A Unified, Standards-Certified Platform for HVAC Equipment Selection and Lifecycle Costing:

Architecture, Computational Methods, and Field Evidence from the MileSoft Suite

Engineering Research Group
MileSoft Software Technologies
https://milesoft.net/

Corresponding authors: *Deepak Daryani* (deepakdaryani@milesoft.net), *Tanuj Daryani* (tanujdaryani@milesoft.net)

Working paper — Version 1.0 — September 8, 2026

Abstract

Heating, ventilation, and air-conditioning (HVAC) equipment selection sits at the origin of the building-energy *performance gap*: the persistent divergence between the performance designers specify and the performance buildings actually deliver. Incumbent practice relies on spreadsheets and disconnected vendor-specific selectors that fragment certified performance data, defeat auditability, and decouple first-cost selection from lifecycle value. This paper presents the architecture, computational methods, and field evidence of the **MileSoft HVAC software suite**—a unified, standards-certified platform organised around a single governed *certified-curve engine* that serves six interoperating modules: air handling units (AHU), fan coil units (FCU), energy recovery ventilators (ERV), chillers, direct-expansion (DX) systems, and lifecycle-cost analysis (LCA). We describe (i) a shared performance-data model and an interpolation–psychrometric core that binds every selection to AHRI-, EUROVENT-, and Indian-Standard–certified ratings; (ii) module-level methods spanning section-by-section AHU configuration and bill-of-materials costing, chilled-water coil rating, air-to-air effectiveness modelling, AHRI 550/590 integrated and non-standard part-load analysis with Indian seasonal efficiency, multi-circuit DX analysis, and discounted lifecycle costing with electrical-infrastructure and carbon accounting; (iii) a governance and lifecycle layer providing versioned libraries, role-based workflows, and audit trails that satisfy the EUROVENT requirement for integrated, version-stamped selection software; and (iv) a management-information and business-intelligence layer. We anchor the platform in the governing standards landscape and report a production deployment at a EUROVENT-certified AHU manufacturer that reduced selection-cycle time by approximately **70%**, eliminated audit findings, **tripled** per-engineer proposal throughput, and improved proposal-to-order conversion by **12%**. Finally, we propose a *capability reference framework* against which selection and lifecycle-costing platforms can be evaluated, positioning the work as both a reproducible engineering blueprint and a benchmark for the field.

Keywords: HVAC equipment selection; air handling units; chillers; energy recovery ventilation; fan coil units; direct expansion; lifecycle cost analysis; psychrometrics; ANSI/AHRI 550/590; AHRI 1060; EUROVENT (EN 1886/EN 13053); IS 16590 / ISEER; ASHRAE 90.1; building-energy performance gap; selection software; engineering governance.

1 Introduction

The built environment is responsible for roughly one-third of global final energy use, and a dominant share of that demand is mediated by heating, ventilation, and air-conditioning (HVAC) plant. It is therefore unsurprising that an extensive body of research, regulation, and capital has been directed at improving HVAC efficiency. What is more surprising—and more consequential—is how reliably buildings fail to deliver the performance promised on paper. This divergence between predicted and measured performance is now well documented as the building-energy *performance gap* [1–3]. Reported magnitudes are sobering: field studies of non-domestic buildings have repeatedly found measured energy use exceeding design predictions—commonly by tens of percent and, in individual cases, by 60 % or more—with the largest discrepancies tending to occur in buildings served by *complex* HVAC systems [1–3].

A critical and often under-appreciated contributor to this gap is *specification uncertainty*—error introduced not during operation but at the moment a designer specifies equipment [3]. When the assumed properties of a coil, a fan, or a recovery device do not correspond to the certified performance of the equipment actually installed, the gap is baked in before the building is even built. Empirical reviews single out exactly the quantities that HVAC selection tools are supposed to fix: assumed specific fan powers that are systematically lower than the values realised in operation, and equipment loads that are underestimated at design time [3]. In other words, a meaningful fraction of the performance gap is an *engineering-tooling* problem, and it is addressable at the point of selection.

1.1 The selection bottleneck

For an HVAC manufacturer or specifying contractor, equipment selection is simultaneously the most technical and the most commercially exposed step in the sales motion. A single air handling unit couples at least five inter-dependent calculations—coil sensible and latent capacity, fan-curve fit at the operating static pressure, motor sizing with future static reserve, filter pressure drop across the loading life, and acoustic levels—each of which depends on the others. Chillers spend the overwhelming majority of their service life at part load, so selecting them on full-load efficiency alone is, quite literally, selecting on a number the machine almost never sees. Energy recovery ventilators, once optional, are now mandated across wide swaths of the commercial building stock and must be defended on effectiveness *across* the annual operating envelope, not at a single rating point. And buyers increasingly demand a multi-year total-cost-of-ownership (TCO) comparison rather than a first-cost quotation.

In practice, this complexity is most often handled by spreadsheets and a patchwork of vendor catalogues. That arrangement fails in characteristic ways. Performance numbers drift as catalogue revisions go unpropagated across dozens of engineers' laptops; two engineers in the same office quote the same project differently; a submittal escapes with a coil that cannot deliver the stated latent capacity; and, when an auditor asks where a given performance figure originated, no clean provenance exists. The cost is not abstract: it manifests as margin leakage, post-order disputes, warranty claims that surface years after handover, and—through specification uncertainty—as a direct contribution to the building performance gap.

1.2 The certification dimension

The standards landscape has, in fact, already concluded that ad-hoc tooling is inadequate. The EUROVENT Certified Performance (ECP) programme for air handling units—the de-facto European benchmark, built on EN 1886 and EN 13053—requires that "the selection software shall be operative as an entity with all unit components integrated in one software," and that components selected by different software or any other means "cannot be certified"; moreover, "each quotation of a certified AHU shall include the date/code/number of the software version used for the selection" [10–12]. The implication is strong: an *integrated, version-stamped, governed* selection engine is not merely a convenience but a *precondition* for issuing certified selections. Parallel requirements hold elsewhere—ANSI/AHRI 550/590 and 551/591 prescribe exactly how chiller full-load and part-load efficiency must be computed and reported [4, 5]; AHRI 1060/1061 (referencing ANSI/ASHRAE 84) governs air-to-air energy-recovery effectiveness [7, 8]; India's IS 16590 defines the seasonal metric (ISEER) for the climate [13, 14]; ASHRAE 90.1 and the EU Ecodesign framework determine *when* energy recovery is required [15, 16]; and the EU F-Gas Regulation now constrains refrigerant selection on the basis of global-warming potential [17].

1.3 Contributions

This paper presents the design and field evidence of the **MileSoft HVAC software suite**, a production platform that treats the above not as six separate selectors but as one governed engineering system. The platform is organised around a single *certified-curve engine*—a shared performance-data store, interpolation and psychrometric core, validation-rule layer, project-lifecycle and versioning service, bill-of-materials (BOM) costing engine, and governance layer—which serves six interoperating modules (AHU, FCU, ERV, chiller, DX, and lifecycle-cost analysis) and an analytics layer. Our contributions are:

1. **A unified architecture** for HVAC selection in which every module draws performance from one certified, version-governed store, so that a library update propagates everywhere at once and every output is traceable to a specific certified rating and software version (Section 3).
2. **Module-level computational methods**, documented against the governing standards, for: section-by-section AHU configuration with granular BOM costing; chilled-water FCU coil rating with sensible/latent decomposition and specific-fan-power accounting; air-to-air ERV effectiveness modelling with dual-stream pressure-drop and fan selection; AHRI 550/590 full-load, IPLV, and NPLV analysis together with IS 16590 ISEER; multi-circuit DX analysis with line-set and ambient corrections; and discounted lifecycle costing with electrical-infrastructure and carbon accounting (Section 4).
3. **Field evidence** from a production deployment at a EUROVENT-certified AHU manufacturer, including reported reductions in selection-cycle time, audit findings, proposal throughput, and conversion (Section 5).
4. **A capability reference framework**—a structured benchmark—against which selection and lifecycle-costing platforms can be evaluated, derived from the standards landscape and the failure modes of incumbent tooling (Section 6).

The remainder of the paper is structured as follows. Section 2 reviews the performance gap, the governing standards, and the limitations of incumbent tooling. Section 3 presents the platform architecture and the certified-curve engine. Section 4 documents the computational methods of each module, illustrated with screens and data from the production system. Section 5 reports the deployment case study. Section 6 discusses how the architecture closes the specification-side of the performance gap and proposes the capability reference framework. Sections 7 and 8 address limitations and future work, and Section 9 concludes. Appendix A lists nomenclature; Appendix B provides worked numerical examples drawn from the production system.

A note on scope and method. This is an applied, systems-oriented research paper: a documented account of a deployed engineering platform, grounded in the governing standards and in a real production dataset, together with a single operator-reported deployment case study. It is not a controlled multi-site trial, and the field metrics reported in Section 5 are operator-reported rather than independently audited (see Section 7). The contribution is the architecture, the standards-aligned methods, and the capability framework—artifacts intended to be reproducible and independently testable.

2 Background and Related Work

A platform that claims to bind selection to certified performance must implement the methods that the certification standards prescribe. This section assembles that technical backbone—the metrics and equations that recur throughout Section 4—and then situates the platform against incumbent tooling.

2.1 Certified performance: metrics and methods

2.1.1 Chiller full-load and part-load efficiency (AHRI 550/590; IS 16590)

ANSI/AHRI Standard 550/590 (I-P), with its SI counterpart 551/591, standardises how the performance of a vapour-compression water-chilling package is rated and reported, so that machines from different manufacturers can be compared on a common basis [4, 5]. Beyond full-load efficiency at standard rating conditions, the standard defines the *Integrated Part-Load Value* (IPLV) and *Non-standard Part-Load Value* (NPLV) as weighted single-number measures of part-load efficiency, evaluated at the 100%, 75%, 50%, and 25% load points. Writing A, B, C, D for the efficiency (in kW/ton, or its reciprocal in SI) at those four loads, the weighted-reciprocal form is

$$\mathrm{IPLV} = \frac{1}{\dfrac{0.01}{A} + \dfrac{0.42}{B} + \dfrac{0.45}{C} + \dfrac{0.12}{D}}, \tag{1}$$

with the weights reflecting the fraction of operating hours an "average" single chiller spends near each load [4, 6]. IPLV uses the standard condenser-relief trajectory fixed by the standard; NPLV applies the *same* equation but at the project's own condenser conditions, and is therefore the more faithful predictor of site energy for a specific installation [6]. Crucially, IPLV/NPLV describe a single stand-alone chiller; the standard itself cautions that a comprehensive hourly analysis remains the most accurate predictor [6]—a caveat the platform answers with bin-energy modelling (Section 4.4).

For the Indian market, the Bureau of Energy Efficiency mandates the *Indian Seasonal Energy Efficiency Ratio* (ISEER) under IS 16590, derived using the ISO 16358-1 methodology over Indian weather bins [13, 14]. ISEER is the ratio of the cooling seasonal total load (CSTL) to the cooling seasonal energy consumption (CSEC),

$$\mathrm{ISEER} = \frac{\mathrm{CSTL}}{\mathrm{CSEC}} = \frac{\sum_j n_j \dot{Q}_j}{\sum_j n_j P_j}, \tag{2}$$

where n_j is the number of hours in temperature bin j and $\dot{Q}_j, P_j$ are the cooling capacity and power in that bin. BEE chiller certification explicitly requires the coefficient of performance at 25/50/75/100% load together with the ISEER value under IS 16590 conditions [13].

2.1.2 Air-to-air energy recovery (AHRI 1060; ASHRAE 84)

ANSI/AHRI Standard 1060 (I-P), with SI counterpart 1061, rates air-to-air exchangers for energy-recovery ventilation, referencing the method of testing in ANSI/ASHRAE Standard 84 [7, 8]. Effectiveness is the ratio of actual transfer to the maximum thermodynamically available transfer between the supply and exhaust streams. For sensible, latent, and total quantities respectively,

$$\varepsilon_s = \frac{\dot{m}_{\min} c_p (t_1 - t_2)}{\dot{m}_{\min} c_p (t_1 - t_3)}, \quad \varepsilon_l = \frac{\dot{m}_{\min} (w_1 - w_2)}{\dot{m}_{\min} (w_1 - w_3)}, \quad \varepsilon_t = \frac{\dot{m}_{\min} (h_1 - h_2)}{\dot{m}_{\min} (h_1 - h_3)}, \tag{3}$$

where stations 1 and 3 denote entering supply (outdoor) and entering exhaust air, station 2 denotes leaving supply air, t, w, h are dry-bulb temperature, humidity ratio, and enthalpy, and $\dot{m}_{\min}$ is the smaller of the two mass-flow rates [7, 8]. Standard 1060 further defines *net* effectivenesses that correct for the exhaust-air transfer ratio (EATR, cross-leakage) and outdoor-air correction factor, and requires that effectiveness be reported alongside supply- and exhaust-side pressure drop—because the energy a recovery device saves can be partly offset by the fan energy required to overcome its pressure drop [7, 9].

