

Clean Panels, Clear Returns: Build-Buy-Partner Choices in Robotic Solar Cleaning

Fleet capacity and lifecycle economics for photovoltaic cleaning procurement

Authored by Chennakeshav Adya

Independent Researcher

October 2026

Abstract

A photovoltaic owner considering robotic cleaning must decide which development, equipment and operating responsibilities it can fund and deliver. This paper examines an internal development programme, an equipment purchase and a contracted partner. The comparison uses a common plant boundary and remaining operating horizon, with explicit treatment of the work and liabilities retained by the owner.

Primary evidence on soiling, module maintenance, abrasion testing and operating practice informs the diligence questions. The paper separates published technical findings, manufacturer statements and proposed investment controls. It provides a reproducible fictional monthly model linking required cleaning passes, fleet capacity, manual fallback, cleaning effectiveness, deployment timing and cash expenditure. Tables and figures show scale, service-fee boundaries, delayed rollout and combined operating stresses.

The calculated ranking is conditional on the assumed service performance and costs. Equal cleaning effectiveness and adverse deployment conditions can change whether investment creates value. A proposed implementation framework connects those uncertainties to pilot acceptance, procurement exceptions, development ownership, funding commitments and replacement-service arrangements. The numerical cases represent no observed plant, supplier quotation, contractual entitlement or production forecast. Applying the method requires site-specific technical and commercial verification.

JEL Classification: G31, G32, L22, L94, O32, Q42

Keywords: photovoltaic cleaning, robotics, procurement, fleet capacity, make or buy, technology partnership, lifecycle cash flow, operating risk

1. Define the cleaning service before choosing its owner

The build, buy and partner decision should begin with a common service specification. This paper proposes comparing the cost of delivering accepted cleaning to the same eligible solar arrays over the same remaining operating period. Building means developing an internal delivery capability. Buying means purchasing an established capability, whose perimeter could be equipment or a business. Partnering means allocating delivery and economics through a negotiated arrangement. Each alternative needs an explicit boundary before its costs and benefits can be compared. A fleet purchase and a company acquisition require separate investment accounts.

IEA PVPS describes soiling as spatially heterogeneous within modules and plants, with measurement and modelling requirements that remain important for cleaning decisions. Its January 2023 summary identifies cleaning robots and their maintenance as potential capital and operating expenditures associated with mitigation. These observations establish reasons to investigate a site; they supply no standard recoverable-energy percentage or preferred ownership model [1].

The proposed service specification should identify the module population, surfaces to be cleaned, accepted residual contamination and permitted operating windows. It should also identify obstructions requiring another method. A robot completing a route has supplied an observable activity. Acceptance requires the agreed evidence that the covered surface received the specified treatment without an unresolved damage event. Additional generated electricity needs its own measurement and settlement analysis. These distinctions allow a committee to compare delivery alternatives without treating every completed pass as a paid energy benefit.

A regional operating disclosure provides a bounded reference point. Sweihan PV Power Company's September 2021 green bond framework states that Noor Abu Dhabi used cleaning robots powered by the plant without water. This is the issuer's description of that project. It provides no comparison of build, buy and partner costs, fleet reliability or transferable investment return [9].

For the proposed comparison, the owner should retain a supported incumbent cleaning case. The alternative's value consists of changes to collected energy receipts and avoidable expenditure, after implementation and continuing costs. Existing costs that remain payable should stay in both cases. Any assumed reduction in manual work needs a dated change to the labour or service commitment. Water savings need an identified avoided payment or supply requirement. The analysis should preserve separate quantities for cleaning output, energy effect and owner cash. This creates a common unit of comparison while allowing each transaction structure to have its own development, mobilisation and financing schedule.

2. Establish compatibility with the installed module population

The first technical boundary is the actual combination of robot, module, mounting arrangement and operating method. In the proposed diligence process, the owner should map these combinations before estimating usable fleet capacity. Record module variants and coatings, support positions, row geometry, gaps, tracker movement and permitted loads. Ask the responsible engineers which combinations need further evidence. A successful demonstration on one array should enter the record with its specific configuration, operating settings and inspection results. Additional configurations remain separate approval questions.

LONGi's April 2025 maintenance manual allows automatic cleaning equipment subject to its operating requirements and protection of the modules. Its arid-and-rainless-area guidance refers to automatic equipment approved by LONGi. The manual also states that it supplies no warranty terms. These are manufacturer instructions for LONGi modules; they establish no approval of a particular robot or replacement for the applicable warranty documentation [4].

Trina Solar's Vertex manual, version J of November 2023, applies to its listed modules. Its dry-cleaning-robot instructions specify soft plastic brush material, avoidance of scratches and attention to robot weight. It states that damage and power attenuation resulting from improper

robot use fall outside its warranty. The manual's conditions require review for the actual module population and operation [5].

The proposed compatibility record should connect each acceptance conclusion to the exact equipment version. A change in wheel material, brush assembly, operating pressure or control logic should prompt a documented assessment of whether earlier evidence remains applicable. The record should identify who can authorise those changes and what the supplier must disclose. A commercial warranty offered by the robot supplier needs separate scrutiny of covered components, exclusions, claim procedure and financial capacity to respond. A module supplier's acceptance letter should be read for its stated scope and conditions.

This boundary affects the ownership decision directly. An internal build may require engineering work across several module types. An equipment purchase may require adaptations to existing rows and parking positions. A partner may propose a restricted initial deployment on eligible blocks. The comparison should price those specific scopes and retain an alternative for excluded arrays. Where compatibility remains unresolved, the associated output and expenditure belong in a conditional case. A fleet-wide value calculation should not assume that the entire installed plant can receive the same treatment on the same date.

3. Price cleaning effectiveness alongside surface durability

Cleaning acceptance should address immediate removal and the condition of the treated surface. The proposed evidence plan should specify the contamination, brush and contact conditions used in testing, together with the measurements taken before and after treatment. It should distinguish a clean-looking surface from demonstrated optical and electrical performance. Repeated operation also requires a stated examination for scratches, coating change, frame damage and other relevant observations. The acceptable scope and duration of this examination should be agreed with the responsible specialists.

Bomber and colleagues' 2024 NREL presentation describes a five-year glass-coupon field study including Dubai and Kuwait City. It compared dry brushing, water spraying, wet sponge and squeegee cleaning with uncleaned specimens, using optical and surface examinations. Its reported experimental population consisted of glass coupons and coatings. The study provides evidence about those specimens and methods; it supplies no commercial fleet availability or site-specific robotic cleaning return [2].

The IEC catalogue lists IEC 62788-7-3:2022 with its 2024 amendment consolidated as edition 1.1, published on 31 July 2024. Its public scope covers methods for evaluating abrasion of photovoltaic surfaces and coatings, using representative specimens or modules. The complete paid standard was not examined here. The cited scope supports identification of an abrasion-testing framework, without establishing a robot's certification or field lifetime [3].

For a proposed investment, the testing dossier should explain how laboratory exposure relates to the actual cleaning method. Record the brush configuration, local contamination used, operating settings, repeat count and examination method. A specialist should explain what the test can establish and which field conditions remain outside it. The financial case should carry separate consequences for immediate cleaning underperformance and possible damage. Combining them into one recovery assumption would conceal the different evidence and expenditure needed to resolve each issue.

The owner can then compare the costs of obtaining further evidence under each commercial route. Building may require prototypes and redesign before wider deployment. Buying may provide existing test records whose applicability still needs assessment. A partner may accept a staged field programme with agreed responsibilities for inspection and exceptions. Those arrangements are proposals for negotiation. Their financial value depends on the actual evidence, terms and counterparty support. The comparison should retain investigation and replacement cash where recovery remains uncertain, with a separate record of any claim that might reimburse it.

4. Measure the incremental service and collected benefit

The proposed measurement plan should be agreed before a pilot begins. Identify treatment and comparison areas, sensor locations, observation periods and the treatment of outages, weather changes and missing data. Preserve the operating information available when each cleaning decision was made. Subsequent observations should remain identifiable as subsequent information. The objective is to estimate the incremental effect of the proposed service under a stated comparison, with an uncertainty range that the investment committee can understand.

