© June 2026 | IJIRT | Volume 13 Issue 1 | ISSN: 2349-6002

Why Pv Module Faults Cost Generation

Deepak Sharma

***Abstract*—Solar photovoltaic plants rarely fail outright. Far more often they bleed energy slowly and quietly, through faults that depress output by a few percent at a time and never trip an alarm. Across a large plant, and across a thirty-year operating life, these losses are enormous and most are recoverable. The difference between a plant that performs near its potential and one that quietly underperforms is rarely the hardware; it is whether faults are found early and acted on, or left to accumulate.**

This white paper examines, for an O&M engineering audience, exactly how module faults cost generation and how that generation is recovered. It traces the physical mechanisms by which the common faults soiling, cell cracks, hotspots, potential-induced degradation, connection and busbar failures, and underperforming strings convert into lost kilowatt-hours. It then sets out the diagnostic methods that detect them and the analytic techniques that catch them early, prioritise them by value, and verify their recovery. The central argument is simple: fault-driven loss is not an unavoidable cost of operation but a recoverable one, and the tools to recover it are well understood and available today.

***Index Terms*—Most fault-driven loss is gradual and invisible to simple alarms; it is found only by actively looking for it in performance data.**

- **Each fault type has a distinct physical mechanism and a distinct, detectable signature in the I–V curve, thermal image, or performance trend.**
- **Recovery follows a repeatable loop: detect, diagnose, prioritise by recoverable value, act, and verify.**
- **Predictive analytics shifts the work from reacting to failures to anticipating them, recovering generation while losses are still small.**
- **A disciplined detection-and-recovery programme typically returns several percent of generation a large, fast payback against modest O&M cost.**

I. INTRODUCTION: THE QUIET COST OF FAULTS

A solar plant's nameplate capacity is a promise made under laboratory conditions. What it actually delivers depends on how well its modules convert the sunlight they receive, day after day, for decades. Between the promise and the delivery sits a chain of losses some unavoidable physics, some recoverable faults. This paper is about the second kind: the faults that quietly erode output and that a well-run O&M programme can find and fix.

The defining characteristic of fault-driven loss is that it is usually invisible to casual observation. A plant losing six percent to a combination of soiling, a handful of weak strings, and a developing connection fault looks, on a sunny day, like a plant that is working. The inverters are online, the meter is spinning, and nothing is obviously wrong. The loss reveals itself only when output is compared against what the plant should be producing under the conditions it is actually experiencing and that comparison is the foundation of everything in this paper.

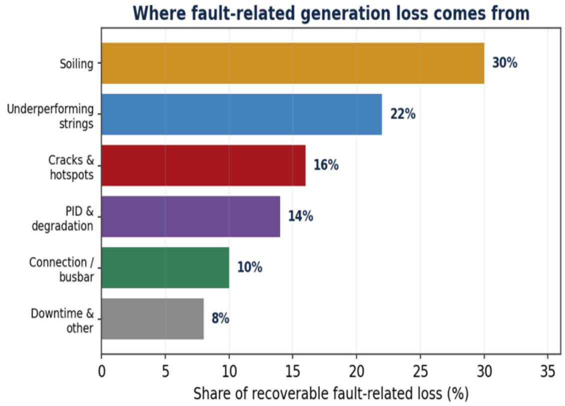

Figure 1 A representative breakdown of where recoverable, fault-related generation loss comes from on a typical utility-scale plant. The exact shares vary by site, climate, and age, but the pattern losses spread across several mechanisms, none dominant is consistent.

As the figure suggests, fault-related loss is rarely concentrated in a single dramatic failure. It is distributed across several mechanisms, each taking its share: soiling, underperforming strings, cracks and hotspots, potential-induced degradation, connection and busbar problems, and downtime. This distribution matters, because it means there is no single fix. Recovering the loss requires a systematic programme that addresses each mechanism in turn, prioritised by the value each returns.