2.1.3 Air handling unit casing and component rating (EN 1886; EN 13053)

The EUROVENT ECP programme rates AHUs against EN 1886 (mechanical performance) and EN 13053 (rating and performance of units, components, and sections) [10–12]. EN 1886 classifies the casing on mechanical strength (D1–D4), air leakage (L1–L4), thermal transmittance (the U-value, classes T1–T5), thermal-bridging factor (TB1–TB5), and filter-bypass leakage; EN 13053 governs the rating of components and the final pressure drops used for fans and filters over the service interval. As noted in Section 1, ECP additionally requires that all components be selected within *one* integrated software entity and that each certified quotation carry the software version code [12]—a governance requirement, not merely a numerical one.

2.1.4 Coil duty and fan power (psychrometrics)

Coil selection in the AHU, FCU, and DX modules rests on moist-air thermodynamics [18]. For an air-side mass-flow $\dot{m}_a$, the total and sensible cooling duties and the water-side duty are

$$\dot{Q}_t = \dot{m}_a\,(h_{\text{in}} - h_{\text{out}}), \qquad \dot{Q}_s = \dot{m}_a\,c_{p,a}\,(T_{\text{in}} - T_{\text{out}}), \qquad \dot{Q} = \dot{m}_w\,c_{p,w}\,\Delta T_w, \tag{4}$$

with the latent duty $\dot{Q}_l = \dot{Q}_t - \dot{Q}_s$ and the sensible heat ratio $\text{SHR} = \dot{Q}_s/\dot{Q}_t$. Fan energy is captured by the specific fan power, $\text{SFP} = P_{\text{fan}}/\dot{V}$, the very quantity that empirical studies identify as routinely underestimated at design time [3]. Section pressure drop accumulates additively along the air path, $\Delta p_{\text{tot}} = \sum_k \Delta p_k$, with filters contributing a value that rises from an initial (IPD) to a final (FPD) pressure drop over their loading life [11].

2.1.5 Lifecycle cost

Lifecycle-cost analysis follows the discounted-cash-flow convention codified in ISO 15686-5 and NIST Handbook 135 [19, 20]. For an analysis horizon of N years and discount rate r,

$$\text{LCC} = C_0 + \sum_{t=1}^{N} \frac{C_{e,t} + C_{m,t} + C_{r,t}}{(1+r)^t}, \tag{5}$$

where C_0 is first cost (equipment, installation, commissioning) and $C_{e,t}, C_{m,t}, C_{r,t}$ are the year-t energy, maintenance, and replacement costs. The net present value of choosing a proposed option over a baseline, and the discounted payback (breakeven year), follow directly from the annual cash-flow differences.

2.2 Codes that drive demand, and constrain refrigerant

Two regulatory currents shape *which* equipment a selection must consider. First, energy-recovery mandates: ANSI/ASHRAE/IES 90.1 (and, by adoption, the IECC) requires exhaust-air energy recovery once a fan system exceeds a supply-airflow threshold (on the order of 5000 CFM) at high outdoor-air fractions, historically at $\geq 50\,\%$ enthalpy recovery, with the 2025 cycle sharply raising the bar toward $80\,\%$ in many climate zones [15]; the EU Ecodesign framework imposes parallel requirements on non-residential ventilation units [16]. Second, refrigerant constraints: the EU F-Gas Regulation (EU) 2024/573 tightens the hydrofluorocarbon phase-down on a global-warming-potential (GWP) basis, with service and placing-on-market restrictions reaching air-conditioning, heat-pump, and chiller equipment through the latter half of the decade [17]. A modern selection platform must therefore expose effectiveness across the annual envelope (for recovery compliance) and refrigerant GWP/charge inventory (for F-Gas exposure)—both of which the MileSoft modules report.

2.3 Incumbent tooling and its failure modes

Against this backdrop, the dominant tools remain (i) general-purpose spreadsheets, (ii) single-vendor catalogue selectors, and (iii) isolated point tools for one equipment class. Each has structural weaknesses. Spreadsheets place certified curves on individual laptops with no version control, no audit trail, and no enforced feasibility rules, so performance silently drifts and provenance is lost—precisely the specification uncertainty implicated in the performance gap [3]. Single-vendor selectors are accurate within a catalogue but cannot be composed across an AHU–chiller–FCU project, do not carry the manufacturer's own branding or BOM logic, and lock data inside one supplier. Point tools solve one equipment class but re-introduce the integration problem the moment a project spans several. None of the three, on its own, satisfies the EUROVENT requirement for a single integrated, version-stamped selection entity [12]. The platform described next is a direct response to these gaps.

3 System Overview and Architecture

3.1 Design goals

The platform is governed by six design goals, each a direct response to a failure mode identified in Section 2:

G1 — Single certified spine.
Every performance number originates from one certified store, so that any output traces to a specific certified rating rather than to an undocumented spreadsheet cell.

G2 — One-push library governance.
When research and development revises a coil, fan, or compressor family, the new performance propagates to *all* active projects at once; there are no stale copies to chase.

G3 — Enforced feasibility.
Configuration rules prevent invalid combinations from ever being quoted (for example, a recovery wheel whose bypass damper would short, or a fan that cannot sit on its operating curve).

G4 — Project lifecycle and versioning.
Selections live inside a versioned project with explicit revisions and submittal-versus-as-built deltas, not in an engineer's downloads folder.

G5 — Integrated costing and output.
Branded submittals, BOM, performance sheets, and price quotations are emitted from the same engine that performed the selection, eliminating manual re-keying.

G6 — Governance, integration, and audit.
Role-based workflows, single sign-on, clean export to enterprise resource planning (ERP) systems, and an end-to-end audit trail make selections governable and certifiable.

Goals G1, G2, G4, and G6 are precisely what the EUROVENT ECP programme demands of certified selection software [12]; the platform treats them as first-class architecture rather than as features bolted onto a calculator.

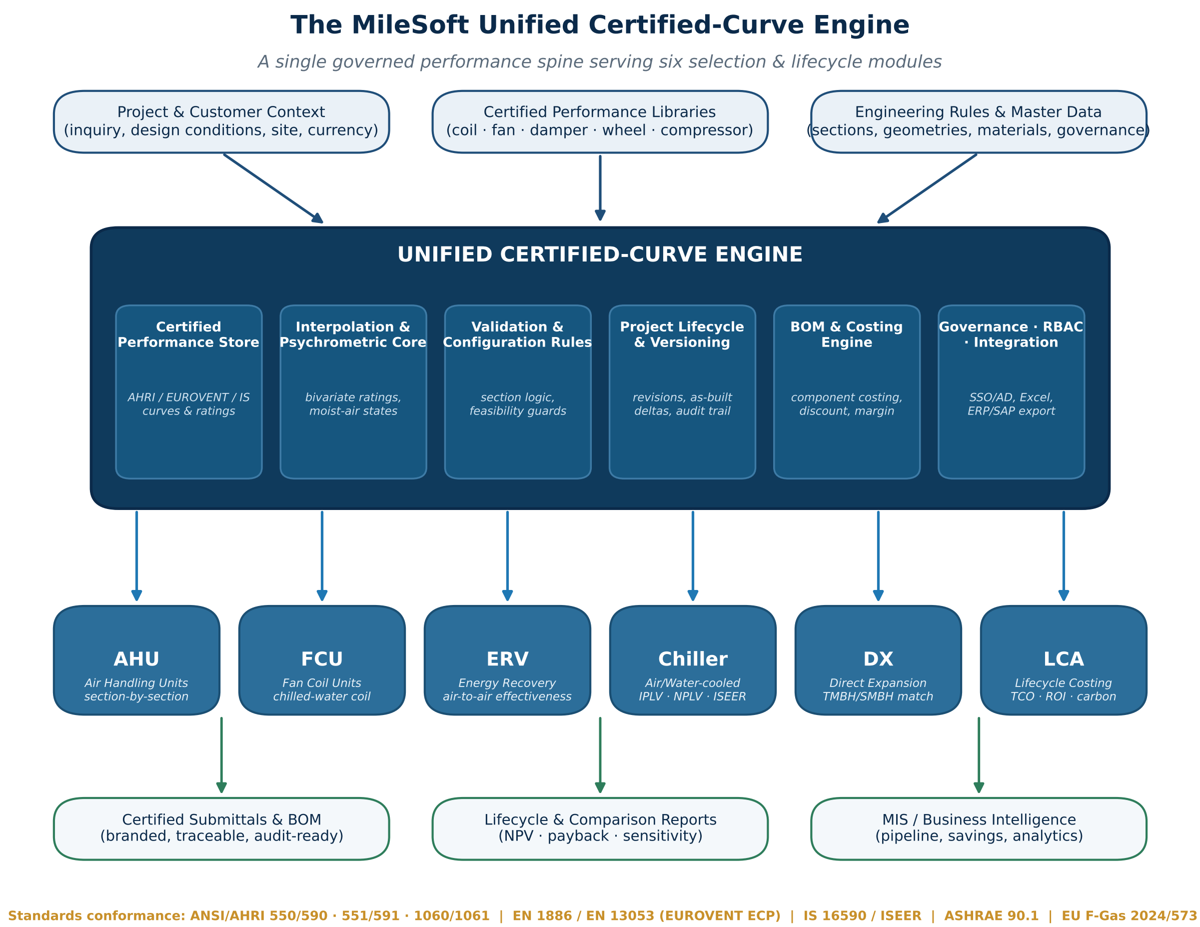

Figure 1. Platform architecture. A single *certified-curve engine*—comprising a certified performance store, an interpolation and psychrometric core, a validation and configuration-rule layer, a project-lifecycle and versioning service, a BOM and costing engine, and a governance/integration layer—ingests project context, certified libraries, and engineering master data, and serves six interoperating modules (AHU, FCU, ERV, chiller, DX, lifecycle-cost analysis). All modules emit traceable submittals and BOMs, lifecycle and comparison reports, and feed a management-information/business-intelligence layer. Standards conformance spans the chiller, recovery, AHU, seasonal-efficiency, energy-recovery-mandate, and refrigerant regimes.

3.2 The certified-curve engine

The engine (centre of Fig. 1) is the heart of the platform. Its six services are:

Certified performance store. A governed repository of certified ratings—coil and fan curves, damper and recovery-wheel data, and compressor/chiller performance maps—each tagged with its certification basis (AHRI, EUROVENT, or IS), its rating conditions, and a version code. This store is the single source of truth referenced by every module (G1).

Interpolation and psychrometric core. A numerical layer that evaluates performance *between* tabulated certified points. For chillers this means bivariate interpolation of capacity and power over leaving-chilled-water and condensing temperature (Section 4.4); for coils and recovery devices it means evaluating moist-air states—enthalpy, humidity ratio, wet-bulb, dew point—so that duty and effectiveness (Eqs. 3–4) are computed at the project's *actual* entering conditions rather than at a catalogue reference point.

Validation and configuration rules. The layer that enforces G3: section-ordering logic for AHUs, feasibility guards on coil/fan/recovery combinations, and range checks on operating conditions, so the engine cannot "successfully" produce an infeasible selection.

Project-lifecycle and versioning service. Implements G4: projects contain versioned units; revisions are tracked; and the deltas between a quoted submittal and an as-built configuration are explicit and auditable.

BOM and costing engine. Implements G5: from a completed selection it derives a granular bill of materials and a price build-up—casing, coil, fan, motor, filters, dampers, doors, accessories, and additional costs—with discount and margin treatment, in the operator's currency.