IEA PVPS's March 2014 analytical-monitoring report describes monitoring equipment and analysis of plant performance, including periodic regression methods for examining temperature and other influences. Its executive summary addresses losses requiring analytical assessment. This supports attention to measurement configuration and explanatory variables. It does not validate a particular robotic-cleaning trial or prescribe a recoverable-energy result [8].

A more recent local example concerns measurement equipment. DEWA's December 2025 announcement describes prototype testing of its Autonomous Soiling Detector at the authority's head-office parking area. DEWA says the device compares actual production with expected production under defined conditions. This is a bounded prototype disclosure. No numerical result from that announcement is adopted as a cleaning-fleet performance assumption [11].

The proposed pilot should retain a complete event record, including aborted routes, interventions, manual assistance and damage inspections. An acceptance rule should state whether a partial route earns any service credit and how repeated attempts are counted. The evaluator should record when a robot was available but the array was inaccessible. These observations allow the committee to understand which constraint limited delivery and which party could address it. A vendor's chosen successful demonstration should remain distinguishable from the full agreed evaluation period.

Energy and cash require additional reconciliation. The owner should identify whether extra module output could reach the export meter and be paid under the actual arrangements. Record any coincident outage, export limit or settlement exclusion before assigning receipts. Count cleaning electricity, technician time and replacement consumables in the intervention account. The proposed evaluation should report cleaning acceptance, incremental exported energy and incremental owner cash separately. A pilot with unresolved attribution can still inform further testing, while its financial conclusion remains conditional on the missing evidence.

5. Convert fleet specifications into a dated service calendar

A proposed fleet plan should work from the plant layout and eligible operating windows. For each block, identify the route, transfer method, charging or parking requirement and access

restrictions. Record how a disabled unit is retrieved and how cleaning continues during repair. The plan should show the people and equipment needed to move, inspect and service the robots. A nominal operating speed supplies one input; the accepted work completed across the calendar is the quantity relevant to an operating commitment.

IEA PVPS's October 2022 climate-specific O&M guidance identifies travel and preparation requirements for hot-and-dry sites and recommends clearly specified responsibilities and negotiated performance indicators. Its summary also discusses the need to account for regional and stakeholder differences. This is guidance for configuring an O&M programme. Its historical cost ranges and statements about technology availability are not used as current robotic-cleaning benchmarks [6].

The proposed fleet calendar should deduct the time allocated to mobilisation, transfers, inspections, repair and other unavailable periods. Shared support resources need an explicit allocation when several sites request service together. A maintenance spare should be identified separately from a machine assumed to produce ordinary cleaning output. The calendar should retain the reason for each deduction so a reviewer can test a different staffing or spares proposal without changing unrelated assumptions.

For investment comparison, useful utilisation should be measured against the compatible and authorised workload. A growing list of potential plants supplies no completed site qualification or signed service demand. Internal deployment should identify the owner's actual plant schedule. An acquired supplier requires diligence on customer commitments and delivery resources. A partner arrangement requires agreement on which customer receives priority during constrained periods. Those records determine whether the fleet can perform the work counted in the cash case.

The proposed downside analysis should include a concentrated service interruption, such as a delayed replacement component or an inaccessible transfer route. It should state the assumed repair duration and available fallback method. Any resulting lost revenue, extra labour and contract payment should be charged once to the appropriate account. The stress is an illustrative test until asset evidence supports its inputs. This approach makes the fleet plan auditable and helps identify whether additional equipment, support capacity or a narrower deployment could change the decision.

6. Assess the capability that an internal build must deliver

An internal build should have a defined operational end state. The proposed assessment includes the robot design, manufacturing or assembly process, control software, field support and responsibility for continuing changes. Identify which elements already exist within the operator and which require recruitment, procurement or development. Record the dependencies between those activities and the date on which accepted service could begin. Development expenditure should remain visible even where an early prototype performs a cleaning pass.

The United States Department of Energy's Adoption Readiness Level framework complements technical readiness with commercial adoption considerations. Its stated risk areas cover value proposition, market acceptance, resource maturity and licence to operate, including matters such as supply chain and workforce. It is a qualitative framework. It provides no readiness score, product endorsement or financing conclusion for the cleaning alternatives considered here [12].

Applied as an author-proposed assessment, that distinction requires separate evidence for technical functioning and repeatable delivery. The internal team should show how it would obtain critical components, verify incoming quality, control design changes and support installed units. A prototype dependency on one engineer or unavailable component should have a stated operating consequence. The owner should assess whether it wishes to retain those responsibilities through the full service period and what resources that commitment requires.

The proposed development budget should distinguish knowledge acquisition from equipment that can supply service. Testing, failed iterations and specialist engineering can be legitimate expenditures whose resulting value depends on subsequent deployment. Their cash dates and remaining commitments should be recorded. A decision to stop development needs a corresponding exit account for unused materials, support obligations and any recoverable assets. The comparison should preserve those costs and conditional recoveries without assuming a sale market.

An internal build may also create reusable knowledge about the operator's arrays and cleaning needs. The committee should describe that potential benefit through identifiable uses and expenditure avoided in a later decision. An unsupported strategic value should remain outside the numerical case. The proposed build decision should therefore name the evidence required for the next commitment, the maximum expenditure authorised at that stage and the operational alternative available if development is delayed. This keeps development risk connected to a dated capital decision.

7. Diligence the capability included in a purchase

A proposed purchase needs a precise asset and service perimeter. For equipment, identify hardware, software access, commissioning, training, spares and continuing support. For a business acquisition, identify which customer agreements, intellectual property, staff, liabilities and operating resources would actually transfer. The committee should reconcile the purchase price with the expenditure still required to deliver accepted service. A transaction label supplies no evidence that those resources are included or available when needed.

TrinaTracker announced the launch of its AURORA cleaning robot on 24 September 2026. The announcement describes route planning, integration with plant information and self-powered cyclic operation as product features. These are manufacturer assertions in a launch release. They establish the disclosed product offering and its claimed characteristics, with no independently verified fleet performance, installed-site compatibility or investment outcome adopted here [10].

For the proposed procurement process, each material assertion should be connected to an inspectable record. Ask for the test configuration, software version, conditions and complete operating period behind a claimed capability. Establish which performance statements enter the contract and which remain descriptive material. The owner should cost any site adaptation and retain a separate decision on the actual acceptance criteria. Evidence from a different module or tracker combination should carry its stated applicability limits.

NREL's 2018 O&M guidance recommends maintaining current service agreements, warranties, equipment information and maintenance records. Its contractor assessment also addresses evidence of financial solvency and a service provider's track record. These are diligence

considerations in the cited guidance. They establish no credit assessment or enforceable claim against a proposed robot supplier [7].

The proposed purchase review should connect technical and commercial exceptions to the investment account. Missing support documentation can require additional diligence or a narrower initial commitment. A supplier's obligation to repair equipment should be examined alongside the cash and time needed for the owner to continue operating during a dispute. In a business acquisition, historical operating records should be tested against the customers and resources included in the transaction. The committee can then compare the delivered capability with the internal-build alternative using a common service scope and separate implementation cash.

8. Allocate partner responsibilities and preserve the value comparison

The proposed partner route should identify the capability each party contributes and the evidence supporting its availability. One party may provide hardware while another supplies site access, technicians or customer relationships. The arrangement needs a defined service perimeter, timetable and decision process. A shared venture, subcontract and service agreement create different financial accounts; the investment comparison should state which structure it assumes and retain the documents needed to establish actual rights.

NREL's 2018 guidance discusses the use of a backup O&M provider with products and training to take over if the incumbent fails to perform. The guidance describes a possible service-continuity arrangement. Its inclusion supplies no evidence that a suitably qualified backup is available for a particular robotic-cleaning fleet [7].