The scope of this paper is the module and the string the generating heart of the plant and the faults that arise there. It is written for the O&M engineer who must find these faults and recover the generation they cost, and it assumes familiarity with basic PV operation while developing the diagnostic and analytic reasoning in full.

II. HOW A FAULT BECOMES LOST ENERGY

To recover fault-driven loss, an engineer must first understand the physical chain that turns a defect into a missing kilowatt-hour. Almost every module fault acts through one of two mechanisms, both visible in the single-diode model that describes a solar cell: it either reduces the current the cell can generate, or it raises the internal resistance that current must overcome. Understanding which mechanism, a fault uses is the key to both detecting it and quantifying its cost.

2.1 The Two Master Mechanisms

A solar cell behaves like a light-driven current source in parallel with a diode, with two parasitic resistances: a series resistance representing the contacts and connections, and a shunt resistance representing leakage paths across the cell. Faults express themselves as changes to these elements.

$$I = I_L - I_0\,[\exp(q\cdot V / (n\cdot k\cdot T)) - 1] - (V + I\cdot R_s)/R_{sh} \quad (1)$$

where:

I_L = light-generated current reduced by soiling, shading, cell inactivity

R_s = series resistance raised by corroded connections, weak solder joints, busbar damage

R_{sh} = shunt resistance lowered by cracks, PID, and leakage paths

A fault that blocks light or disables cells lowers the generated current. A fault that degrades a connection raises series resistance, which bends the I–V curve down near the operating voltage and reduces the power that can be drawn. A fault that creates a leakage path lowers shunt resistance, which bleeds current away near short-circuit. Each leaves a characteristic mark on the I–V curve, and reading that mark is how an engineer identifies the fault and estimates its cost.

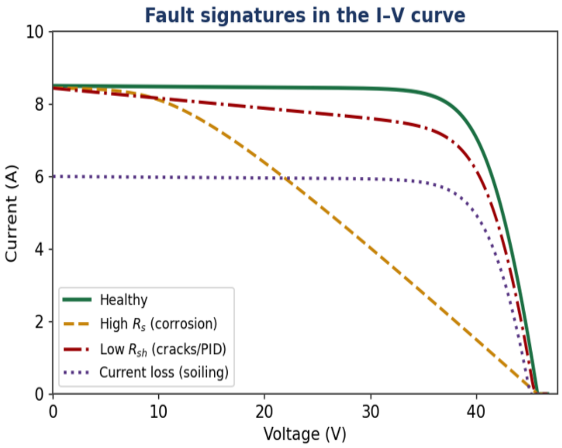

Figure 2 The signatures of the two master mechanisms in the I–V curve. Rising series resistance (corrosion, weak joints) bends the curve down near the operating voltage; falling shunt resistance (cracks, PID) slopes it near short-circuit; a current loss (soiling, mismatch) drops the whole curve. The shape names the fault.

2.2 From Fault to Lost Kilowatt-Hours

A fault's cost in energy is the gap between the power the affected modules should produce under the prevailing conditions and the power they actually produce, integrated over time. Because the expected power can be calculated from the measured irradiance and temperature, this gap is directly measurable which means a fault's cost can be quantified in kilowatts and, with an energy price, in currency. This is the basis on which faults are later prioritised.

$$P_{expected} = P_{stc} \cdot (G / 1000) \cdot [1 + \gamma\cdot(T_{mod} - 25)] \quad (2)$$

The shortfall between expected and measured power, summed over the hours the fault persists, is the lost energy. A string running five percent below expectation in a sunny region loses meaningful energy every single day the fault is left unaddressed and because many faults worsen over time, the daily loss grows until the fault is found. This compounding is why early detection is worth so much: the same fault costs far less if caught in its first weeks than if left to develop for years.