Governance and integration layer. Implements G6: role-based access (sales engineer configures, technical lead approves, procurement generates BOMs), single sign-on against the corporate directory, structured export to ERP/spreadsheet systems, and the audit trail that ties every output to a user, a timestamp, a certified rating, and a software version.

3.3 The module portfolio

Six modules consume the engine. Table 1 summarises each: its equipment class, deployment form factor, the principal engineering computation it performs, the governing standards it answers, and its primary outputs. The modules are deliberately heterogeneous in form factor—a Windows desktop application for the computation-heavy chiller rating, and browser-based applications for the workflow-heavy AHU, FCU, ERV, DX, and lifecycle modules—while sharing the same certified data and governance discipline.

3.4 Deployment, multi-tenancy, and integration

The web modules are deployed as multi-user, multi-branch applications: the production AHU system, for example, is organised by regional branch and financial year, tracks inquiries through an explicit allocation workflow, and is operated by a sizeable named-engineer team (Section 4.1). The platform is currency- and market-aware—production deployments span the Indian market and the Gulf Cooperation Council (GCC), with selections and offers denominated accordingly (for example, in Indian rupees or UAE dirhams) and with application categories appropriate to each market. Integration follows G5–G6: single sign-on against the corporate directory, structured export to spreadsheets and ERP systems (including SAP and Oracle), so that engineering and procurement do not maintain parallel sources of truth. The combination of a shared engine with per-module form factors lets the platform place computational depth where it is needed (the desktop chiller engine) and collaborative workflow where *that* is needed (the web AHU/DX/ERV/FCU/LCA applications) without fragmenting the certified data.

Table 1. The six MileSoft modules and the certified-curve engine they share. Computations and standards are detailed in Section 4.

Module	Form factor	Principal computation	Governing standards	Primary outputs
AHU selection	Web	Section-by-section configuration; coil/fan/filter/damper sizing; psychrometric duty; granular BOM cost build-up	EN 1886, EN 13053 (EUROVENT ECP)	Certified submittal, BOM, branded offer
FCU rating & selection	Web	Chilled-water coil rating; sensible/latent split; coil geometry; SFP; bulk selection	AHRI/EUROVENT coil rating	Performance sheet, dimensional data, BOM
ERV selection	Web	Air-to-air effectiveness (sensible/latent/total); dual-stream pressure drop; supply/return fan selection	AHRI 1060/1061, ASHRAE 84; ASHRAE 90.1 (mandate)	Effectiveness report, leaving states, fan data
Chiller rating	Desktop	Bivariate AHRI rating interpolation; staged part-load; full-load, IPLV, NPLV, ISEER	ANSI/AHRI 550/590, 551/591; IS 16590	Rating report (IPLV, NPLV, ISEER); GA; wiring
DX selection	Web	Total/sensible capacity matching at entering/ambient conditions; multi-circuit; line-set; refrigerant inventory	AHRI unitary rating; EU F-Gas (GWP)	Selection offer, capacity report, BI
Lifecycle cost (LCA)	Web	Discounted multi-year TCO; energy bins; ROI/breakeven/NPV; electrical-infrastructure and carbon accounting	ISO 15686-5, NIST HB 135	TCO comparison, payback, sensitivity

3.5 Governance, versioning, and auditability

The governance layer is what elevates the platform from a calculator to a certifiable engineering system. Three properties matter. First, *single-source propagation*: because all modules read one certified store, a library revision is atomic across the estate (G2). Second, *provenance*: every quoted number is bound to a certified rating and a software version, satisfying the EUROVENT requirement that each certified quotation record the selection-software version [12]. Third, *separation of duties*: role-based workflows ensure that configuration, approval, and procurement are distinct, auditable steps. Together these yield the outcome reported in Section 5—selections that are "audit-ready out of the box," with the audit trail a structural consequence of the architecture rather than a manual reconstruction.

4 Computational Methods by Module

This section documents the engineering method of each module, illustrated throughout with screens and data from the live production system. Unless noted, figures are unmodified production screenshots; numerical values quoted in the text are those of the reference selection shown in the corresponding figure.

4.1 Air handling unit selection

The AHU module is the platform's highest-volume tool and its closest embodiment of design goals G1–G6. It treats an AHU the way a factory builds one: as an ordered assembly of sections, each configured against certified component data, costed component-by-component, and quoted within a governed project.

4.1.1 Section-by-section configuration

An AHU is assembled from a governed catalogue of section types—in the production master these include inlet/mixing box, pre-filter, bag filter, fine filter, HEPA filter, coil, fan section, empty plenum, and empty heater section, among others. Each section carries its own engineering dialogue. The mixing-box dialogue (Fig. 2) configures fresh- and return-air dampers, fresh-air filter grade (e.g. G4), an additional damper with its own airflow in CFM, louvre, section length in millimetres, and door type; the HEPA dialogue specifies filter grade (e.g. H14), mounting frame, media, and—critically—the initial, final, and clean pressure drops (IPD/FPD/PD, in pascals) that enter the cumulative air-path pressure drop $\Delta p_{\text{tot}} = \sum_k \Delta p_k$ (Section 2). The pre-filter dialogue supports primary and secondary stages each with configuration, grade, depth, and media. Casing geometry is sized from an airflow-to-dimension configuration master (for example, a "blow-through, double" arrangement mapping each design airflow in CFM to a casing length, width, and height), and unit types (floor-mounted, ceiling-suspended, vertical) are governed centrally—so that the casing thermal and mechanical behaviour can be held within the EN 1886 classes that EUROVENT certifies [10, 11]. Configuration rules (G3) prevent infeasible section orderings from being quoted.

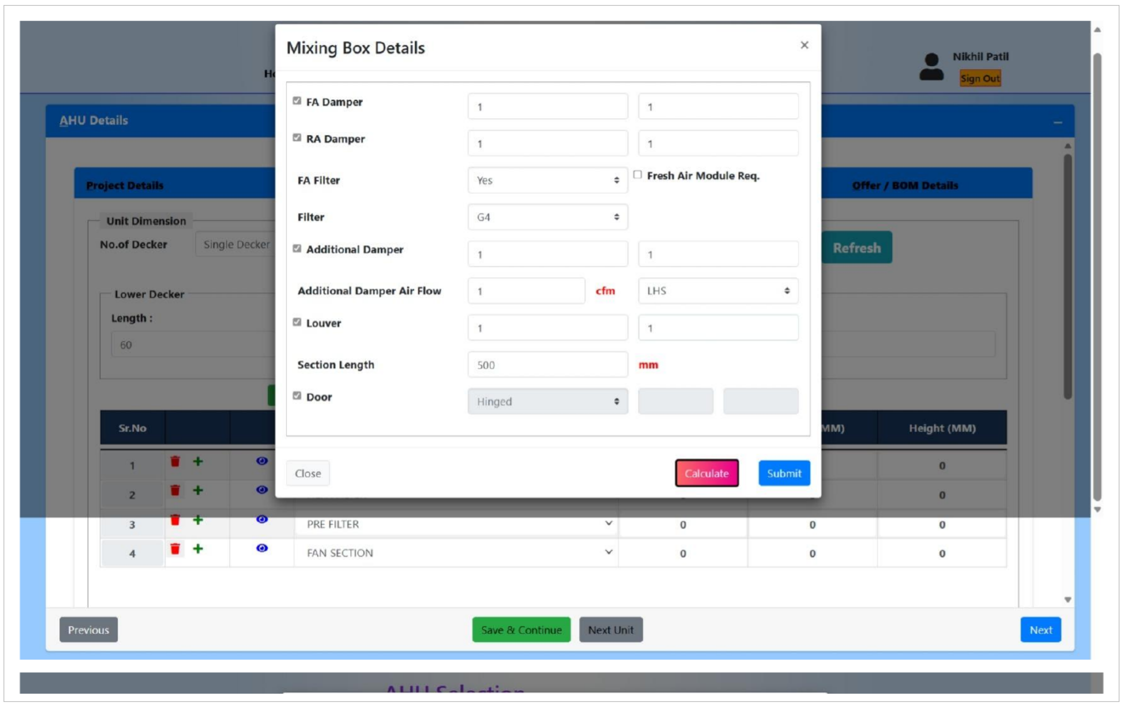

Figure 2. AHU section-level configuration (mixing-box dialogue). Each section—mixing box, filters, coil, fan—carries a dedicated engineering dialogue; filter dialogues capture the initial and final pressure drops (in pascals) that accumulate along the air path. Section ordering and component feasibility are rule-enforced (design goal G3).

4.1.2 Granular bill of materials and costing

Once the sections are configured, the BOM and costing engine (G5) derives a component-level price build-up. In the reference unit of Fig. 3 the build-up resolves casing, coil, damper, filter, fan, motor, door, eliminator, additional accessories, canopy, and cowl into a per-tag unit price and a project total (here a unit price of 534,850 at a quantity of two, for a unit total of 1,069,700, in the operator's currency), with explicit discount and margin fields. Because this BOM is emitted by the same engine that performed the selection, the submittal, performance sheet, and price quotation are mutually consistent by construction—there is no separate spreadsheet to fall out of step. The unit dialogue also captures the engineering qualifications that protect margin and manage risk: structured *notes*, *deviations*, and *exclusions* (for example, recording the gauges, valves, and piping accessories that are *not* in scope), which flow into the branded offer.

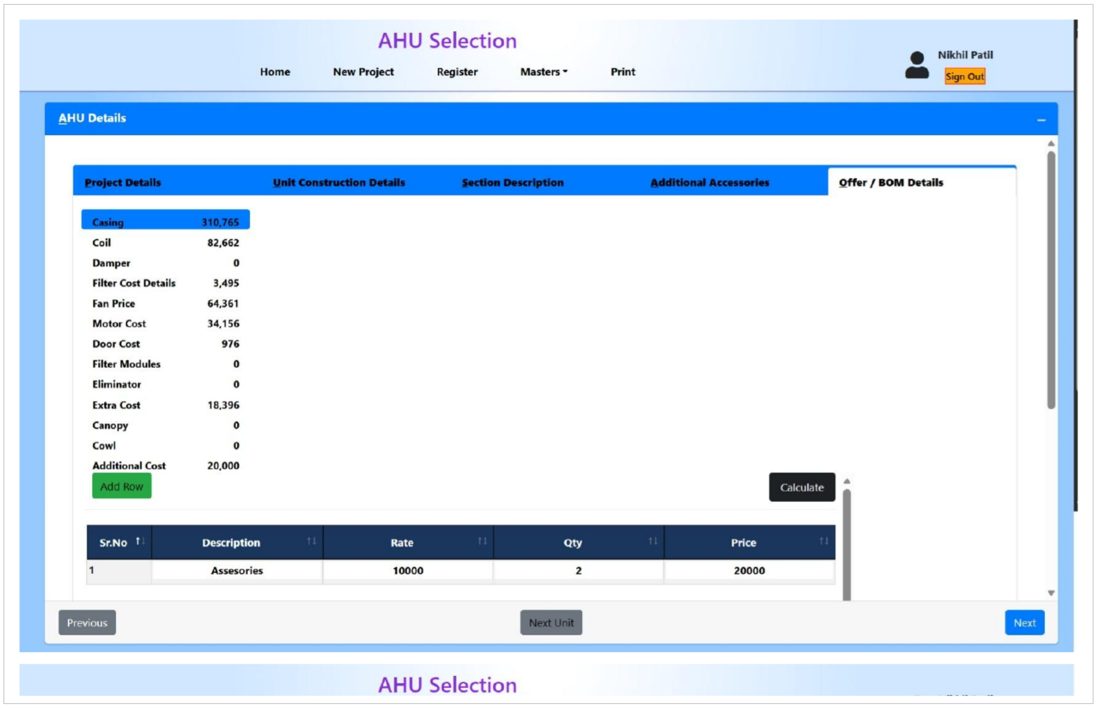

Figure 3. Granular BOM and price build-up for an AHU. The costing engine resolves casing, coil, fan, motor, filters, dampers, doors, and accessories into a per-tag unit price and project total, with discount and margin treatment, emitted from the same engine that performed the selection (design goal G5).