For the author-proposed comparison, partner economics should follow the obligations retained by the owner. Identify payments for accepted service, minimum commitments, mobilisation, exceptions and termination. Show the cash needed if service is interrupted while the parties resolve responsibility. A proposed right to obtain software, spares or another provider needs review for its practical operation and actual enforceability. Any possible reimbursement should remain separate from the gross expenditure that the owner may have to fund first.

The three routes can then be compared against the same eligible workload and evidence-based service outcome. Their differences include the timing of deployment, capital committed before acceptance, support cost, retained responsibilities and exit consequences. A higher assumed cleaning effectiveness should have its own technical evidence requirement. A lower payment should have a defined scope and commercial basis. The proposed framework assigns no universal premium to ownership, outsourcing or partnership.

The final decision record should identify the selected scope, conditions still unresolved and changes that would require a new investment review. Preserve the accepted comparison, implementation budget and fallback service. As operation begins, reconcile actual accepted cleaning and paid cash with the assumptions that supported approval. Record any difference together with its evidence and proposed response. The resulting choice remains specific to the operator's assets, resources and negotiated arrangements, with no claim of a generally optimal route or an observed regional return.

9. Define the economic boundary before comparing bids

The proposed comparison begins with an existing photovoltaic plant and a fixed remaining operating period. It evaluates the incremental cash consequences of changing its cleaning arrangements. Plant acquisition value, ordinary project debt and electricity revenue common to every option sit outside the cleaning investment calculation. This boundary allows the committee to identify the cash flows that a procurement decision can actually change. A broader acquisition or refinancing model can incorporate those changes after the operating analysis has been accepted.

Each alternative needs a complete description of the service being purchased. The description should specify eligible rows, required cleaning passes, acceptable operating windows, inspection obligations and the response to work that the robots cannot complete. The owner should also identify excluded surfaces and operations. An offer covering accessible module rows can leave perimeter modules, transition points and exceptional deposits with a separate team. The retained work belongs in the proposed operating budget and staffing plan. Its treatment should remain visible when bids are ranked.

Development expenditure, installed fleet cost and an annual service fee occur at different times. We therefore propose recording commitments against the same decision date and discounting the incremental cash flow through the same end date. The development route carries its initial expenditure even if the pilot fails and subsequent commercial deployment is cancelled. The equipment purchase includes installation and commissioning before service begins. A partnership includes mobilisation, agreed fees, retained owner supervision and exit obligations. Tax, financing and currency assumptions should be added only when their treatment has been established for the actual transaction.

The value of additional generation depends on whether the additional electricity can be sold and collected. For an actual site, the model should reconcile the proposed operating benefit to the relevant contractual payment rules and dispatch constraints. A flat illustrative electricity value provides a transparent calculation exercise. It cannot establish the revenue available under an unidentified power-purchase agreement. A reader applying the method should replace that input with the appropriate incremental settlement calculation and retain the underlying commercial evidence.

Working capital requires a separate funding schedule. A service may be paid in advance, equipment may require a deposit and a repair claim may remain unpaid while replacement parts are purchased. The present-value comparison should be accompanied by the largest cumulative cash commitment before benefits are received. That calculation lets the investment committee assess whether the preferred route can be funded through commissioning and an adverse operating event. The numerical illustration in this paper excludes project financing and tax and identifies its other simplifications alongside the results.

10. Allocate development costs to an approved fleet

An internal development programme should have a named beneficiary fleet. Its approved scope might cover one plant, several contracted plants or a separately funded technology business. These boundaries create different sources of recovery for design expenditure. An operator applying the method should allocate development cost only to sites with an approved, feasible

deployment plan. Potential sales to unrelated owners can be examined in a separate business case after commercial access, customer requirements and delivery costs have been investigated.

The proposed approval sequence separates prototype expenditure, a controlled field pilot and commercial fleet procurement. Before committing to the next stage, the committee should receive the agreed test results, the outstanding failure list and an updated estimate of the remaining cash requirement. A decision to stop preserves the expenditure already incurred and avoids only the expenditure that remains cancellable. Showing both amounts makes the abandonment decision intelligible to the board and avoids presenting a cancelled programme as a costless experiment.

The development team also needs a costed support plan. The suggested scope includes design ownership, source files, production documentation, component procurement, field repair and change control. An attractive prototype leaves unanswered questions if the site team cannot obtain a replacement controller, reproduce an approved component or diagnose a failure after the original engineers leave. These are proposed diligence questions for an operating capability. They provide no evidence that a particular supplier or internal team currently meets the requirements.

Portfolio expansion changes both expenditure and operational obligations. A robot fleet purchased for an additional plant can introduce new row geometries, access restrictions, module surfaces or software integrations. The proposed scale test therefore requires both a financial calculation and a compatibility assessment for each added site. Shared engineering expenditure can be spread over a larger approved portfolio when the additional installations use that engineering. Site-specific adaptation and support costs remain attached to the sites that create them.

An equipment purchase also transfers a defined package of capabilities. The owner should identify which tasks remain dependent on the supplier and which can be performed with internal staff or an alternative maintainer. A quoted hardware price should be reconciled to installation, software access, spares, operator training and continuing support. The development and purchase cases can then be compared on the service they deliver and the obligations the owner retains. The model below tests portfolio size and deployment timing using explicitly fictional inputs.

11. Translate pilot measurements into a cash claim

We propose defining the pilot's decision question before equipment is installed. A useful question is whether a particular cleaning arrangement can complete the specified work within the allowed calendar, sustain the required surface condition and produce an economically usable improvement. The pilot protocol should identify comparison blocks, measurement intervals, operating exclusions and the tests that determine whether an installation can proceed. Each proposed threshold should be agreed by the parties responsible for engineering acceptance and the investment decision.

The measurement record should connect a completed cleaning task to its location, start and finish time, equipment identity and subsequent inspection. The production record should separately capture the electricity measurement and relevant operating conditions. A measured production change requires assessment of the other changes occurring during the comparison. A cleaning log alone provides no measure of incremental collected revenue. The proposed evidence chain therefore keeps task completion, surface condition, electricity and cash settlement as separately reconcilable records.

Comparison blocks should be selected before the trial results are known. The proposed design records the reasons for their inclusion, the planned cleaning treatment and any change to that treatment. If a block is excluded after a fault, rain event or instrumentation failure, the exclusion and its effect on the result should remain visible. This approach gives the committee a way to distinguish a favourable result under selected conditions from a service that can be delivered across the intended operating scope.

An operating event also needs a complete labour record. A robot may finish its route after a technician resets it, carries it across a gap or replaces a consumable. The proposed pilot records those interventions, travel time, parts and supervision against the same work order. The commercial model should then apply the expected staffing arrangement and cost assumptions openly. A short trial with unusually intensive supplier support should be described with that qualification when its results are used to support a longer contract.

Uncertainty should travel into the procurement decision. The committee can examine a range for completed work, recovered generation and intervention cost and identify the combinations under which the recommendation changes. This proposed test does not require an unsupported probability distribution. It requires clear conditional results and a practical way to obtain the missing evidence. If a modest change in a material input reverses the decision, the committee should identify whether a further trial or a contractual provision can resolve the exposure before committing the full amount.

12. Price the continuity obligations that remain with the owner

The proposed continuity assessment identifies the actions required when the equipment or its supplier becomes unavailable. It should distinguish an individual hardware fault, a fleet-wide software restriction and the failure of the contracted service provider. For each event, the owner should specify the work that stops, the temporary operating arrangement, the parts or permissions needed for recovery and the cash payable before normal service resumes. These are scenario design choices; the paper assigns no observed frequency to the events.

Access rights need an operational test. A contract may refer to data delivery, replacement support or source-code arrangements, but the committee should identify the exact files, credentials, documentation and physical items that the replacement operator would need. Their availability should be checked through the agreed acceptance process. Legal advisers should assess the proposed rights, restrictions and enforceability for the actual contract. The commercial assessment can then cost the recovery plan that those reviewed rights permit.