The Cost of Delay

A developing fault is rarely static. A weak connection grows more resistive as it heats; a crack propagates;

PID deepens; soiling accumulates. Each day a fault goes undetected, it typically costs a little more than the day before. The value of detection is therefore not merely the loss caught on the day it is found, but the larger, growing loss prevented over all the days that would have followed. This is the core economic argument for active detection over passive operation.

III. THE PRINCIPAL FAULTS AND WHAT THEY COST

This section examines the faults responsible for most recoverable loss, each through the same lens: the mechanism by which it costs generation, the signature by which it is detected, and the typical scale of its impact. Together they account for the great majority of fault-driven loss on a working plant.

3.1 Soiling

Soiling the accumulation of dust, dirt, and debris on the module surface is the most universal loss mechanism. It works by the first master mechanism: it blocks light, reducing the generated current uniformly across the affected modules. Its cost ranges from a percent or two in clean, rainy climates to well over ten percent in dry, dusty regions during long rain-free periods.

Mechanism: blocks incident light, lowering generated current; the whole I–V curve drops proportionally.

Signature: a gradual, plant-wide decline in temperature-corrected performance that recovers sharply when the modules are cleaned or rained on the characteristic saw-tooth.

Scale: 1–10%+ depending on climate and cleaning interval; recoverable in full by cleaning.

3.2 Underperforming Strings

Individual strings that lag their peers through a weak module, partial shading, a developing connection fault, or a conducting bypass diode are among the most common and most recoverable losses. Because the fault is localised to specific strings, the rest of the plant masks it in any aggregate metric; it is found only by comparing strings against one another or against their expected output.

Mechanism: varies by root cause reduced current, raised series resistance, or a bypassed substring but always shows as a string producing below its peers.

Signature: a persistent shortfall of one string against the median of its peer group, or against its irradiance-corrected expected power.

Scale: the affected strings may each lose 5–20%; the plant-level impact depends on how many strings are affected.

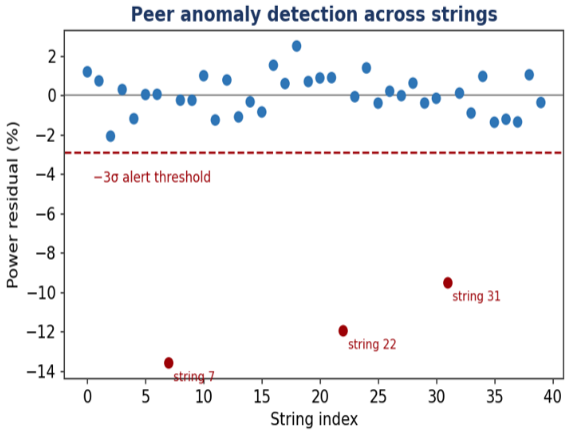

Figure 3 Peer comparison across the strings of one inverter. Most cluster near the group median; the few that fall well below it are flagged for inspection. Because peer strings share the same weather, this comparison needs no irradiance sensor and isolates the underperformers directly.

3.3 Cell Cracks and Hotspots

Mechanical cracks in cells, and the hotspots they often cause, degrade output and in the case of hotspots threaten the module itself. A crack can isolate part of a cell, reducing its contribution; if the isolated region is forced to carry current it cannot generate, it dissipates that power as heat, creating a hotspot that can scorch the encapsulant and crack the glass. The electrical loss is compounded by a safety and longevity risk.

Mechanism: cracks lower shunt resistance and inactivate cell area; hotspots add destructive local heating.

Signature: dark regions in electroluminescence imaging; localised temperature anomalies in thermal imaging.

Scale: a few percent per affected module electrically, but with a disproportionate risk to module life and safety.