4.1.3 Inquiry workflow at production scale

The module is not a single-user calculator but a governed, multi-branch workflow application. Figure 4 shows the production dashboard: inquiries are tracked by branch and financial year through an explicit state machine—*pending for allocation*, *allocated but offer not created*, *new inquiry allocated*, and *offer created*—with per-engineer attribution, one-click export to spreadsheet, and a team-statistics view. At the time of capture the system held on the order of 2.8×10^3 inquiries across the estate, operated by a named-engineer team of roughly seventy-five users, with role-based separation between configuration, approval, and procurement (G6). This is the operational substrate that makes "audit-ready out of the box" a structural property: every offer is bound to an inquiry, an engineer, a timestamp, and the certified component data and software version used (G1, G6) [12].

4.2 Fan coil unit rating and selection

Terminal units are frequently the weakest link in HVAC selection—small enough that they are often sized from catalogue tables and a rule-of-thumb correction, yet numerous enough that an error repeated across a building surfaces as a comfort complaint and a warranty claim. The FCU module brings the same certified-curve rigour to terminal units that the AHU module brings to central plant, sharing the same engine (G1) so that a coil-family revision benefits FCUs and AHUs simultaneously (G2).

A correct FCU selection is a coil-rating problem evaluated at the room's *actual* entering air condition

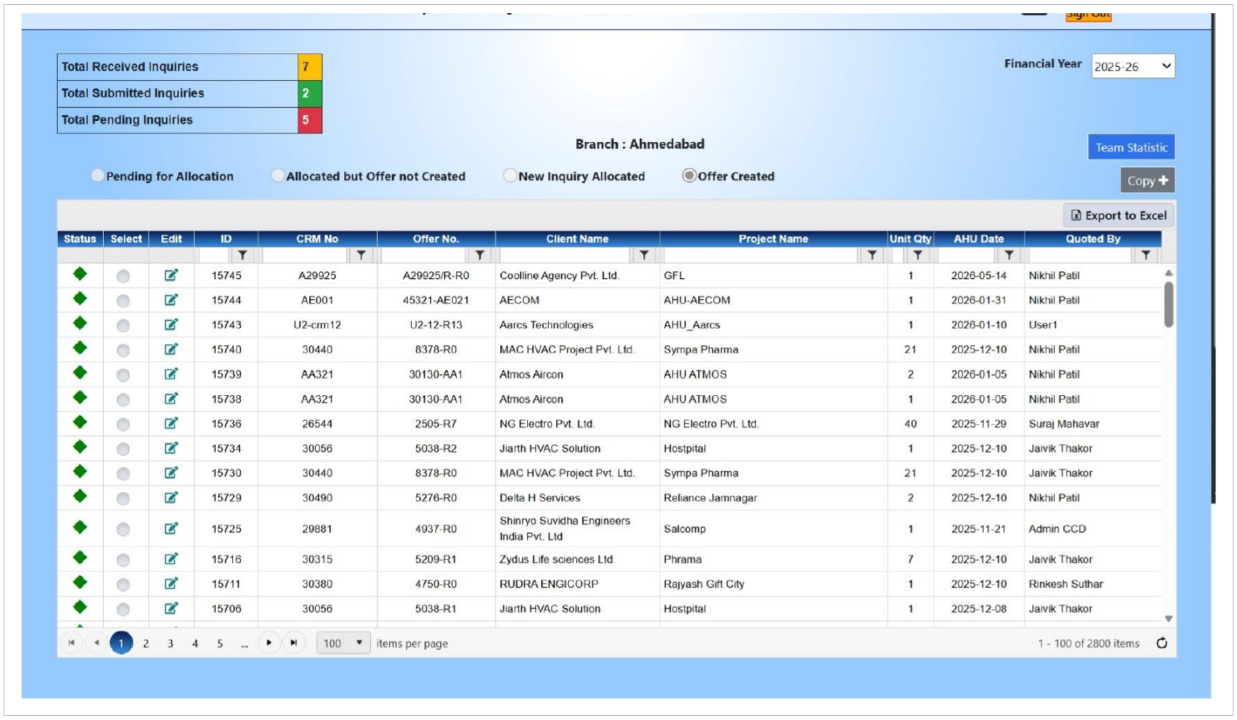

Figure 4. Production AHU inquiry dashboard. Inquiries are governed by branch and financial year through an explicit allocation state machine with per-engineer attribution and spreadsheet export. The captured estate holds on the order of 2.8×10^3 inquiries operated by a team of roughly seventy-five engineers (design goals G4, G6).

rather than a catalogue reference point. The rating dialogue (Fig. 5) takes the air entering condition (in the reference selection, 25 °C dry-bulb, 18.67 °C wet-bulb, 54.88 % relative humidity), the chilled-water condition (entering 5.5 °C, leaving 14.4 °C, $\Delta T_w = 8.90\,°\mathrm{C}$), and the coil geometry (tube pattern, three rows, three-circuit, fin spacing 2.11 fpi, face velocity 1.12 m/s), and returns the full rated state via the psychrometric core (Eq. 4). For the reference unit this yields a total capacity of 1.19 kW and a sensible capacity of 0.97 kW (a sensible heat ratio of $\mathrm{SHR} \approx 0.82$), an off-coil condition of 17.2 °C dry-bulb / 15.46 °C wet-bulb, a water velocity of 0.15 m/s with a water-side pressure drop of 0.567 kPa, and an input power of 61.79 W at an external static pressure of 50 Pa—from which the specific fan power $\mathrm{SFP} = P/\dot{V}$ is reported (here 0.61). Reporting SFP at selection time directly attacks the specific-fan-power underestimation that empirical studies tie to the performance gap [3]. The result row also fixes the constructional and electrical specification (ducted single-skin construction, a $\frac{1}{2}$-inch aluminium filter, a 220–240 V single-phase 50 Hz supply). At project scale the module supports *bulk selection*—a room list is selected in one pass, with units that require engineering judgement flagged—and acoustic guardrails that prevent the engine from “successfully” selecting a unit that would breach a configured noise budget. A deep master-data layer (casing material, coating, and thickness; coil casing, coating, and deck insulation; construction; drain-pan coating; and so on) keeps the constructional library governed (G2).

4.3 Energy recovery ventilator selection

Energy recovery has moved from optional to code-mandated across much of the commercial stock, with ASHRAE 90.1 triggering recovery above a supply-airflow threshold at high outdoor-air fractions and raising the required enthalpy-recovery ratio toward 80 % in the 2025 cycle [15]. The engineering consequence is that an ERV must be defended on effectiveness *and* on the fan-power penalty its pressure drop imposes—across the

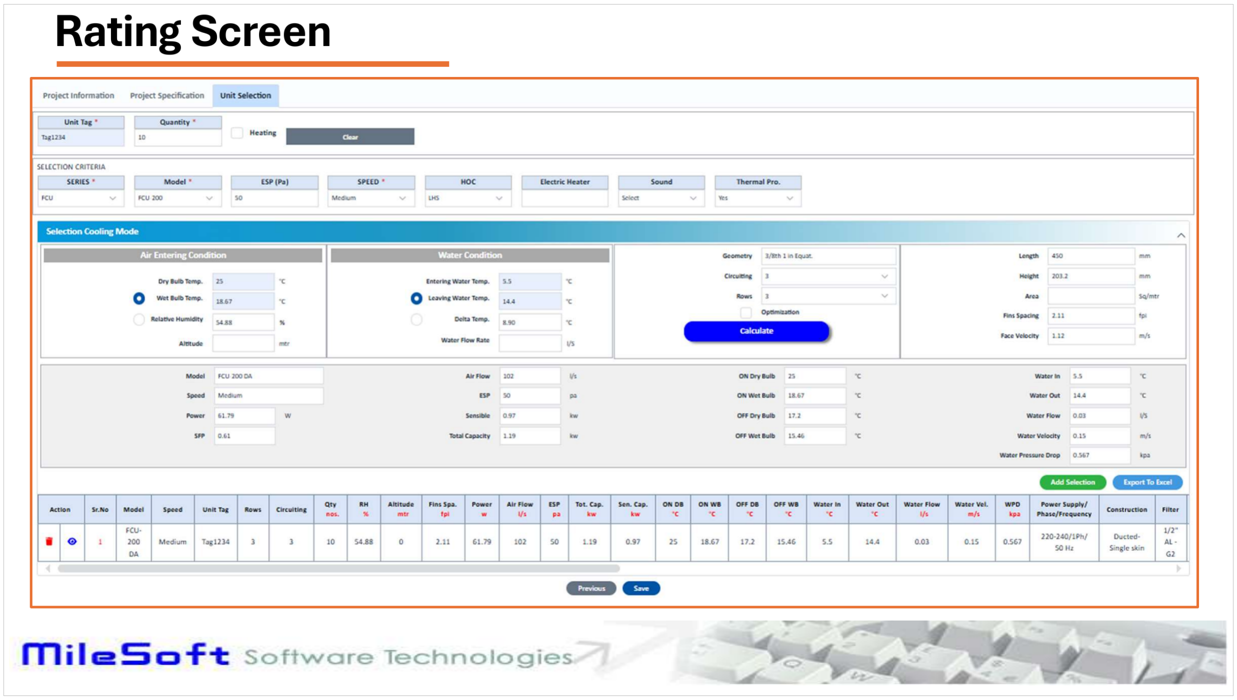

Figure 5. FCU chilled-water coil rating. The engine evaluates the coil at the actual entering air and water conditions and the specified geometry (rows, circuiting, fin spacing), returning total and sensible capacity, off-coil state, water velocity and pressure drop, input power, and specific fan power. Reporting SFP at selection time addresses a documented source of the performance gap [3].

operating envelope, not at a single point.

The ERV module is a psychrometric effectiveness engine implementing the AHRI 1060/ASHRAE 84 definitions (Eq. 3). The dialogue (Fig. 6) takes the fresh-air and return-air states and external static pressures (in the reference selection, fresh air at 250 CFM, 34 °C dry-bulb / 18 °C wet-bulb; return air at 200 CFM, 24 °C dry-bulb / 17 °C wet-bulb; 100 Pa external static on each side) and returns the full leaving state on both streams. For the reference unit the temperature effectiveness is 79.3 %, with the supply (fresh) air leaving at 27.9 °C dry-bulb / 17.2 °C wet-bulb / 34 % relative humidity (absolute humidity 9 g/kg, enthalpy 46.5 kJ/kg) and the exhaust air leaving at 31.7 °C dry-bulb (enthalpy 50 kJ/kg); the engine reports supply- and exhaust-side pressure drops separately (56 Pa and 44 Pa) and face velocities, exactly as Standard 1060 requires so that the recovery benefit can be netted against the fan-power penalty [7, 9]. Because pressure drop loads *both* the supply and exhaust fans, the module sizes both: it reports, for the supply and return fans, the external and total static pressures, working and maximum speeds, static efficiency, supply voltage/frequency, and power. A cooling-coil option supports post-recovery conditioning, and projects are retained in an auditable log. By computing the leaving states from certified effectiveness at the project's own entering conditions—rather than quoting a single summer-design number—the module gives the ventilation engineer a defensible answer to "what ERV did you select, and why," across the annual envelope the code now scrutinises.