A service fee can be tested against a conditional ceiling. The ceiling is the fee at which the specified partnership produces the same incremental present value as the chosen feasible alternative, holding the stated service assumptions constant. It changes when mobilisation, retained owner costs, operating failure or exit costs change. The model's ceiling therefore represents a negotiating calculation under its fictional assumptions. An actual willingness to pay also depends on the owner's funding capacity, risk policy and confidence in the promised service.

Purchasing the supplier itself would require a separate transaction analysis. Its enterprise value would need support from the acquired business's revenue, costs, obligations and assets, together with a credible integration plan. Savings at the owner's plants would form one explicitly identified contribution. This paper's equipment-purchase case does not value a robotics company

or assign a control premium. The distinction gives the reader a clear scope for using its results in an acquisition discussion without treating fleet procurement expenditure as a company valuation.

The final recommendation should specify an approved scope and the conditions for changing it. A committee might approve a limited equipment purchase, a cancellable development stage or a service contract with a defined recovery arrangement. Each proposal should identify its remaining evidence requirement, responsible decision-maker and cash commitment. After deployment, the review should compare completed work, interventions and collected benefits against the approved assumptions. A revised operating decision can then be traced to the evidence that changed it.

13. Define a common service before pricing procurement

The worked case concerns an existing, fictional 200 MWdc photovoltaic plant with ten years remaining in its operating contract. Its owner considers buying a turnkey robot fleet, developing a proprietary fleet or retaining a contracted cleaning partner. A fully resourced manual programme is the common counterfactual. All amounts, equipment quantities, performance relationships and contract terms below are author-created assumptions. They describe a transparent decision experiment and provide no evidence about an actual supplier, site, tariff or equipment approval. The comparison values only changes attributable to cleaning procurement. Existing land, generation equipment and electrical access remain common assets with their existing obligations.

The synthetic operating year requires two cleaning passes in quieter months, three in intermediate months and four in the busiest months. This creates 7,200 MW-passes annually. One MW-pass means completing the stipulated cleaning treatment across one MWdc of the plant; it is a workload unit, with no claim of uniform physical area or terrain across actual sites. Each month permits twelve service nights. Both owned fleets contain eighty active robots and eight spare units. Every active unit is assigned 0.9 MW per service night. The stipulated service factors are 95% for purchased equipment and 92% for proprietary equipment. Spare units incur purchase and renewal expenditure but do not provide an additional simultaneous operating shift.

These assumptions produce 820.8 MW-passes of purchased-fleet monthly capacity and 794.88 for the proprietary fleet. Peak demand is 800. Purchased robots therefore complete the full specified workload. The proprietary fleet completes 99.36% in the four peak months and requires manual completion of the remaining 20.48 MW-passes annually. A partner supplies 98% of required work, leaving 144 MW-passes for the owner's manual fallback. These differences are exposed before financial ranking. A comparison based only on equipment counts would overlook the binding proprietary-fleet limit and the retained labour requirement under the service contract.

The programme assumes that manual personnel, access and equipment are available whenever called in the base case. That assumption is economically consequential and is relaxed in the partner-disruption tests. Monthly capacity cannot demonstrate that cleaning occurs before a particular dusty interval, that adjacent rows are accessible or that one spare covers simultaneous failures. Actual deployment needs an executable route and shift schedule, compatible module and mounting systems, and evidence supporting productive hours. The calculated capacity thresholds are screening quantities for that work. They do not certify availability, reliability or a completed engineering design.

The numerical illustration uses fictional inputs and proposed decision records. Its plant, fleet throughput, cleaning effectiveness, costs, electricity value and contract assumptions are not observations or forecasts. The finite ten-year comparison excludes project debt, tax and a terminal extension. Manual fallback is an explicit capacity and cost assumption. Actual deployment requires site-specific engineering, warranty, settlement and contractual verification. Calculated table values use up to three decimal places with round-half-up; years and months remain integers. Floating-point noise is removed at twelve decimal places before display rounding; the model retains its underlying precision.

Table 1. Comparable fictional procurement commitments

Choice	First service year	Active robots	Spares	Initial fleet USDm	Development USDm
Manual	1	0	0	0.000	0.000
Buy	1	80	8	2.714	0.000
Build	2	80	8	1.834	1.600
Partner	1	0	0	0.000	0.000


Existing plant, land and electrical connection are common. Manual fallback is expressly resourced in the base case. No actual delivery or compatibility is established.

Table 2. Monthly work demand and service coverage

Month	Required MW-passes	Buy robot %	Build robot %	Partner robot %
1	400.000	100.000	100.000	98.000
2	400.000	100.000	100.000	98.000
3	600.000	100.000	100.000	98.000
4	600.000	100.000	100.000	98.000
5	800.000	100.000	99.360	98.000
6	800.000	100.000	99.360	98.000
7	800.000	100.000	99.360	98.000
8	800.000	100.000	99.360	98.000
9	600.000	100.000	100.000	98.000
10	600.000	100.000	100.000	98.000
11	400.000	100.000	100.000	98.000
12	400.000	100.000	100.000	98.000

Buy/partner first operating year and build first operating year shown. Twelve service nights each month; coverage is a synthetic workload cap, not a reliability or access certificate.

Figure 1. Seasonal workload and completed robot service under the fictional operating window



All inputs fictional. Twelve service nights per month, stated active fleets and service factors. Build is shown in its first operating year; remaining work uses the specified manual fallback.

14. Reconcile completed work with electricity and operating cost

First-year clean potential is stipulated at 400,000 MWh, allocated through twelve monthly weights. A separate 98.5% ordinary photovoltaic availability factor applies before cleaning losses. Output declines by 0.5% annually across every alternative. The monthly unclean soiling-loss fraction ranges from 4% to 12%. Manual cleaning removes 60% of that stipulated loss, while completed robotic work removes 90%. These are hypothetical comparative recovery factors. They are deliberately tested at equal recovery later because a procurement decision should not inherit a performance advantage without supporting measurements.

For each month, robot and manual completed-work shares determine the recovery fraction. Unserved work receives no recovery. Residual soiling is deducted from ordinary-available energy, followed by the weighted cleaning interruption: 0.1% for robotic work and 0.2% for manual work. This sequential calculation assigns each lost MWh once. It does not represent a daily dirt stock, storm chronology or hourly dispatch simulation. In particular, completing half the monthly workload is assumed to deliver half the stipulated recovery contribution. A real site could respond differently if cleaning is concentrated in the most productive or most affected rows.

The manual first-year export is 380,896.600 MWh. The purchased fleet exports 390,524.065 MWh and the partner case 390,331.334 MWh. Proprietary service begins in year two, when it exports 388,536.440 MWh. Its comparison uses the year-two manual baseline after the same

degradation; comparing that quantity directly with year-one manual export would mix operating performance with calendar timing. The model retains monthly baseline and alternative energy so that this distinction can be reconstructed.

Manual treatment costs USD 80 per MW-pass. Owned robots incur USD 8 per completed MW-pass, annual fixed support of USD 120,000 for purchase or USD 90,000 for build, and maintenance equal to 2% of fleet acquisition cost. Cleaning electricity is purchased separately at USD 100/MWh, using 0.01 MWh per manual MW-pass and 0.02 per robotic MW-pass. This expenditure is charged once to cash and is not also deducted from exported solar electricity. The manual programme costs USD 0.583 million annually. Purchased-fleet operating expenditure is USD 0.241 million and proprietary service costs USD 0.195 million in its first operating year. Partner expenditure includes its fee, owner support, purchased cleaning power and remaining manual work, totalling USD 0.669 million annually in the base case.

Table 3. First operating-year energy and cash service

Choice	Year	Export MWh	Manual MW-passes	Owner operating USDm
Manual	1	380,896.600	7,200.000	0.583
Buy	1	390,524.065	0.000	0.241
Build	2	388,536.440	20.480	0.195
Partner	1	390,331.334	144.000	0.669

Build starts later and its gross energy has one year of the same degradation. Extra export and cost are compared against that same calendar-year manual counterfactual.