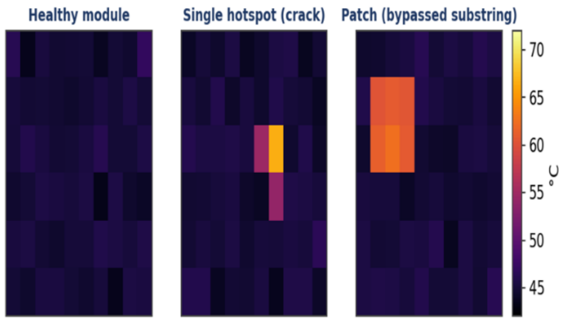

Figure 4 Thermal signatures. A healthy module shows uniform temperature (left); a single hotspot points to a cracked or reverse-biased cell (centre); a warm patch indicates a bypassed substring (right). Thermal imaging makes these faults visible across a whole plant quickly.

3.4 Potential-Induced Degradation

Potential-induced degradation (PID) is among the most consequential and least visible faults in large plants. Driven by the high system voltages of modern arrays and the potential difference between cells and the grounded frame, it creates leakage paths that lower shunt resistance and can strip a significant fraction of a module's power over months while showing nothing on a casual visual inspection. It tends to concentrate at the high-voltage end of strings.

Mechanism: voltage-driven leakage lowers shunt resistance, bleeding away current and power.

Signature: a slow power loss concentrated at the high-voltage end of strings; a characteristic mottled pattern in electroluminescence.

Scale: can reach tens of percent in severe, untreated cases; often partially recoverable if caught early.

3.5 Connection and Busbar Faults

Every current-carrying joint in a plant is a potential site of rising resistance. Corroded connectors, weak solder joints, and degraded busbars raise series resistance, reducing power and because the dissipated heat is proportional to the square of the current generating localised heating that can escalate to busbar burning, a fire risk. These faults grow worse over time as heating accelerates the degradation.

Mechanism: raised series resistance reduces power and dissipates heat (I^2R) at the joint.

Signature: reduced fill factor and a curve bent down near operating voltage in I–V tracing; a hot line or point in thermal imaging.

Scale: from a fraction of a percent early on to total string loss and a safety hazard if left to burn.

3.6 The Loss Picture as a Whole

Stepping back from the individual faults, it is useful to see the whole loss picture how the plant's nameplate potential is reduced, step by step, to delivered energy. Fault-related losses sit alongside the unavoidable physical losses (temperature, conversion) in a chain from DC nameplate to AC delivered. The recoverable fraction soiling, underperforming strings, connection losses, avoidable downtime is the target of a recovery programme.

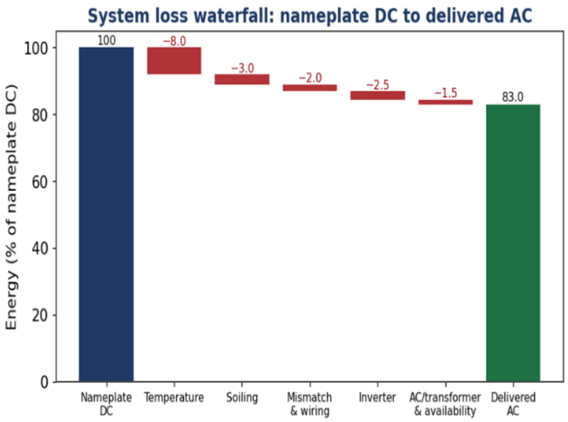

Figure 5 The system loss waterfall, from nameplate DC potential to delivered AC energy. Temperature loss is largely fixed by physics and siting; soiling, mismatch, and downtime are substantially recoverable. The recoverable losses are where an O&M programme acts.

IV. DETECTION: FINDING THE LOSS

A fault cannot be recovered until it is found, and the defining challenge of fault-driven loss is that most of it is invisible to the simple monitoring most plants rely on. Detecting it requires actively comparing what the plant produces against what it should produce, at every scale from the whole plant down to the individual module. This section sets out the detection methods, from the instrument-based inspection of individual modules to the data-driven monitoring of the whole plant.