4.4 Chiller rating

The chiller module is the platform's most computation-intensive component and is deployed as a Windows desktop application. It implements the full ANSI/AHRI 550/590 (and SI 551/591) rating method together with the Indian ISEER metric, so that a single rating run produces every efficiency figure a consultant, an auditor, or

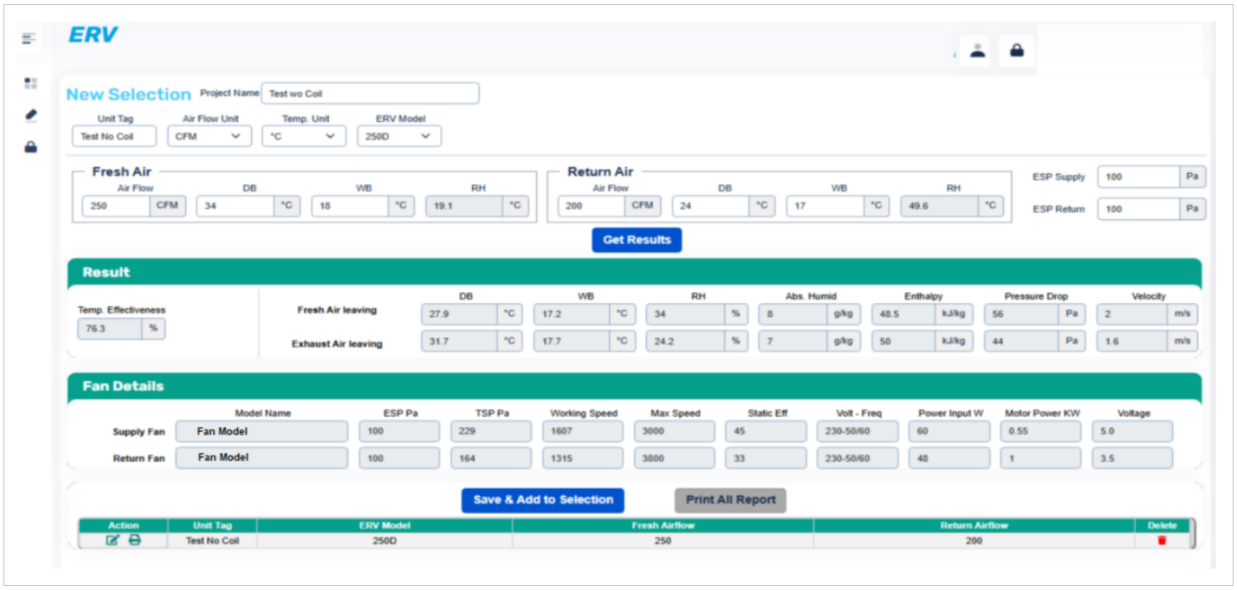

Figure 6. ERV psychrometric effectiveness engine. From the fresh- and return-air states and external static pressures, the engine computes temperature effectiveness and the full leaving state on both streams (dry-/wet-bulb, relative and absolute humidity, enthalpy), reports supply- and exhaust-side pressure drops separately per AHRI 1060, and sizes both the supply and return fans.

a BEE submission may demand [4, 5, 13].

4.4.1 Performance model and master data

Each chiller model is described by a structured master record spanning general data (compressor type—screw or scroll—refrigerant, oil, evaporator/condenser models, connection sizes, fan data), electrical and physical data, certified performance ratings, separate part-load, fan-power, and compressor-power multipliers, and explicit AHRI and IS rating-condition sets. Part-load behaviour is represented over staged load bands (in the reference model, six stages spanning 0–25, 26–38, 39–50, 51–75, 76–88, and 89–100% with optional tonnage modulation), capturing the unloading characteristics of variable-speed machines. The certified performance itself is a *bivariate* map: Fig. 7 shows capacity (in tons) and power (in kW) tabulated against leaving chilled-water temperature and condensing temperature (for example, at a leaving chilled-water temperature of 44°F and a condensing temperature of 55°F the model delivers 182.12 TR at 85.2 kW). The interpolation core evaluates capacity and power *between* these certified points, and the master allows new leaving-water and ambient rows to be added so the certified envelope can be extended without leaving the governed store (G1, G2).

4.4.2 Rating, IPLV, NPLV, and ISEER in one pass

The rating screen (Fig. 8) takes the project duty—in the reference rating, an ambient of 32 °C, cooler outlet/inlet of 7/12 °C, a flow of 53.15 L/s, an evaporator pressure drop of 146.52 kPa, and a fouling factor of 1.8×10^{-5} m^2 K/W—and produces the AHRI part-load table directly. For the reference machine this table reads 301.7 TR at 332.5 kW (a unit EER of 10.89, or 1.102 kW/ton) at 100 % load and 35 °C ambient; 226.3 TR at 173.1 kW (0.765 kW/ton) at 75 %; 150.9 TR at 93.2 kW (0.618) at 50 %; and 75.4 TR at 39.9 kW (0.529) at 25 %. Substituting these four kW/ton values into Eq. 1 yields the integrated part-load value the screen reports as IPLV.SI = 0.6609, while the parallel reducing-condenser-water (ISEER) table yields ISEER.SI = 0.8007 (Eq. 2); a worked reconstruction of the IPLV figure from the four points is given in Appendix B. The engine

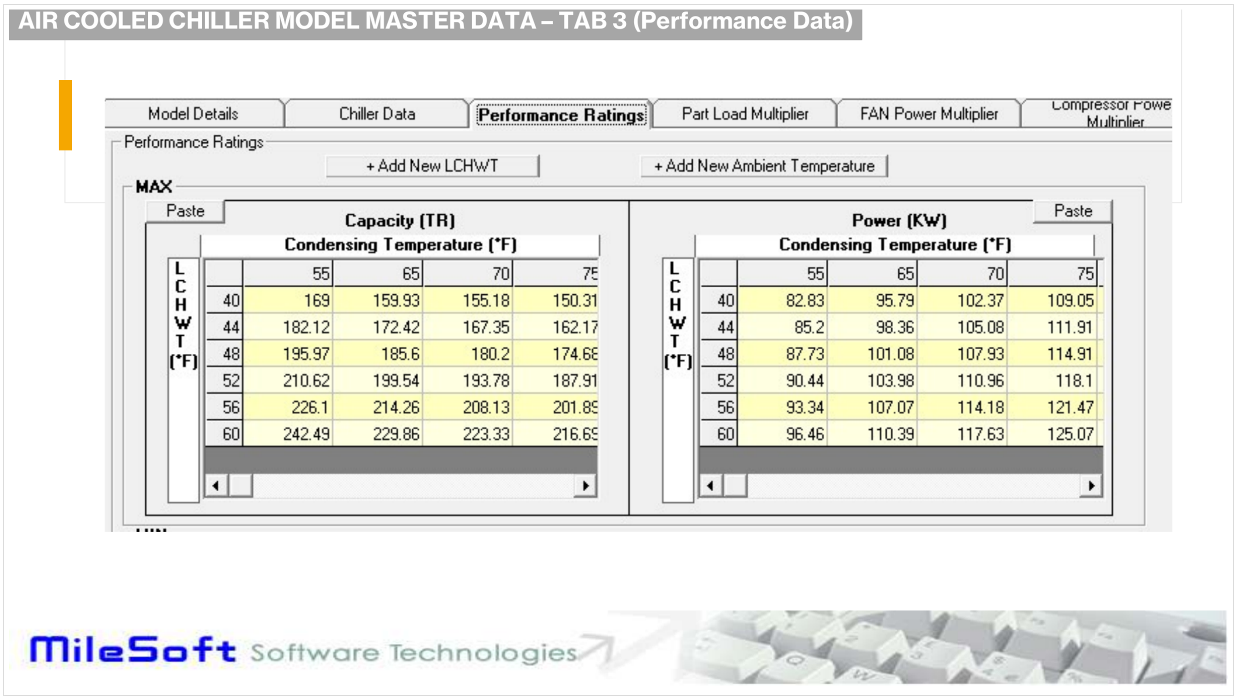

Figure 7. Certified bivariate chiller performance map. Capacity (tons) and power (kW) are tabulated against leaving chilled-water temperature (rows) and condensing temperature (columns); the engine's interpolation core evaluates performance between certified points. The certified envelope is extensible within the governed store.

computes full-load efficiency, IPLV, NPLV, and ISEER in a single pass, with NPLV configurable to the project's own condenser-water trajectory and ambient bin profile—the number that, as the standard itself notes, best predicts site energy for a specific installation [6]. Output formats are selectable (AHRI, NPLV, NPLV 10-point, PLV, PLV 10-point, UPLV, ISEER), and the rating bundles the general-arrangement drawing, wiring diagram, and model photograph, with the calculation methodology disclosed so the submittal survives consultant review. Refrigerant data carried in the master (here R134a) supports GWP/TEWI reporting for F-Gas exposure [17].

4.5 Direct-expansion selection

Direct-expansion systems—splits, packaged rooftops, multi-circuit units, and variable refrigerant flow—are deceptively complex once refrigerant charge, line-set effects, and multi-circuit balancing are accounted for, yet most tooling stops at "pick a model from a catalogue." The DX module is a browser-based selector that couples engineering selection with a light customer-relationship workflow. A three-step wizard (customer information $\rightarrow$ unit selection $\rightarrow$ selection preview) captures both the engineering brief and commercial context—project type and stage, win probability, and expected closing date—so that selection and pipeline live together.

The selection itself is a capacity-matching problem at the project's entering and ambient conditions. As Fig. 9 shows, the engine matches required total and sensible capacity (in the reference selection, a required 30 MBH total and 25 MBH sensible—an implied SHR ≈ 0.83—at 1000 CFM design airflow, an indoor coil entering condition of 80/67°F dry-/wet-bulb, and a 115°F ambient) and returns conforming models with a final price in the operator's currency (here UAE dirhams, reflecting the GCC deployment). Consistent with the platform's broader method, the module is designed to model capacity, EER, and degradation circuit-by-circuit for multi-circuit units, to correct performance and refrigerant charge for actual line-set length and vertical rise, to sweep capacity and energy across the ambient envelope rather than at a single rating point, and to carry a

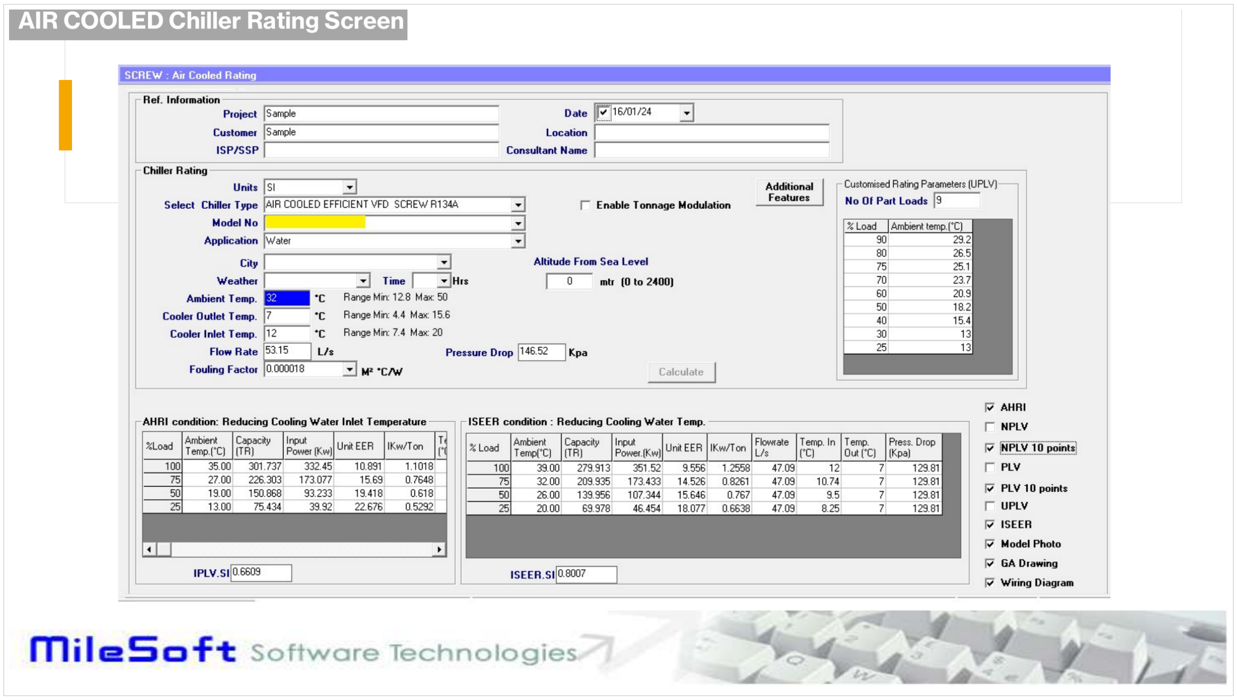

Figure 8. Chiller rating output. From the project duty, the engine produces the AHRI part-load table at the 100/75/50/25% points and computes full-load efficiency together with IPLV (0.6609 here), NPLV, and ISEER (0.8007 here) in one pass, with NPLV configurable to the project's condenser-water and ambient bin profile. Output formats and accompanying drawings are selectable.

refrigerant inventory with GWP for environmental compliance [17]—all drawing coil performance from the same certified library as the AHU and FCU modules (G1).

4.6 Lifecycle cost analysis

First-cost selling is giving way to total-cost-of-ownership procurement: commercial tenders now routinely request a multi-year lifecycle cost comparison, and the team that can produce a clean, defensible one wins more work. The LCA module—deployed as a guided nine-step wizard supporting single-unit, multi-unit, and new-construction scenarios—computes the discounted lifecycle cost of Eq. 5 and, decisively, draws its energy numbers from the *actual* selection rather than a parallel estimate (G1): when the upstream AHU, chiller, or DX selection changes, the lifecycle analysis updates with it.

