance and a weekly summary auto-generating key figures and narrative findings for HR proceedings.

RESULT

The fuel bill fell by 10% without a corresponding drop in mileage, a movement that could not be attributed to improved driving alone, and was consistent with deterrence of card misuse.

The methodology was formalised into a fleet-level reporting tool, initially covering Yorkshire Water's 118-vehicle fleet, then expanded across an additional 46 contracts in the CGU division, including United Utilities, Severn Trent, Welsh Water, Northern Powergrid, Western Power Distribution, Scottish Water, Northumbrian Water, Affinity Water, and several central government accounts. Additional layers: V Power overspend tracking, MPG cost impact conversion to weekly £ per vehicle, and driver behaviour score correlation, distinguishing fuel cost increases driven by poor driving from those with no legitimate explanation.

11 — YORKSHIRE WATER WWSA · PRE-POWER-PLATFORM

Contract P&L *Reporting Suite*.

A structural redesign that left the raw export untouched, performed transformations in segregated columns, and compressed weekly cost feedback from over a week to under two days.

3 days → 0 <small>MANUAL RECONCILIATION PER WEEK</small>	Mon <small>DISTRIBUTION DAY, 10AM</small>	Fri → Tue <small>WEEKLY P&L CALL</small>	11 <small>SOURCE EXPORT TABS</small>
--	---	--	--

Each Amey contract had a finance department producing the official monthly P&L, a legally accurate account of the prior month's costs, delivered at the start of the following month. By the time it landed, costs were weeks old. The Business Efficiency team's scope included producing a parallel weekly P&L from available cost exports, giving contract managers current-week visibility they could act on rather than waiting for the official pack to confirm what had already happened.

The Yorkshire Water WWSA contract produced a weekly P&L breaking costs across **South West, South East, North** geographic areas, each subdivided by **Civil, Jetting, Non-Core** workstreams, with separate overhead and disallowed columns and comparison against the current RF1 forecast. Output covered every meaningful cost line: DLO and DSP revenue, direct and overtime labour headcount and cost, materials by category, plant hire, vehicle hire, fuel, subcontract, mobile phones, traffic management, tipping, fines, and provisions.

THE INHERITED PROCESS

Owned by a senior manager who would begin each week on Monday, spend three full days collating and reconciling exports, and distribute by end of day Wednesday. Managers reviewed Thursday; the weekly call took place Friday. The process followed a manual specifying what to do with each export once pasted: which columns to add, which to reformat, which values to recode. Each intervention was applied directly to the pasted source data, modifying the raw export, so formulas could not be left in place and reused; the manual had to be followed from the beginning each week.

THE STRUCTURAL FIX: PRE-POWER-PLATFORM


On taking ownership, the fix was structural rather than technological, since Power Platform tools were not yet available. On every export tab (Vehicle Hire, Plant Hire, GR, PO, Store Movements, Weekly Timesheet Data, Mobile Phones, Subcontract, Revenue, Overheads, Fuel), a set of **yellow-highlighted columns** was added to the right of the raw data. These columns contained formulas performing all required transformations and lookups *without touching the source data itself*. The raw export was pasted in and left intact. All processing happened in segregated yellow columns.

A central **Cost Feed Sheet** pulled from those columns using SUMIFS across all source tabs, aggregating each cost category into the correct P&L account code and operational area. The output produced the formatted weekly P&L automatically, including week-on-week tracking and comparison against the RF1 forecast.

THE COMPRESSION

Weekly P&L distributed by **10am Monday** for the week just closed. Managers had cost visibility on Monday morning. The weekly call moved from Friday to Tuesday midday, **compressing the feedback loop from over a week to less than two days**. Cost problems surfaced close to real time, so action could be taken within the same week rather than discovering the issue when the official monthly P&L arrived the following month.

Is the next one
yours?



robert.bridge@icloud.com
[linkedin.com/in/robertbridgeuk](https://www.linkedin.com/in/robertbridgeuk)
London, UK