4.1 The Detection Principle

Every detection method in PV rests on one move: normalise out the weather, then compare against an expectation. A raw power reading is meaningless without the irradiance and temperature that produced it; corrected for those, it can be compared against the datasheet, against peer strings, or against the plant's own history. What changes from method to method is the reference the datasheet for a module, the peer group for a string, the modelled expected energy for a plant not the underlying logic.

The Three Things Measured Together

Every performance measurement that can detect a fault captures three quantities at the same instant:

- Electrical output power, or the full I–V curve.
- Plane-of-array irradiance the sunlight actually striking the modules.
- Module temperature measured at the cell, not the ambient air.

Without all three, the reading cannot be corrected, and an uncorrected reading cannot reveal a fault

4.2 Instrument-Based Detection

At the module and string level, a small set of instruments reveals the faults of Section 3 directly. Each instrument is suited to particular faults, and together they form a complete diagnostic toolkit.

Method	Faults Revealed	What It Shows
Thermal (IR) imaging	Hotspots, cracks, bypassed substrings, hot connections	Temperature anomalies fast, whole-plant, often by drone
I–V curve tracing	Series/shunt resistance faults, degradation, mismatch	The quantitative curve shape names the fault and its cost
Electroluminescence	Micro-cracks, inactive cells, PID patterns	Dark regions in the silicon's infrared emission
Insulation testing	Moisture ingress, backsheet/insulation faults	Insulation resistance a safety and health indicator
Visual inspection	Soiling, breakage, discolouration, snail trails, burning	What the eye can see always the first step

4.3 Data-Driven Detection

Instrument-based methods find faults when an engineer goes looking; data-driven methods find them continuously; from the telemetry the plant already produces. This is where the recovery of fault-driven loss scales from occasional inspection to constant vigilance. Three techniques carry most of the load.

Expected-power residuals: compare each string's measured power against the power physically expected from the measured irradiance and temperature; a persistent shortfall is the fault signal, quantified directly in lost kilowatts.

Peer comparison: compare each string against the median of its peers on the same inverter; because peers share the weather, this isolates a lagging string without needing any irradiance sensor.

Trend and change-point analysis: watch the temperature-corrected performance over time, distinguishing the saw-tooth of soiling, the slow drift of degradation, and the abrupt step of a developing fault.

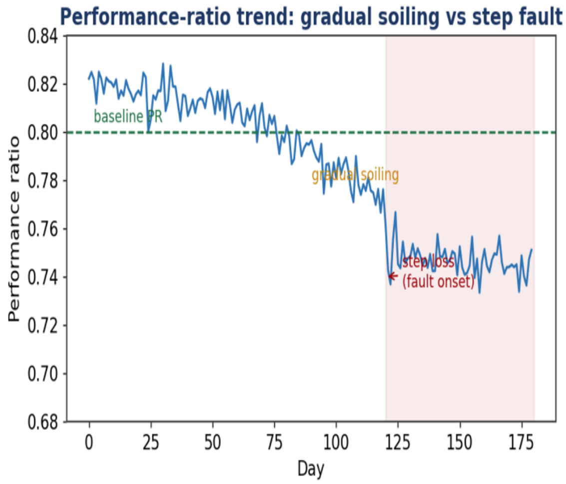

Figure 6 A performance-ratio trend reads two faults at once: the gradual downward drift of soiling, and the abrupt step of a fault's onset. Distinguishing these patterns and acting on them is the essence of data-driven detection.

4.4 From Detection to Anticipation

The most valuable detection happens before a fault has done much damage. Because many faults develop gradually a connection slowly heating, soiling accumulating, degradation deepening their early signatures are present in the data long before they become severe. Predictive analytics exploits this: by trending the slow signals and flagging deviations as they emerge, it shifts the work from reacting to failures that have already cost generation, to anticipating them while the loss is still small. This is the difference between finding a fault that has cost a season's energy and catching the same fault in its first weeks.