The wizard builds the comparison in stages: existing and proposed unit cost details (equipment cost, service-cost percentage, and installation cost per option); a year-by-year energy model; return-on-investment and breakeven; an electrical-infrastructure analysis; and a carbon account. Figure 10 shows the multi-year energy model—a fifteen-year lifecycle energy and cost analysis in which each existing and proposed option carries an annual consumption with year-on-year escalation (in the reference case, an existing unit at roughly 5880 kW h to 5998 kW h per year against a proposed option at roughly 4380 kW h to 4468 kW h per year), an annual derating factor capturing efficiency loss with age, and energy priced over the horizon. Figure 11 shows the resulting return-on-investment schedule: a year-by-year accounting of existing-versus-proposed service cost, energy-cost saving, and remaining (un-recovered) cost, with the breakeven year flagged explicitly (here, recovery in year 6) and each subsequent year marked as net saving. From these cash flows the module reports net present value and discounted payback for the chief-financial-officer conversation, together with sensitivity

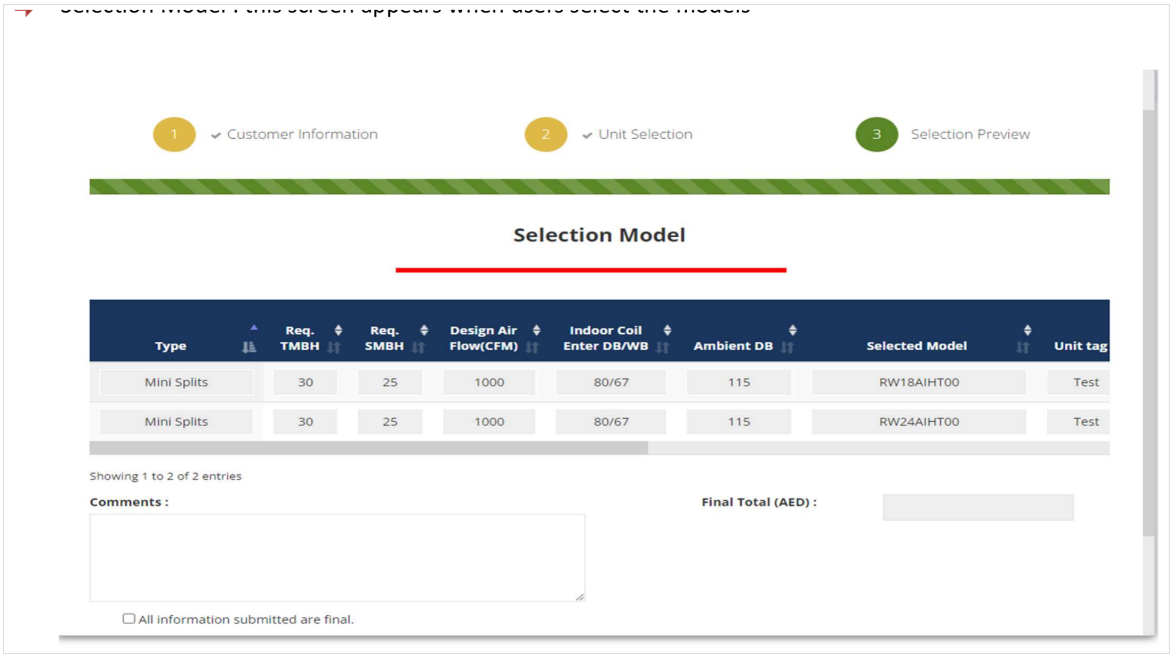

Figure 9. DX selection. The engine matches required total and sensible capacity (MBH) at the specified indoor entering and ambient conditions and design airflow, returning conforming models with project pricing. Selection is embedded in a customer-relationship workflow that captures project stage, win probability, and expected close.

analysis over the energy-escalation rate.

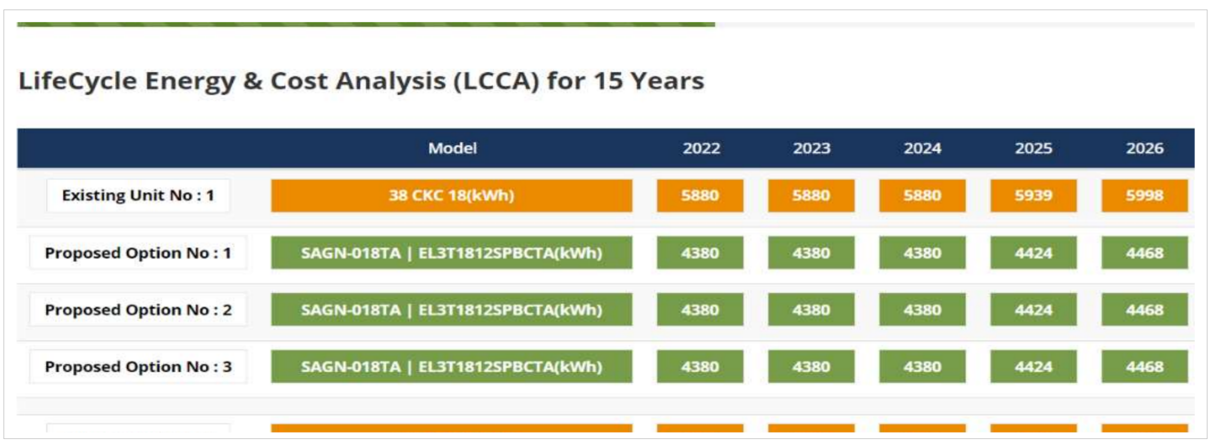

Figure 10. Fifteen-year lifecycle energy and cost analysis. Each existing and proposed option carries an annual consumption with year-on-year escalation and an age-related derating factor; energy is priced over the horizon. The energy figures derive from the actual upstream selection rather than a parallel estimate (design goal G1).

Two further dimensions distinguish the module. The *electrical-infrastructure* analysis translates connected load into utility-side consequence: with a power factor pf and a transformer utilisation u, the apparent power and transformer sizing are $\text{kVA} = L/pf$ and $\text{kVA}_{\text{tx}} = \text{kVA}/u$ (in the reference configuration $pf = 0.8$ and $u = 0.85$), so that a more efficient selection's reduced connected load is monetised as deferred transformer capacity. The *carbon* account exposes the operational-emissions consequence of the choice through a parameterised model $\text{CO}_2 = L \times f_{\text{ef}} \times h$, where L is the connected load, f_{ef} is a grid emission factor, and h an annual operating-hours

TAB 5 – ROI Analysis

UnitNo	OptionNo	Year	Existing Unit Service Cost (AED)	Proposed Unit Service Cost (AED)	Service Cost Savings (AED)	Energy Cost Savings (AED)	Remaining Cost (AED)	Recovery Status
Existing Uni	Proposed Option No : 3							
		2021	1000	0	1000	450	6550	
		2022	1100	400	700	464	5386	
		2023	1210	420	790	491	4105	
		2024	1331	441	890	541	2674	
		2025	1464	463	1001	612	1061	
		2026	1610	486	1124	718	-781	Recovered in Year 5
		2027	1772	510	1262	864	-2907	Savings
		2028	1949	536	1413	1066	-5386	Savings
		2029	2144	563	1581	1348	-8315	Savings
		2030	2358	591	1767	1746	-11828	Savings

Figure 11. Return-on-investment and breakeven schedule. A year-by-year accounting of service cost, energy-cost saving, and remaining un-recovered cost identifies the breakeven year explicitly (here year 6), from which net present value and discounted payback follow (Eq. 5).

scaling (in the reference configuration $f_{ef} = 0.41$ and $h = 2500$). The refrigerant cost and GWP trajectory can likewise be carried for buyers facing F-Gas phase-down [17]. Together these turn the lifecycle module into the business case that the engineering selections substantiate.

4.7 Management information and business intelligence

Because every selection is a governed record, the platform can aggregate engineering activity into portfolio intelligence—the management-information layer of Fig. 1. The DX module's business-intelligence view (Fig. 12) summarises pipeline value across the GCC by country (Bahrain, Kuwait, Oman, Saudi Arabia, the UAE), by project type (new construction versus replacement), by project stage (hibernation, job-in-hand, tender, won), by application (commercial, educational, mosque, residential), and by sales engineer, with table/chart toggling, date filtering, and export. The lifecycle module's management view (Fig. 12) tracks realised energy outcomes across the portfolio—per record, the country, customer, project stage, segment, and the existing-versus-proposed annual energy cost with the annual estimated saving. This layer is what converts per-selection engineering rigour into a managed sales-engineering operation: the same governed data that makes a single submittal auditable makes the whole pipeline measurable.

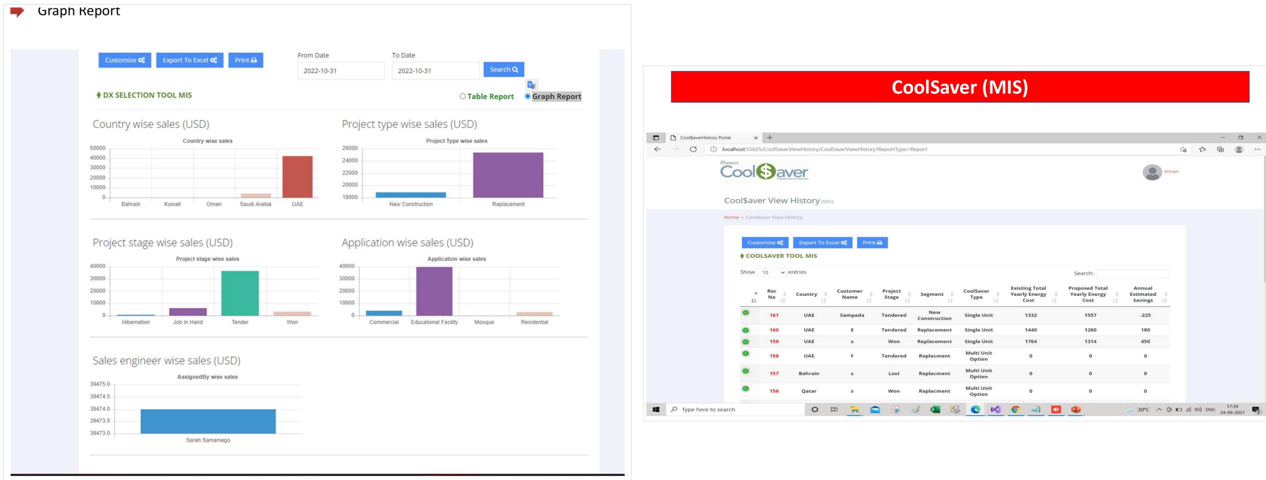

Figure 12. Management-information and business-intelligence layer. Left: DX pipeline analytics across the GCC by country, project type, stage, application, and engineer. Right: lifecycle portfolio view tracking realised existing-versus-proposed annual energy cost and estimated savings per record. Governed selection records make the entire pipeline measurable.

5 Field Evidence: A Production Deployment

We report a production deployment of the platform at a EUROVENT-certified air-handling-unit manufacturer operating in India.[1] The deployment is presented as a single, in-depth case study—evidence of the architecture's effect in practice, not a controlled multi-site trial.

5.1 Context and prior workflow

The manufacturer is a certified AHU producer for whom every issued selection is, in effect, a certified engineering document: under the EUROVENT ECP programme, a certified AHU selection must be produced by integrated software and must record the software version used [10–12]. Prior to deployment, however, the selection workflow was the very arrangement Section 2 identifies as inadequate—general-purpose spreadsheets for the engineering calculation and a word processor for the submittal. That arrangement carried the structural liabilities catalogued earlier: certified performance lived in copies scattered across engineers' machines with no version control; submittals and the underlying calculations could fall out of step; the bill of materials was re-keyed by hand; and, when audited, the provenance of a given performance figure had to be reconstructed manually. For a certified manufacturer, these are not mere inefficiencies—they are direct exposure on the certification itself.