15. Price commitments and retained cash on one calendar

Purchased equipment costs USD 28,000 per unit, including the eight spare units, with USD 0.250 million of site capital. Its USD 2.714 million commitment occurs at the decision date and service starts in year one. Proprietary development consumes USD 1.200 million immediately. The pilot costs USD 0.400 million in year one, followed by an assumed successful commercial gate and USD 1.834 million fleet and site commitment at that year-end. Proprietary units cost USD 18,000 each. The lower unit cost is a fictional commercial assumption, with no inference about achievable manufacturing margins, certification or intellectual-property value.

Owned fleets incur renewal expenditure equal to 65% of their equipment purchase cost after five operating years: year five for purchase and year six for build. Site capital is excluded from the renewal base. The assumed renewal maintains the stated service capability without increasing cleaning recovery. Any extra replacement interruption is unmodelled beyond the service and cleaning-interruption factors already specified. Year-ten closure costs USD 0.150 million for owned alternatives. There is no equipment resale, contract extension or terminal recovery. The partner instead incurs USD 0.050 million of final handback expenditure.

Every operating receipt and payment is assumed to settle at year-end. Incremental energy earns USD 40/MWh, and the discount rate is 8%. Existing plant revenues and costs cancel through comparison with manual operation on the same calendar. Incremental NPV is USD 0.955 million for purchase, USD 0.428 million for build and USD 1.884 million for partnership. These values already include development, renewal and closure. Adding a second deduction for capital recovery would double count expenditure. They also provide no valuation of the wider plant or evidence that an investor can acquire it at those amounts.

The retained incremental-cash ledger records a peak deficit of USD 2.714 million for purchase and USD 3.434 million for build. Partner annual incremental cash is positive throughout the base case. This annual ledger does not establish zero working-capital need: deposits, monthly invoice timing and payment disputes are omitted. No borrowing, tax, inflation, sponsor commitment or distribution restriction is modelled. Positive NPV expresses conditional operating value under the stipulated cash timing. Actual funding capacity and investor distributions require separate financing and contract accounts.

Table 4. Incremental annual cash over the fixed remaining term

Year	Buy USDm	Build USDm	Partner USDm
0	-2.714	-1.200	0.000
1	0.727	-2.234	0.291
2	0.725	0.770	0.290
3	0.723	0.768	0.288
4	0.721	0.766	0.286
5	-0.882	0.764	0.284
6	0.717	-0.267	0.282
7	0.716	0.760	0.280
8	0.714	0.759	0.278
9	0.712	0.757	0.277
10	0.560	0.605	0.225


Real USD million, t0 capital then end-year cash. Manual baseline has zero incremental cash. Development, renewal and final closure are included; no residual sale or extension.

Table 5. Decision-date value and conditional break-even tests

Choice	NPV USDm	Extra t0 cost USDm	Break-even USD/MWh	Peak deficit USDm
Buy	0.955	0.955	24.922	2.714
Build	0.428	0.428	32.112	3.434
Partner	1.884	1.884	9.659	0.000

The energy price column is the break-even minimum, not an offered tariff. Negative extra-cost capacity means a shortfall. Peak deficit retains incremental cash only and is not a whole-plant funding requirement.

Figure 2. Retained incremental cash after development, fleet commitment and renewal



Undiscounted incremental cash relative to manual operation. No interim distributions, financing, intrayear liquidity or whole-plant funding claim. Fixed year-ten closure with no residual sale.

16. Test scale and the value of an accepted service price

The partner contract contains a fixed USD 0.180 million annual fee and USD 60 per completed robotic MW-pass. Owner oversight adds USD 0.040 million annually. Variable fees follow actual completed work; the fixed fee remains payable in the disruption scenario. No performance compensation, security recovery or termination damages are assumed. This explicit payment construction connects the physical service quantity to owner cash and gives a basis for testing a bid with an otherwise comparable scope.

At the base workload, the variable fee reaches USD 99.794 per MW-pass before the partner's incremental NPV falls to zero, holding the fixed fee and every other assumption unchanged. That is a ceiling relative to continued manual cleaning. A lower threshold governs the choice between partnership and equipment purchase: the partner loses its calculated advantage over purchase when its variable fee reaches approximately USD 79.614 per MW-pass. A buyer considering both alternatives should use the relevant competing option, with its own risks and expenditure, when assessing negotiation room. Neither boundary is an actual supplier price or a recommended contractual rate.

Scale tests preserve the same service density and unit productivity at 50, 100, 200, 400 and 800 MWdc. All three automation alternatives have negative incremental NPV at 50 MWdc. At 100 MWdc, partnership is positive by USD 0.192 million while purchase remains slightly negative and proprietary build remains substantially negative. At 400 MWdc, build reaches USD 3.248 million and exceeds purchase at USD 3.036 million, although partnership remains higher at USD 5.268 million. Fixed development expenditure is spread over more service volume, which explains the change between the owned alternatives.


These scale results assume an accessible operating programme with unchanged route productivity. A real multi-site fleet can require transport, local spares, travel time, communications integration and separate contractual interfaces. The model includes none of those additions. It therefore supports questions about minimum service volume and fixed-cost exposure, with no evidence of a universal commercial scale threshold. Integer fleet requirements also matter. Under the stipulated peak workload, purchased service needs at least 78 active units and proprietary service at least 81 to avoid manual completion. The base fleets remain at eighty; the model does not silently add the extra proprietary robot or claim that eighty is an optimal purchase quantity.

Table 6. Portfolio scale under the same service density

MWdc	Buy NPV USDm	Build NPV USDm	Partner NPV USDm	Owned active/spare units
50.000	-0.605	-1.687	-0.654	20/2
100.000	-0.085	-0.982	0.192	40/4
200.000	0.955	0.428	1.884	80/8
400.000	3.036	3.248	5.268	160/16
800.000	7.196	8.887	12.035	320/32

This assumes one accessible operating programme with unchanged unit productivity. Cross-site travel, controls and deployment overhead are unmodelled; this is not an observed fleet scale curve.

Figure 3. Scale sensitivity at unchanged route productivity and service density



Decision-date real USD million, 8% discount. Synthetic same-programme scale cases exclude transport between sites, extra deployment overhead and observed vendor economies.

17. Preserve the rollout gate and specify failure recovery

The proprietary development decision is separated into dated commitments. In the stop case, the owner pays development at year zero and the pilot at year one, declines commercial fleet expenditure, and retains manual cleaning for the remaining term. The resulting incremental NPV is negative USD 1.570 million. That amount prices the exercised stop path under known fictional

cash flows. It does not assign a success probability, calculate a probability-weighted option value or establish that stopping after the pilot is the correct decision in every information state.

Delays retain the common year-ten end date. Fleet commitment moves to the year before service and the associated renewal follows the delayed operating age. Development and pilot costs remain at years zero and one. Manual work continues before rollout, so the model changes actual operating benefits and expenditure without adding a separate foregone-benefit charge. A one-year additional delay makes the build case slightly negative at USD 0.058 million below manual. A two-year delay produces USD 0.203 million for purchase, negative USD 0.507 million for build and USD 1.366 million for partnership. These tests expose how much of a finite remaining operating term each commitment consumes.

The partner-disruption scenario provides no robot service in year four, retains the fixed fee, and incurs USD 0.150 million of transition cash. Recovery in year five is stipulated; there is no assumed legal right or demonstrated replacement contractor. Full manual fallback leaves incremental NPV at USD 1.402 million. Limiting fallback to one quarter of missing work produces 5,400 unserved MW-passes, year-four export of 362,107.063 MWh and NPV of USD 1.328 million. With no manual fallback, 7,200 MW-passes remain unserved and NPV is USD 1.303 million.