V. RECOVERY: A REPEATABLE METHODOLOGY

Detection creates knowledge; recovery turns it into regained generation. The recovery of fault-driven loss follows a repeatable loop that applies to any fault, at any scale: detect the deviation, diagnose its cause, prioritise it by recoverable value, act, and verify the recovery. Run consistently, this loop transforms a plant from one that loses generation invisibly to one that recovers it systematically.

5.1 The Recovery Loop

1. Detect a metric, alert, or inspection flags a deviation from expected performance.
2. Diagnose the deviation's signature (I–V shape, thermal pattern, trend) identifies the fault and its root cause.
3. Prioritise the fault is ranked by the value it returns against the cost to fix it, so the most valuable work is done first.
4. Act the appropriate correction is applied: cleaning, replacement, re-termination, or PID mitigation.
5. Verify a follow-up measurement confirms the generation has been recovered, and the outcome is recorded.

5.2 Prioritising by Recoverable Value

A real plant surfaces more candidate faults than a crew can address at once, so prioritisation is essential and it must be driven by recoverable value, not by how alarming a fault looks. A dramatic-looking cracked module that costs little generation should rank below an unremarkable string quietly running six percent low, because the goal is recovered energy per unit of effort, not visual severity.

$$\text{priority score} = (\text{recoverable annual value} \times \text{urgency}) / \text{cost to fix} \qquad (3)$$

Computing a recoverable value for each fault from its measured power shortfall, the site's sun-hours, and the energy price turns a disordered list of problems into a ranked work plan. The crew works the list from the top, and each intervention is the most valuable one available at that moment. This single discipline is often the largest practical improvement a plant can make to its O&M effectiveness.

Why Appearance Misleads

The faults that cost the most generation are frequently the least dramatic. A whole-plant soiling loss, a handful of strings each a few percent low, a slowly rising connection resistance none looks urgent, yet together they dwarf the occasional shattered module that draws the eye. Prioritising by measured recoverable value, rather than by appearance, is what directs the crew's limited time to where it actually recovers the most energy.

5.3 Matching Action to Fault

Each fault has a corresponding recovery action, and applying the right one and addressing the root cause, not just the symptom is what makes the recovery durable.

Fault	Recovery Action	Durability
Soiling	Clean on an economically optimised schedule	Recurs; manage by schedule
Underperforming string	Diagnose and fix root cause (module, connection, diode)	Durable once cause is fixed
Cell crack / hotspot	Replace affected module; address mechanical cause	Durable; prevents escalation
PID	Apply PID mitigation; address grounding/voltage stress	Often partially recoverable
Connection / busbar	Re-terminate or replace; fix moisture/torque cause	Durable; prevents fire risk

5.4 The Economics of Recovery

The case for a detection-and-recovery programme is ultimately economic, and it is compelling. Recovered generation accumulates continuously once a programme is running, while the costs are largely one-time interventions. Because each recovered percentage point of performance is earned every year for the remaining life of the plant, even modest recovery compounds into a large lifetime value against a small recurring O&M cost.

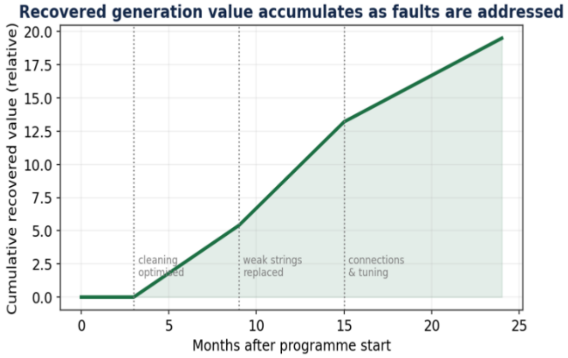

Figure 7 Recovered generation value accumulates as successive interventions land cleaning optimised, weak strings replaced, connections and tuning addressed. Each step adds to a cumulative recovery that, sustained over the plant's life, far exceeds the cost of the programme.