5.2 What was deployed

The manufacturer adopted the platform's AHU module (with the certified-curve engine, configuration rules, BOM/costing, and the governed inquiry workflow of Section 4.1) across a multi-branch, roughly seventy-five-engineer organisation. The deployment integrated with the enterprise stack along the G5–G6 dimensions: single sign-on against the corporate directory, and structured BOM export into the enterprise resource-planning system, so that engineering and procurement shared one source of truth. Selections were bound to certified component

[1] The deployment site is a EUROVENT-certified AHU original-equipment manufacturer. The operational metrics in this section are *operator-reported*, drawn from the deployment, and have not been independently audited by the authors; see Section 7 for the corresponding threats to validity.

libraries and a recorded software version, satisfying the ECP provenance requirement by construction.

5.3 Reported outcomes

The manufacturer reports the outcomes summarised in Table 2. Selection-cycle time fell by approximately 70 %; per-engineer proposal throughput tripled; the proposal-to-order conversion rate improved by 12 %; and, since rollout, the organisation reports *zero* audit findings against its issued selections. The mechanisms tie directly to the architecture. The cycle-time and throughput gains follow from eliminating manual re-keying and from the one-push library governance that removes the cross-checking overhead of reconciling scattered spreadsheets (G2, G5). The elimination of audit findings is the structural consequence of single-source provenance and recorded software versions (G1, G6): the audit trail no longer has to be reconstructed because it is a by-product of the workflow [12]. The conversion uplift is consistent with two effects the platform enables—faster, more professional, and internally consistent submittals, and the ability to attach a defensible lifecycle-cost case (Section 4.6) to a first-cost quotation.

Table 2. Operator-reported outcomes following platform deployment at a EUROVENT-certified AHU manufacturer. Metrics are self-reported and not independently audited (Section 7).

Dimension	Reported change	Architectural mechanism
Selection-cycle time	$\approx$70 % reduction	Eliminated re-keying; one-push library governance (G2, G5)
Proposals per engineer per month	$\approx 3\times$ increase	Section-based configuration; bulk workflow; shared engine (G1, G5)
Proposal-to-order conversion	+12 %	Consistent professional submittals; attachable lifecycle-cost case (G5; Sec. 4.6)
Audit findings since rollout	*Zero*	Single-source provenance; recorded certified rating and software version (G1, G6)

We emphasise the evidentiary status of these figures: they are reported by a single deploying organisation and are not independently verified. They are nonetheless informative. The direction and rough magnitude of every metric is what the architecture predicts, and the *zero-audit-findings* outcome in particular is a qualitative state-change—from a workflow in which provenance is reconstructed to one in which it is guaranteed—that is difficult to achieve by incremental improvement of a spreadsheet process and is the most direct corroboration of the platform's central design claim.

6 Discussion

6.1 Closing the specification side of the performance gap

The performance gap is the product of many causes spanning design, construction, commissioning, and operation [1, 3]. The platform does not claim to address all of them; it targets the one that is squarely an engineering-tooling problem—*specification uncertainty*. The mechanism is concrete. When a coil's duty is evaluated from certified data at the project's actual entering condition (Section 4.2) rather than read off a generic table, when a fan's specific fan power is reported at selection time rather than assumed (Sections 4.1–4.2), when an ERV's effectiveness and *both* fan penalties are computed across the operating envelope (Section 4.3), and when a chiller is judged on IPLV/NPLV/ISEER rather than full-load efficiency alone (Section 4.4), the assumptions that enter the building's energy model are tied to the equipment that will actually be installed. Specific fan power and load underestimation—the two specification-side culprits that empirical reviews name explicitly [3]—are precisely the quantities the platform forces into the open. The platform cannot guarantee a closed gap, but it removes a category of error that begins at the keyboard.

6.2 Auditability and certification as architecture, not afterthought

The deployment's zero-audit-findings outcome (Section 5) illustrates a broader thesis: in a certified-manufacturing context, auditability must be a property of the *architecture*, not a manual exercise performed after the fact. The EUROVENT ECP requirement—integrated software, components selected as one entity, software version recorded on every certified quotation [12]—is, read literally, a software-architecture specification. A spreadsheet estate cannot satisfy it except by heroic manual discipline; a single governed engine satisfies it by construction. This reframing matters beyond EUROVENT: as energy codes tighten and as embodied-carbon and refrigerant-GWP disclosures become contractual [15, 17], the ability to bind every quoted number to a certified, versioned source becomes a competitive necessity rather than a compliance nicety.

6.3 From first cost to lifecycle value

By drawing the lifecycle module's energy figures from the actual selection (Section 4.6), the platform structurally couples the engineering decision to the financial one. This is consequential for the performance gap from the demand side: a procurement process that can see and price the fifteen-year energy and carbon consequences of a selection (Figs. 10–11) is far more likely to choose the efficient option than one comparing first costs alone. The electrical-infrastructure and carbon accounts extend the argument—reduced connected load is monetised as deferred transformer capacity, and operational emissions are made explicit—so that efficiency is defended in the buyer's own financial language.

6.4 A capability reference framework for selection platforms

The paper's final contribution is to generalise from one platform to a *benchmark*: a structured set of capabilities that a world-class HVAC selection-and-lifecycle platform should provide, derived from the governing standards (Section 2) and the failure modes of incumbent tooling. We propose eleven criteria, grouped into four families:

Performance integrity.

(C1) Certified-performance traceability—every output bound to a specific certified rating; **(C2)** one-push library governance—revisions propagate atomically; **(C3)** evaluation at actual project conditions via an interpolation/psychrometric core, not catalogue reference points.

Engineering breadth and correctness.

(C4) Cross-class composition—AHU, chiller, FCU, ERV, and DX selected within one project; **(C5)** enforced feasibility—infeasible configurations cannot be quoted; **(C6)** standards conformance across the relevant regimes (AHRI 550/590 & 1060, EN 1886/13053, IS 16590, ASHRAE 90.1, F-Gas).

Commercial and lifecycle integration.

(C7) Integrated BOM and costing from the same engine; **(C8)** lifecycle-cost integration drawing energy from the actual selection; **(C9)** enterprise integration—single sign-on and clean ERP export.

Governance and intelligence.

(C10) Project lifecycle, versioning, and audit provenance; **(C11)** portfolio analytics over governed records.

Table 3 scores the four incumbent approaches of Section 2 against these criteria. The intent is not to claim that no tool offers any capability, but to make explicit *which* capabilities each structural approach can and cannot provide. Spreadsheets, however expertly built, cannot provide traceability, governance, enforced feasibility, or audit provenance, because those are properties of a governed system rather than of a calculation. Single-vendor selectors provide deep correctness *within* a catalogue but cannot compose across equipment classes or carry a manufacturer's own BOM and branding. Point tools solve one class well but re-create the integration gap. A unified certified platform is the only structural approach that can satisfy all eleven—and, by the same token, the only one that can satisfy the EUROVENT integrated-software requirement [12]. We offer the framework as a reusable benchmark: a checklist against which any selection platform, including this one, can be honestly evaluated.

Table 3. Capability reference framework. Structural ability of each incumbent approach to satisfy the eleven criteria. ● = structurally provided; ○ = partial / catalogue-bounded; — = not structurally provided.

Capability criterion	Spread-sheets	Single-vendor selector	Point tools	Unified certified platform
C1 Certified-performance traceability	—	●	○	●
C2 One-push library governance	—	○	—	●
C3 Evaluation at actual conditions	○	●	○	●
C4 Cross-class composition	○	—	—	●
C5 Enforced feasibility	—	○	○	●
C6 Standards conformance (broad)	○	○	○	●
C7 Integrated BOM and costing	○	—	—	●
C8 Lifecycle-cost integration	○	—	○	●
C9 Enterprise integration (SSO/ERP)	—	○	—	●
C10 Project lifecycle, versioning, audit	—	○	—	●
C11 Portfolio analytics	—	○	—	●

6.5 Generalisability

Although instantiated for AHU, FCU, ERV, chiller, DX, and lifecycle costing across the Indian and GCC markets, the architecture is not specific to those classes or markets. The certified-curve engine is agnostic to equipment type: any device whose certified performance can be expressed as governed curves or rating maps—cooling towers, pumps, variable-refrigerant-flow systems, dedicated outdoor-air systems—can be added as a module drawing on the same store, governance, and lifecycle services. Likewise, the standards layer is parameterised: the same IPLV/NPLV machinery serves AHRI markets and, with the ISEER bin set, the Indian market. The framework of Table 3 is fully general.

7 Threats to Validity and Limitations

We state the limitations of this work plainly, as they bound the strength of its claims.

Evidentiary strength of the field results. The outcomes in Section 5 derive from a *single* deploying organisation and are *operator-reported* rather than independently audited. There is no control group, and the study is observational; we therefore cannot exclude confounders such as concurrent process improvement, team learning, or selection and reporting bias, nor a Hawthorne effect during rollout. The reported magnitudes (e.g. "$\approx$70% faster") are sensitive to how the pre-deployment baseline is defined. The metrics should be read as a directionally consistent, mechanistically plausible case study, not as a quantitative effect size with confidence bounds.

Construct and external validity. The case site is one EUROVENT-certified manufacturer in one market; deployments in other segments, scales, or regulatory regimes may differ. The platform's per-module data points (Section 4) are illustrative reference selections from the production system, not a randomised sample of selections.

Dependence on library governance. The platform's correctness is conditional on its certified library being correctly populated and maintained. The architecture guarantees *propagation and provenance*—that whatever is in the governed store is applied consistently and traceably—but it cannot certify the upstream test data itself; that remains the manufacturer's metrological responsibility.

Modelling simplifications. Several models are intentionally parameterised rather than first-principles. The carbon account uses an emission-factor-times-hours parameterisation (Section 4.6) rather than an hourly grid-marginal analysis, and IPLV/NPLV is, by the standard's own admission, a weighted four-point surrogate for a full hourly simulation [6]. These are appropriate for selection-stage decision support but are not substitutes for detailed annual energy modelling.

Standards drift. The conformance claims are tied to specific standard editions (e.g. AHRI 550/590, ASHRAE 90.1 cycles, EU F-Gas 2024/573). These evolve; sustained conformance requires the platform to track revisions—a maintenance obligation, not a one-time property.

No public benchmark artefact. We propose a capability reference framework (Section 6) but do not, in this paper, release an open scoring rubric or dataset against which third parties can reproduce the comparison. We regard this as the most important item of future work.

8 Future Work

Several directions would strengthen both the platform and the evidence for it.

Independent, multi-site evaluation. The clearest next step is a controlled, multi-organisation study with defined baselines and, ideally, pre-registered metrics—together with the release of an anonymised benchmark dataset and an open scoring rubric for the Table 3 criteria, so that the comparison becomes independently reproducible.

Building-information-modelling interoperability. Bidirectional exchange with BIM authoring tools via open schemas (gbXML, IFC) would let selections round-trip with the building model: loads and geometry flow in, selected equipment with certified performance flows back, eliminating a re-entry boundary that is itself a source of specification error.

Cloud-native re-platforming and an open API. Migrating the computation-heavy desktop chiller engine to a cloud-native, API-first service would unify the suite's form factor, enable headless integration into customers' own quotation pipelines, and simplify the atomic library propagation that underpins G2.

Machine-learning-assisted selection. The governed corpus of historical, accepted selections is a training signal: models could recommend likely-optimal configurations from the project brief, flag selections that deviate from historical norms, and detect margin anomalies—augmenting, not replacing, the deterministic certified engine.

A digital-twin feedback loop. Most ambitiously, ingesting metered operating data and comparing realised against predicted performance would close the loop from selection to operation—turning the platform from a predictor into a learning system and attacking the performance gap from both ends [1, 2].

Hourly energy and marginal-carbon modelling, and refrigerant transition. Extending the lifecycle module to 8760-hour simulation and grid-marginal carbon accounting, and adding refrigerant-transition tooling (drop-in GWP comparison, TEWI/LCCP) as the F-Gas phase-down advances [17], would deepen the financial and environmental case the platform can make.

9 Conclusion

HVAC equipment selection is where the building-energy performance gap is seeded and where it can first be addressed. Yet the dominant tools—spreadsheets and disconnected vendor selectors—fragment certified performance, defeat auditability, and sever first-cost selection from lifecycle value, in direct tension with what tightening energy codes and certification programmes now demand. This paper presented the MileSoft HVAC suite as a structural answer: a unified platform organised around a single governed certified-curve engine that serves six interoperating modules—AHU, FCU, ERV, chiller, DX, and lifecycle costing—so that every output is evaluated at the project's actual conditions, bound to a certified and versioned source, costed and quoted from the same engine, and retained as a governed, auditable record.