These results require careful interpretation. Less fallback lowers electricity recovery and also avoids manual spending. At the assumed tariff, those opposing effects keep the additional NPV loss relatively small. That calculation does not establish that leaving equipment dirty is technically acceptable or contractually permitted. Persistent damage, access restrictions, penalties and emergency replacement are excluded. A transaction team should add supported consequences to the ledger and test prolonged disruption before accepting the apparent financial resilience. Contract lock-in is represented narrowly by the retained fixed fee and transition expenditure, with no claim that those terms exhaust actual switching costs.

Table 7. Delayed service with fixed year-ten closure

Delay years	Buy NPV USDm	Build NPV USDm	Partner NPV USDm
0	0.955	0.428	1.884
1	0.564	-0.058	1.614
2	0.203	-0.507	1.366
3	-0.129	-0.921	1.138


Fleet commitment shifts to the year before service. Build development and pilot cash remain at years zero and one. Manual service continues before rollout. No separate lost-benefit penalty is charged.

Table 8. Partner disruption in year four

Fallback share	Unserved MW-passes	Year 4 export MWh	NPV USDm	Variable fee ceiling USD/MW-pass
Full manual fallback	0.000	375,211.671	1.402	93.256
Quarter manual fallback	5,400.000	362,107.063	1.328	91.506
No manual fallback	7,200.000	357,729.744	1.303	90.916
No disruption	0.000	384,505.590	1.884	99.794

No robot service in year four; stipulated year-five recovery, retained fixed fee and USD 0.150 million transition cash. Variable fees follow completed robot work. No damages or actual termination entitlement is assumed.

Figure 4. Variable service-fee capacity after a stipulated year-four disruption



Fixed annual fee remains payable; variable fees follow completed work. Year-five recovery is assumed, and no damages or actual counterparty recovery is credited. Fee ceilings compare with manual continuation.

18. Turn sensitivity into evidence requirements

The comparison is particularly sensitive to the assumed recovery advantage. Equalising robot and manual recovery at 60% makes every automation alternative negative: purchase loses USD 1.479 million, build USD 1.656 million and partnership USD 0.501 million. The small difference in cleaning interruption remains. These results identify measurement of incremental recovery as a decision-critical workstream. The base ranking cannot substitute for a controlled comparison under representative access, soiling and production conditions.

Other tests isolate commercial and execution exposure. At USD 25/MWh, purchase is approximately break-even at USD 0.005 million, build is negative USD 0.386 million and partnership remains positive USD 0.953 million. Increasing equipment and site capital by 25% makes both owned alternatives negative. Reducing service capacity to 75% leaves purchase positive by USD 0.271 million and makes build negative by USD 0.218 million. The model's calculated break-even electricity prices are USD 24.922/MWh for purchase, USD 32.112 for build and USD 9.659 for partnership. These are conditional minimum receipts under the stated scope, with no claim about an actual power contract.

The joint adverse case combines 25% higher capital, 75% service capacity, fallback limited to one quarter of missing work, robot recovery of 75% and a USD 25/MWh electricity price. All alternatives become negative, including partnership at negative USD 0.021 million. No probability is assigned. This outcome supports a conditional recommendation: partnership leads the base numerical comparison, while an acceptable commitment still requires evidence on

recovery, productive service capacity, actual fee scope and fallback arrangements. A favourable expected operating narrative alone would leave these specific decision breaks unresolved.

The supporting calculation package retains 56 cases, 616 annual rows and 6,720 monthly rows. Forty-digit decimal arithmetic applies the stated workload, energy and cash equations before display rounding. Calculated checks reconcile energy losses, service shares, dated cash, fee boundaries and the pilot-stop ledger. These author checks support reproducibility and do not constitute independent validation of assumed inputs. A separate reviewer can reconstruct each case from the machine-readable inputs and equation specification without using the author's calculation functions. The scope remains a finite cleaning-procurement experiment; equipment approval, actual damage risk, intellectual property, enforceability and financing require evidence beyond this model.

Table 9. Joint and separate decision stresses

Assumption change	Buy NPV USDm	Build NPV USDm	Partner NPV USDm
Equal cleaning recovery	-1.479	-1.656	-0.501
Energy price 25	0.005	-0.386	0.953
Capital plus 25%	-0.078	-0.205	1.884
Service capacity 75%	0.271	-0.218	1.038
Two-year rollout delay	0.203	-0.507	1.366
Discount 12%	0.398	-0.131	1.593
Correlated adverse case	-2.242	-2.095	-0.021


Every stress has the same manual comparator except the stated electricity price. The joint case changes capital, service capacity, fallback, restoration and price without assigning a probability.

Table 10. Pilot gate and service-capacity decision points

Test	Input or quantity	Calculated result	Unit
Stop after pilot	Development plus pilot	1.600	USDm nominal
Stop after pilot	Incremental NPV	-1.570	USDm
Buy peak workload	Active units required	78.000	robots
Build peak workload	Active units required	81.000	robots
Partner full service	Required monthly peak	800.000	MW-passes
Partner base fee ceiling	Fixed fee remains USD 0.180m/year	99.794	USD/MW-pass
Partner equals buy	Fixed fee remains USD 0.180m/year	79.614	USD/MW-pass

Integer capacity thresholds use the stipulated 200 MWdc site and four passes in peak months. Stop is an exercised hypothetical gate, with no probability-weighted option valuation.

Figure 5. Incremental value under rollout, recovery and combined stresses



All alternatives share the same finite operating calendar. Joint stress changes capital, service capacity, fallback, recovery and electricity price; no empirical probabilities or actual contract terms are supplied.

19. Procure a pilot that resolves the ownership decision

A useful robotic-cleaning pilot should answer a defined investment question: which responsibilities can the asset owner economically retain, purchase or share at the intended sites? The proposed test specification should identify the module and tracker population, row transitions, access restrictions, charging arrangements and intervention crew. An apparently simple equipment demonstration can leave those interfaces untested. Procurement should therefore require a priced schedule of the activities performed by the vendor, the owner's staff and subcontractors during the trial. The investment team can then reconstruct the resources needed to repeat the work without temporary demonstration support. Any staffing, spares or supervisory service supplied without charge should appear in the commercial comparison at an expressly estimated replacement cost.

The tender should give every bidder the same test specification and require a returned exception schedule. Proposed deviations should identify their price consequence and the party responsible for the omitted activity. Procurement should retain the original bid alongside each clarification so that a later technical concession remains traceable to its commercial effect. The award recommendation should show whether a bidder has accepted the test, proposed a narrower envelope or priced an additional investigation. Unpriced exceptions should remain open conditions of award.

The proposed procurement gate would link payment to distinct deliverables: installation readiness, verified operation within the agreed test envelope, and transfer of usable records and

procedures. The supplier should be invited to price a failed-test exit, a redesign and a restricted deployment as separate outcomes. A pilot contract also needs an agreed method for handling disputed measurements and restoring the test area. Before approving a larger order, finance should reconcile every recurring intervention with the proposed build, buy or partner cost base. If a successful demonstration depends on an expert who cannot support the fleet, the decision should include the cost and lead time of replacing that dependency. Pilot acceptance should authorise only the configuration and operating envelope actually evaluated.

Table 11. Proposed pilot purchase and acceptance record

Evidence	Decision owner	Release condition
Test envelope and reference areas	Technical lead	Measurement design agreed
Interventions and free support	Finance lead	Repeat-service costs reconciled
Untested conditions	Asset manager	Rollout scope restricted
Pilot exit and restoration	Procurement lead	Exit obligations priced

Author-proposed controls; pilot results require site-specific measurement and review.

20. Assign development ownership and the cost of failure

A build strategy needs an agreed register of pre-existing technology and newly funded development. The proposed register should identify each contributor, the deliverable being commissioned and the permitted use of the resulting material under the negotiated agreement. Procurement should flag any component whose onward use depends on a separate third-party consent. The development budget should include the cost of resolving those dependencies before the affected design is released. Legal advisers should review ownership, licensing and employee or subcontractor arrangements against the actual documents; the investment paper should record unresolved items explicitly.