The Value of One Point, Sustained

Consider a 10 MW plant generating roughly 15 million kWh a year. A single percentage point of performance is about 150,000 kWh worth a meaningful sum every year at any realistic energy price. A recovery programme that returns several points, and holds them, therefore returns that value many times over across a thirty-year life, against intervention costs that are a small fraction of it. The payback is typically measured in months, which is why recovering fault-driven loss is among the highest-return activities in PV operations. (Figures here are illustrative; actual values depend on plant size, climate, and energy price.)

VI. ILLUSTRATIVE RECOVERY: A TIRED PLANT

To show the methodology in operation, consider a representative case an older plant whose performance has settled well below its achievable level, not through any single failure but through years of accumulated, unaddressed faults. The example is illustrative, with figures chosen to show the method rather than to describe any specific plant.

A plant performing at 78% performance ratio, against an achievable 84–86%, presents a six-to-eight-point gap with no single obvious cause. Rather than guess, the operator runs the full detection suite expected-power residuals, peer comparison, a soiling estimate, a thermal survey, and sample I–V sweeps and apportions the gap among its causes. Each is then addressed in order of recoverable value.

Intervention	Cause Addressed	Recovery
Full cleaning	Heavy accumulated soiling	+2.6 points
Replace weak strings	Degraded / faulted modules	+1.4 points
Re-terminate DC connections	Series-resistance losses	+0.7 points
Inverter tuning	Sub-optimal operation	+0.5 points

Each intervention is verified before the next is credited, so the gains are real and additive. Together they lift the plant from 78% to roughly 83% recovering most of the gap, with the remainder attributable to the irreducible degradation of an older fleet. The pattern is typical: cleaning, the cheapest action, returns the largest single gain; the diffuse connection losses, invisible in any single string, are recovered by systematic re-termination; and the whole is documented intervention by intervention. No single heroic fix was involved only the disciplined application of the recovery loop, fault by fault, in order of value.

The Lesson of The Case

A tired plant is recovered not by one dramatic repair but by the patient, prioritised application of many small ones, each verified before the next. The recovery follows recoverable value cleaning first, then weak strings, then connections, then tuning and the cumulative result transforms the asset's economics. This is the recovery methodology of Section 5 applied end to end on a real plant.

VII. CONCLUSION

Module faults cost generation through a small number of well-understood physical mechanisms, each leaving a detectable signature and each carrying a quantifiable cost in lost energy. The loss is real and, across a large plant over a long life, substantial but it is overwhelmingly recoverable. What separates a plant that performs near its potential from one that quietly underperforms is not better hardware but a disciplined programme of detection and recovery: actively comparing output against expectation, finding faults while they are small, prioritising them by recoverable value, acting, and verifying.

For the O&M engineer, the implications are practical and immediate. Most fault-driven loss will never announce itself; it must be sought in the data. The instruments and analytics to find it are mature and available. The recovery loop is simple and repeatable. And the economics are decisive: recovered generation, earned every year for the life of the plant, repays the modest cost of the programme many times over. The faults that cost generation are, for the most part, faults that can be recovered and recovering them is among the highest-return work in solar operations.

The Central Message

Fault-driven generation loss is not an unavoidable cost of operating a PV plant. It is a recoverable one. The mechanisms are understood, the detection methods are mature, the recovery loop is repeatable, and the economics are compelling. A plant run on evidence its faults sought, prioritised, fixed, and verified delivers across its life far more than one left to lose generation in silence. That difference, compounded over thirty years, is the prize this paper has described.

About this paper

This white paper is drawn from the technical handbook Solar PV Modules: The Complete Engineering Handbook by Deepak Sharma, which develops the faults, diagnostics, analytics, and recovery methods summarised here in full depth. The methods described reflect established PV engineering practice; specific procedures, thresholds, and acceptance criteria should be confirmed against current standards, manufacturer documentation, and applicable codes before being relied upon in the field.