We documented the computational method of each module against the governing standards—section-by-section AHU configuration and BOM costing; chilled-water coil rating with specific-fan-power accounting; AHRI 1060 air-to-air effectiveness; AHRI 550/590 full-load, IPLV, NPLV, and IS 16590 ISEER analysis; multi-circuit DX matching; and discounted lifecycle costing with electrical-infrastructure and carbon accounting—illustrating each with data from the live production system. We reported a production deployment at a EUROVENT-certified manufacturer that, by its own account, cut selection-cycle time by roughly 70 %, tripled per-engineer throughput, lifted conversion by 12 %, and—most tellingly—eliminated audit findings, a qualitative state-change consistent with the architecture's central claim. Finally, we abstracted from the single system to a capability reference framework: eleven criteria, grounded in the standards and in the failure modes of incumbent tooling, against which any selection platform can be evaluated, and which only a unified, governed architecture can structurally satisfy.

The central lesson generalises beyond any one product. As certification, energy, carbon, and refrigerant requirements increasingly demand that every quoted performance figure be traceable to a certified, versioned source, *auditability and integration cease to be features and become architecture*. Selection software built on that premise does not merely speed up quotation; it closes the specification-side of the performance gap, and it makes the efficient choice the defensible one. We offer the architecture, the methods, and the framework as a reproducible blueprint and an open benchmark for that goal.

Acknowledgments

The authors thank the engineering and product teams whose production system is documented here, and the deploying organisation whose reported operational data informed the case study. This work was authored by Deepak Daryani and Tanuj Daryani, MileSoft Software Technologies.

Data and Software Availability

The platform is a commercial system; the production screens reproduced in this paper are illustrative. Inquiries regarding the platform may be directed to the corresponding author. The capability reference framework (Section 6) is offered for open use.

Declaration of Competing Interest

The authors are affiliated with the developer of the platform described. This work is presented as an applied engineering and systems study; the field metrics in Section 5 are operator-reported and identified as such throughout.

References

[1] P. de Wilde, "The gap between predicted and measured energy performance of buildings: A framework for investigation," *Automation in Construction*, vol. 41, pp. 40–49, 2014.

[2] A. C. Menezes, A. Cripps, D. Bouchlaghem, and R. Buswell, "Predicted vs. actual energy performance of non-domestic buildings: Using post-occupancy evaluation data to reduce the performance gap," *Applied Energy*, vol. 97, pp. 355–364, 2012.

[3] C. van Dronkelaar, M. Dowson, E. Burman, C. Spataru, and D. Mumovic, "A review of the energy performance gap and its underlying causes in non-domestic buildings," *Frontiers in Mechanical Engineering*, vol. 1, art. 17, 2016.

[4] Air-Conditioning, Heating, and Refrigeration Institute, *ANSI/AHRI Standard 550/590 (I-P): Performance Rating of Water-Chilling and Heat Pump Water-Heating Packages Using the Vapor Compression Cycle*, AHRI, Arlington, VA, 2023.

[5] Air-Conditioning, Heating, and Refrigeration Institute, *ANSI/AHRI Standard 551/591 (SI): Performance Rating of Water-Chilling and Heat Pump Water-Heating Packages Using the Vapor Compression Cycle*, AHRI, Arlington, VA, 2023.

[6] Trane, "Part-load efficiency metrics for chillers: IPLV and NPLV," *Engineers Newsletter* (chiller part-load performance and ANSI/AHRI 550/590), Trane, La Crosse, WI.

[7] Air-Conditioning, Heating, and Refrigeration Institute, *ANSI/AHRI Standard 1060 (I-P): Performance Rating of Air-to-Air Exchangers for Energy Recovery Ventilation Equipment* (SI counterpart: Standard 1061), AHRI, Arlington, VA, 2018.

[8] American Society of Heating, Refrigerating and Air-Conditioning Engineers, *ANSI/ASHRAE Standard 84: Method of Testing Air-to-Air Heat/Energy Exchangers*, ASHRAE, Atlanta, GA, 2020.

[9] Air-Conditioning, Heating, and Refrigeration Institute, *AHRI Guideline V: Calculating the Efficiency of Energy Recovery Ventilation and Its Effect on Efficiency and Sizing of Building HVAC Systems*, AHRI, Arlington, VA, 2011.

[10] European Committee for Standardization, *EN 1886: Ventilation for Buildings—Air Handling Units—Mechanical Performance*, CEN, Brussels, 2007.

[11] European Committee for Standardization, *EN 13053: Ventilation for Buildings—Air Handling Units—Rating and Performance for Units, Components and Sections*, CEN, Brussels, 2019.

[12] Eurovent Certita Certification, *Air Handling Units Certification Programme—Certification Rules and Operational Manual (Eurovent Certified Performance)*, Paris, France (current edition).

[13] Bureau of Energy Efficiency / Bureau of Indian Standards, *IS 16590 and BEE Schedule: Indian Seasonal Energy Efficiency Ratio (ISEER) for Cooling Equipment*, Government of India, New Delhi.

[14] International Organization for Standardization, *ISO 16358-1: Air-Cooled Air Conditioners and Air-to-Air Heat Pumps—Testing and Calculating Methods for Seasonal Performance Factors—Part 1: Cooling Seasonal Performance Factor*, ISO, Geneva, 2013.

[15] American Society of Heating, Refrigerating and Air-Conditioning Engineers, *ANSI/ASHRAE/IES Standard 90.1: Energy Standard for Buildings Except Low-Rise Residential Buildings* (energy-recovery provisions, §6.5.6), ASHRAE, Atlanta, GA, 2022.

[16] European Commission, *Commission Regulation (EU) No 1253/2014 Implementing Directive 2009/125/EC with Regard to Ecodesign Requirements for Ventilation Units*, Official Journal of the European Union, 2014.

[17] European Parliament and Council, *Regulation (EU) 2024/573 on Fluorinated Greenhouse Gases (F-Gas Regulation)*, Official Journal of the European Union, 2024.

[18] American Society of Heating, Refrigerating and Air-Conditioning Engineers, *ASHRAE Handbook—Fundamentals*, ch. 1 (Psychrometrics), ASHRAE, Atlanta, GA, 2021.

[19] International Organization for Standardization, *ISO 15686-5: Buildings and Constructed Assets—Service Life Planning—Part 5: Life-Cycle Costing*, ISO, Geneva, 2017.

[20] S. K. Fuller and S. R. Petersen, *NIST Handbook 135: Life-Cycle Costing Manual for the Federal Energy Management Program*, National Institute of Standards and Technology, Gaithersburg, MD (updated edition).

A Nomenclature

Roman and Greek symbols

Symbol	Description	Symbol	Description
A, B, C, D	part-load efficiency at 100/75/50/25% load (kW/ton)	$\dot{Q}_t, \dot{Q}_s, \dot{Q}_l$	total, sensible, latent cooling duty (kW)
$c_p, c_{p,a}, c_{p,w}$	specific heat (air, water)	r	annual discount rate
C_0	first cost (equipment, install, commissioning)	SHR	sensible heat ratio, $\dot{Q}_s/\dot{Q}_t$
$C_{e,t}, C_{m,t}, C_{r,t}$	year-t energy, maintenance, replacement cost	SFP	specific fan power, $P/\dot{V}$
f_{ef}	grid carbon emission factor	t, T	temperature (°C, °F)
h	moist-air enthalpy (kJ/kg)	u	transformer utilisation factor
h (LCA)	annual operating-hours scaling	$\dot{V}$	volumetric airflow
L	connected electrical load	w	humidity ratio (g/kg)
$\dot{m}_a, \dot{m}_w, \dot{m}_{\min}$	mass-flow (air, water, minimum stream)	ΔT_w	water-side temperature rise
n_j	hours in temperature bin j	$\varepsilon_s, \varepsilon_l, \varepsilon_t$	sensible, latent, total effectiveness
P	power (fan or input)	Δp_k	pressure drop of section k (Pa)
pf	power factor	N	lifecycle analysis horizon (years)

Abbreviations

AHU—air handling unit; FCU—fan coil unit; ERV—energy recovery ventilator; DX—direct expansion; LCA—lifecycle cost analysis; LCC—lifecycle cost; BOM—bill of materials; IPLV—integrated part-load value; NPLV—non-standard part-load value; ISEER—Indian seasonal energy efficiency ratio; CSTL/CSEC—cooling seasonal total load / energy consumption; EER—energy efficiency ratio; SHR—sensible heat ratio; SFP—specific fan power; EATR—exhaust-air transfer ratio; GWP—global-warming potential; TEWI/LCCP—total equivalent warming impact / life-cycle climate performance; ECP—Eurovent Certified Performance; SSO—single sign-on; ERP—enterprise resource planning; MIS—management information system; BI—business intelligence; GCC—Gulf Cooperation Council; BEE—Bureau of Energy Efficiency.

B Worked Numerical Examples

The following examples reconstruct reported outputs from the production system using the equations of Section 2, both to document the methods concretely and to let a reader verify them independently.

B.1 Chiller IPLV from the four part-load points

The chiller rating of Fig. 8 reports part-load efficiencies (in kW/ton) of $A = 1.1018$, $B = 0.7648$, $C = 0.618$, and $D = 0.5292$ at the 100/75/50/25% load points. Substituting into the integrated part-load value of Eq. (1),

$$\begin{aligned}\text{IPLV} &= \frac{1}{\frac{0.01}{1.1018}+\frac{0.42}{0.7648}+\frac{0.45}{0.618}+\frac{0.12}{0.5292}} \\ &= \frac{1}{0.00908+0.54916+0.72816+0.22676} = \frac{1}{1.51316} = 0.661 \text{ kW/ton},\end{aligned}$$

reproducing the value 0.6609 reported by the engine (lower is better). The weights $(0.01, 0.42, 0.45, 0.12)$ place 87 % of the emphasis at or below half load, which is why selecting on full-load efficiency alone is misleading—the reference machine is roughly twice as efficient (in kW/ton) at 25 % load as at 100 %.

B.2 FCU sensible heat ratio and energy balance

The FCU rating of Fig. 5 reports a total capacity $\dot{Q}_t = 1.19\,\text{kW}$ and sensible capacity $\dot{Q}_s = 0.97\,\text{kW}$. The sensible heat ratio is therefore

$$\text{SHR} = \frac{\dot{Q}_s}{\dot{Q}_t} = \frac{0.97}{1.19} = 0.815 \approx 0.82,$$

matching the reference selection and indicating a latent duty of $\dot{Q}_l = \dot{Q}_t - \dot{Q}_s = 0.22\,\text{kW}$ (about 18 % of total), consistent with the moderate-humidity entering condition (25 °C dry-bulb, 18.67 °C wet-bulb). On the water side, with the reported temperature rise $\Delta T_w = 8.9\,°\text{C}$, Eq. (4) requires a water mass-flow $\dot{m}_w = \dot{Q}_t/(c_{p,w}\Delta T_w) \approx 1.19/(4.186 \times 8.9) = 0.032\,\text{kg/s}$, i.e. about 0.032 L/s, consistent with the low reported water velocity (0.15 m/s) and pressure drop (0.567 kPa) for a small terminal coil.

B.3 Lifecycle breakeven

The lifecycle analysis of Figs. 10–11 compares an existing unit consuming approximately 5880 kW h in year 1 against a proposed option consuming approximately 4380 kW h, an annual energy saving of $\Delta E \approx 1500\,\text{kWh}$ that grows modestly as both options derate with age. Writing ΔC_0 for the first-cost premium of the proposed option and p for the energy tariff, the discounted cumulative saving after n years is $\sum_{t=1}^{n} p\,\Delta E_t/(1+r)^t$; the breakeven year $n^\star$ is the smallest n for which this sum first exceeds ΔC_0 (equivalently, the year in which the remaining un-recovered cost in Eq. (5) reaches zero). For the reference case the engine identifies $n^\star = 6$, after which every subsequent year is net positive—so that over the fifteen-year horizon the proposed option's higher first cost is recovered with nine further years of savings, the result that the net-present-value and payback outputs summarise for the financial decision-maker.