Warranty discussions should connect individual promises to the physical interfaces where loss could arise. Proposed schedules should distinguish the robot itself, module surfaces, mounting equipment, tracker movement, electrical connections and third-party property. The owner should request the applicable module and tracker documentation and written responses addressing the proposed cleaning configuration. Silence should remain unresolved in the decision record. Procurement and legal teams should examine exclusions, claim deadlines, evidence requirements, financial limits and the counterparty responsible for each obligation in the actual proposed documents. The investment case should retain gross repair and interruption exposure until a recovery mechanism has been reviewed and a justified collection assumption separately approved.

For an internally developed system, design release should require a controlled bill of materials, software version and authorised modification process. For a purchased system, the same record supports a determination of whether a later change falls inside the tested configuration. A supplier may propose replacing a brush, changing a control setting or moving to another component manufacturer; each change should have an identified approval route and a proportionate retest requirement. The commercial comparison should include investigation time and replacement logistics during unresolved failures. Proposed liability allocation should also specify who preserves evidence and who directs immediate protective work. Contract negotiation can allocate responsibilities, while the operating plan must still identify the people and cash available to act when equipment is unavailable.

Table 12. Proposed development and warranty interface register

Interface	Evidence required	Decision response
Design and software continuity	Usable technical handover	Price unresolved dependencies
Module and tracker compatibility	Applicable documents and responses	Hold affected deployment
Fault and damage recovery	Actual warranty terms	Retain gross cash exposure
Configuration changes	Version and retest record	Approve bounded release

No warranty entitlement, legal remedy or compatibility approval is assumed.

21. Buy an operating capability or govern a shared one

Acquiring a cleaning business requires evidence about the continuity of its operating capability. The proposed diligence perimeter should include customer-level service records, installation histories, staff responsibilities, parts sources and the practical dependencies behind reported fleet operation. Finance should reconcile customer invoices, collections, credits and unpaid remedial work with the service logs for the same contracts. Management estimates of service margin should identify central engineering support, travel, replacement equipment and development costs omitted from site accounts. Where customer relationships depend on named individuals or unrecorded modifications, the acquisition plan should price a documented transfer process and retain the associated execution risk in the investment memorandum.

A joint venture or operating partnership introduces a different decision problem: how the parties will authorise work when their commercial interests diverge. Proposed reserved decisions should cover design changes, new-site acceptance, customer credit, related-party procurement and spending beyond the agreed support budget. Each decision needs a workable escalation route and an operating fallback during disagreement. The asset owner should understand which party can commit the venture to an additional service obligation and who supplies cash when a customer disputes an invoice. A shared business plan should expose capital calls, losses and withdrawal obligations under the proposed agreements. An expectation of partner support should remain conditional until the relevant commitment and funding source have been established.

The transaction structure should also define the evidence required to recognise capability at closing. A proposed acquisition holdback could be discussed against specific unresolved deliverables, subject to negotiation and legal review; its existence would not establish future recoverability. For a partnership, an initial service area can provide a bounded basis for testing governance and operating interfaces before additional sites are admitted. The approval paper should distinguish acquired assets, licensed technology, access to staff and future development promises. Each item should have an owner, delivery date and consequence if absent. This makes the build-buy-partner comparison sensitive to the capability actually transferred and to the work the investor must still perform after signing.

Table 13. Proposed acquisition and partnership decision controls

Decision	Required record	Accountable owner
Acquire service capability	Logs, invoices and margin bridge	Investment lead
Transfer critical capability	Staff and technical handover	Operating executive
Approve shared commitments	Reserved decisions and authority	Venture board
Fund disputed service	Capital calls and liquidity terms	Finance lead

Proposed diligence and negotiation subjects; actual transaction terms determine rights and obligations.

22. Allocate cash through disputes, termination and replacement

The proposed service agreement should describe how operating records become invoices and how invoices become cash. Equipment availability, completed cleaning and measured energy improvement are different possible payment bases, each requiring its own acceptance evidence. The parties should define the treatment of interrupted work, repeat cleaning, owner-caused access restrictions and disputed data. Finance should model the proposed billing sequence using the expected certification and payment dates, with a separate case for a dispute. Deposits, retentions and performance security should be recorded according to their actual availability and purpose. Restricted balances should remain outside ordinary operating cash until the conditions for using them have been verified.

Termination diligence should focus on the service that must continue after the notice. The proposed exit schedule should identify equipment ownership, removal responsibilities, access to operating records, outstanding repairs and the handover of site-specific instructions. A replacement provider may need different charging equipment, row modifications or staff training. Those transition requirements should enter the owner's contingency budget even when the outgoing contractor is alleged to be responsible. The commercial team should also establish which outstanding amounts remain disputed and which payments are necessary to keep the site operating during transition. Recoveries should be represented as separately justified cash flows with their own timing and uncertainty.

The proposed step-in schedule should attach responsibility for notices, evidence custody and the handover meeting to named roles. An unsigned draft provision should remain a negotiation item. The owner should nominate the executive who decides whether to fund temporary operations and the finance officer who confirms available liquidity. During a disputed exit, the proposed daily record should distinguish protective work already authorised from additional work awaiting approval. Counsel should assess the actual access and termination provisions before the operating team treats them as executable instructions.

Table 14. Proposed cash and replacement-service map

Event	Cash evidence	Operating response
Accepted service	Certification and payment dates	Reconcile invoice basis
Disputed invoice	Liquidity and restricted balances	Fund approved continuity
Contract termination	Exit and replacement costs	Plan site handover
Proposed step-in	Access and licence review	Test executable replacement

Termination and step-in options require review of the actual agreements and operating prerequisites.

23. Mobilise the fleet around site readiness and support capacity

Mobilisation should be controlled through a site-by-site readiness record. The proposed record should identify the exact deployment configuration, permitted access windows, storage and charging arrangements, communications coverage and the responsible operating crew. Technical acceptance should include the interfaces between cleaning equipment and the installed module and tracker population. A portfolio contract may cover sites with different layouts and operating restrictions; each site should therefore retain its own acceptance conditions. Procurement should separate delivery of machines from completion of installation, training and operational handover. The planned service start should follow the dependencies established in that record, with unresolved conditions reflected in the cash schedule.

Support capacity needs a corresponding mobilisation test. The operator should demonstrate how faults are reported, triaged and assigned, where replacement parts are held and how a stranded or damaged unit is recovered. The initial deployment plan should include responsibility for work outside normal service hours and for simultaneous problems at more than one location. These are proposed planning tests, not assumptions about measured failure frequency. Finance should obtain dated costs for the staff, vehicles, spares and supplier support needed to deliver the selected response arrangement. If those resources serve several contracts, the allocation method should expose competing demands and avoid treating the same support capacity as fully available to every customer.

Rollout approval should follow a bounded sequence of installed capacity and evidence releases. The proposed first release would authorise a defined site group and fund the support resources associated with it. Expansion would depend on reviewed service records, unresolved defect exposure and confirmation that the support organisation can absorb the additional scope. The operating executive should retain authority to pause expansion when a new configuration or recurring fault requires investigation. Procurement commitments should be aligned with that authority so that a pause does not silently trigger avoidable equipment or subcontractor obligations. The mobilisation report should state what is operational, what has been delivered but remains conditional, and what expenditure remains cancellable under the reviewed agreements.

Table 15. Proposed mobilisation release gates

Gate	Evidence	Decision owner
Site readiness	Access, interfaces and charging	Site manager
Support readiness	Staff, parts and recovery plan	Operating executive
Initial service acceptance	Configuration-specific records	Technical lead
Fleet expansion	Capacity and defects review	Investment committee

Illustrative approval sequence; no assumed deployment timetable or failure frequency is supplied.

24. Convert diligence into an executable investment decision

The final investment paper should present a decision map built from the unresolved obligations of each route. For a build proposal, the decisive evidence may concern development completion and support continuity. An acquisition may depend on customer retention, the transfer of technical staff and the quality of historic service margins. A partnership may depend on agreed funding responsibilities and a workable deadlock process. The proposed comparison should use the same site scope and required service outcomes, then show the remaining tasks and liabilities for the investor under each route. A low quoted equipment price should enter the relevant cash schedule alongside the costs required to deliver an accepted operating service.

Each material diligence finding should lead to a specific decision response. A missing compatibility approval could restrict the initial site scope. An unresolved software dependency could require a continuity arrangement before commitment. A service-margin reconciliation failure could change the acquisition price or suspend approval of the affected earnings assumption. The investment committee should see the evidence, the proposed response, the accountable owner and the last date on which the decision can be changed without creating a further commitment. These proposed controls allow a recommendation to remain conditional while making the conditions testable. Findings that have no identified decision consequence should be reconsidered for relevance or explicitly retained as monitoring items.

Approval should conclude with a funded implementation mandate. The mandate should identify the expenditure authorised immediately, the subsequent releases, the people responsible for acceptance and the contingency cash available for transition or interruption. It should retain an unapproved-items register covering technical, contractual and commercial assumptions that remain open. Performance reporting should reconcile operating delivery with invoices, cash collected, support spending and unresolved claims. After a material change in configuration, customer scope or contract terms, the responsible executive should update the comparison using the revised obligations. The resulting decision is a documented choice about ownership, execution and cash exposure at the reviewed sites. Its validity remains conditional on the evidence and commitments that supported approval.

Table 16. Proposed diligence-to-investment decision map

Finding	Possible response	Acceptance evidence
Compatibility unresolved	Restrict initial scope	Reviewed site approval
Software dependency open	Require continuity arrangement	Usable handover test
Margin reconciliation incomplete	Revisit price assumption	Contract-level cash bridge
Partner funding conditional	Limit commitments	Approved funding terms

Responses are proposed decision controls and remain subject to technical, financial and legal review.

Frequently Asked Questions

What does buying mean in the numerical comparison?

The numerical buy case purchases an equipment fleet and retains defined operating responsibilities with the owner. It excludes acquiring the supplier's business. A corporate acquisition requires separate diligence on earnings, customers, staff, intellectual property, liabilities and integration costs.

When should an owner consider developing its own system?

The proposed development decision requires a funded engineering and support plan, an approved beneficiary fleet and cancellable stages before commercial rollout. The committee should compare its full development commitment, pilot outcome and deployment delay with feasible purchased and contracted alternatives.

How should cleaning capacity enter the investment model?

Translate the required passes at each site into work within the available service window. Compare that work with the stipulated active fleet and service availability. Cost the remaining manual work, identify its capacity limit and test the generation consequences when the work cannot be completed.

Can the illustrative fee ceiling be used as a supplier quotation?

The fee ceiling is a conditional calculation using the paper's fictional service scope and cash assumptions. It is no supplier quotation or market benchmark. An actual negotiation needs the contracted scope, retained owner costs, settlement rules, failure response and funding requirements.

What should be approved before a pilot expands?

The proposed approval should identify the tested configuration and site envelope, unresolved compatibility questions, support resources, costed interventions and outstanding obligations. Expansion should be limited to the scope supported by reviewed evidence and the funding authorised for delivery and contingency work.

How does the framework address service-provider failure?

It specifies the work interrupted, the temporary service available, replacement costs and the permissions and equipment needed for recovery. Contract rights require review of the actual agreements. The investment comparison should retain the owner's gross funding requirement and separately justify any assumed recovery.

References

- [1] IEA PVPS Task 13. Soiling Losses: Impact on the Performance of Photovoltaic Power Plants. January 2023; Report T13-21:2022. Official summary, publication month and Key Takeaway 3 on spatial heterogeneity and measurement. <https://iea-pvps.org/key-topics/soiling-losses-impact-on-the-performance-of-photovoltaic-power-plants/>
- [2] Bomber et al.; National Renewable Energy Laboratory. Soiling, Cleaning, and Abrasion: The Results of the Five-Year Photovoltaic Glass Coating Field Study. 2024. Official abstract and presentation metadata; NREL/PR-5K00-90174. Supplementary PDF also inspected; glass-coupon scope. <https://research-hub.nrl.gov/en/publications/soiling-cleaning-and-abrasion-the-results-of-the-five-year-photov/>
- [3] International Electrotechnical Commission. IEC 62788-7-3:2022+AMD 1:2024 CSV: Measurement procedures for materials used in photovoltaic modules - Part 7-3: Accelerated stress tests - Methods of abrasion of PV module external surfaces. 31 July 2024; edition 1.1. Public scope and product metadata; complete paid standard not inspected. <https://webstore.iec.ch/en/publication/99014>
- [4] LONGi Solar Technology Co., Ltd. PV Module Operation and Maintenance Manual. April 2025; version 03. Disclaimer, printed p. 4; sections 2.2.2 and 2.2.5, printed pp. 17-18; revision, p. 26. PDF spreads contain two printed pages. https://eu-files.longi.com/api/public/dl/tfxflagP/EN/31_Maintenance_Manual_EN_V3.0.pdf?inline=true
- [5] Trina Solar Co., Ltd. Trina Solar User Manual: Vertex Series Modules. November 2023; UM-M-0002, version J. Applicable module list; disclaimer, printed p. 0; section 7.3, Method D Cleaning robot, printed p. 33; version history, p. 33. https://static.trinasolar.com/sites/default/files/UM-M-0002_J_210UserManual_EN.pdf
- [6] IEA PVPS Task 13; Jahn et al. Guidelines for Operation and Maintenance of Photovoltaic Power Plants in Different Climates. October 2022. Executive summary, PDF pp. 1-2: negotiated O&M scope and KPIs; hot-and-dry logistics. T13-25:2022. https://iea-pvps.org/wp-content/uploads/2022/11/Final-Executive-Summary-IEA-PVPS-T13-25_2022.pdf
- [7] NREL, Sandia National Laboratories, SunSpec Alliance and SuNLaMP PV O&M Best Practices Working Group. Best Practices for Operation and Maintenance of Photovoltaic and Energy Storage Systems, Third Edition. December 2018. Section 6.4, printed pp. 49-50; sections 8.3 and 8.5, printed pp. 71 and 73. NREL/TP-7A40-73822. <https://docs.nrl.gov/docs/fy19osti/73822.pdf>
- [8] IEA PVPS Task 13; Woyte et al. Analytical Monitoring of Grid-connected Photovoltaic Systems: Good Practices for Monitoring and Performance Analysis. March 2014. Title page, PDF p. 3; executive summary, printed p. III, monitoring equipment and periodic regression analysis. T13-03:2014. https://iea-pvps.org/wp-content/uploads/2020/01/IEA-PVPS_T13-D2_3_Analytical_Monitoring_of_PV_Systems_Final.pdf
- [9] Sweihan PV Power Company. Noor Abu Dhabi Solar PV Project: Green Bond Framework. September 2021. Section 1.2, printed p. 2, issuer statement on plant-powered waterless cleaning robots; dated cover. <https://noorabudhabi.ae/wp-content/uploads/2022/04/Noor-Abu-Dhabi-Green-Bond-Framework.pdf>
- [10] TrinaTracker; Trinasolar. TrinaTracker Launches Two Robots to Expand its Smart PV Ecosystem. 24 September 2026. Dated announcement and AURORA cleaning-robot subsection; manufacturer claims only. <https://www.trinasolar.com/en-glb/newsroom202609240931/>
- [11] Dubai Electricity and Water Authority. DEWA's Autonomous Soiling Detector improves solar PV system performance and cuts dust-related energy losses. 30 December 2025. Prototype location and actual-versus-expected production paragraphs; no quantitative outcome adopted. <https://www.dewa.gov.ae/en/about-us/media-publications/latest-news/2025/12/dewas-autonomous-soiling-detector>
- [12] United States Department of Energy. Adoption Readiness Levels. Undated; accessed 9 October 2026. Introduction, framework description and four core risk areas; qualitative commercialization framework. <https://www.energy.gov/arl/adoption-readiness-levels>

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

Disclaimer

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.