THE ROOF COMES FIRST · WHITEPAPER

What does an inspection actually prove?

Detection, the inspection gap and the combustible roof beneath rooftop PV

CENTRAL QUESTION

What protects a combustible roof when a relevant fault is not detectable, is not detected or is not corrected before it becomes an ignition source?

Olivier Langejan

Chief Commercial Officer, AllShield

ol@allshield.com | allshield.com/en/olivier-langejan

Prepared for insurers, risk engineers, asset owners, inspectors and PV designers.

Version 1.03 | 1 November 2026

Contents

Author's perspective....3
Executive summary....4
1. Why this question matters....6
2. What inspection does well....8
3. What an inspection actually establishes....10
4. What may remain unseen....12
5. The inspection gap....14
6. Detection between inspections....16
7. Consequence-limiting design at roof level....18
8. What this means in practice....20
9. What we still do not know....22
Conclusion....24
References....25
Method and sources....26
REVISION HISTORY....28
THE SERIES....29

Author’s perspective

This series is a synthesis of published research, fire tests, standards, approvals and insurer guidance on fire safety and rooftop PV, written from an explicit roof-first perspective: I look at the complete commercial roof system and the building beneath it rather than at individual product classifications. Within that perspective I summarise what the evidence shows, add my own technical analysis and, where relevant, propose design principles. Those proposals are marked as mine and kept separate from what has been demonstrated.

My background is in architecture and building engineering. Over the past three years I have discussed this subject with engineers, risk specialists, insurers and test institutes, and I have been involved in setting up and supervising fire-test programmes for AllShield's own roof systems. That experience shapes the questions I ask; the papers test those questions against external evidence. How the sources were selected, and what remains open, is described at the end of each paper.

These papers are prepared by Olivier Langejan as lead author with support from a small team. His own professional focus is building and roof design; specialist findings are drawn from the cited evidence, with AI assistance disclosed separately in Method and sources.

I am affiliated with AllShield, which develops BarrierSheet and other roof fire-protection solutions. That commercial relationship is relevant and is therefore stated explicitly. The analysis does not assume that an AllShield product provides the effects discussed. The starting point is the physical or economic problem. Where an AllShield product is discussed, it is treated as one product-specific example and its performance should be read only within the boundaries of the cited tests, approvals and roof assemblies.

In this paper I apply that roof-first approach to one specific question: what does a satisfactory inspection actually prove about fire risk, and what protects a combustible roof when a relevant fault is not found in time? This paper was written as part of my role as Chief Commercial Officer at AllShield. The subject touches directly on AllShield’s commercial interests: if inspection and monitoring cannot remove the dependence on detection, a protective function at roof level becomes part of the discussion, and AllShield supplies such measures. For that reason, this paper does not argue that inspection should be reduced or that money should be diverted from inspection to roof protection. It contains no product case study and uses no AllShield product as evidence. The argument in Chapters 1–9 is product-neutral.

One relationship is relevant to the evidence cited here. The comparison of national PV fire statistics used in Chapter 9 was led by researchers at ZAG [15], where AllShield commissioned roof-PV fire testing. That study is used only for its findings on incident reporting. The document behind this statement is listed after the references.

None of the conclusions depends on an AllShield product being selected. Whether a particular product performs the consequence-limiting function is a separate evidence question, addressed in *What does a fire test actually prove?* and *Beyond flame spread*, and an existing roof construction may already provide an acceptable risk position.

Olivier Langejan
Chief Commercial Officer, AllShield

Executive summary

Inspection is one of the most important controls available for rooftop photovoltaic systems. It identifies installation defects, damaged components and thermal abnormalities, and where deficiencies are corrected in time it interrupts an escalation pathway before a defect becomes an ignition source. This paper does not question that value. It asks a narrower question: what does a satisfactory inspection actually prove about fire risk, and what protects a combustible roof when a relevant fault is not found in time?

The evidence sets two boundaries. The first is detectability. Thermography does not detect an electrical defect directly; it detects a thermal signature associated with it, which is why IEC TS 62446-3 prescribes minimum irradiance and operating current for a valid measurement. Controlled experiments show that some defects produce strong thermal signatures while others cannot be distinguished reliably from healthy cells, and that irradiance matters more than camera resolution. Connectors and cable joints beneath the modules add an accessibility boundary: a component outside the field of view cannot be measured, however good the camera. Under the Dutch Scope 12 scheme such concealed connections are documented with photographs from installation, a sensible practice that also means their assessment rests on records rather than on observation at the time of inspection. A study of 6,276 connectors harvested from rooftops confirms that severe connector degradation occurs in the field, and shows that fully melted connectors can measure lower resistance than less damaged ones.

The second boundary is time. A satisfactory report describes an assessed condition; it does not preserve that condition until the next inspection. Connections age non-linearly, cables degrade, inverters are replaced and roofs are repaired, and every intervention adds new workmanship. A fifteen-year-old system may contain a connection made six months ago. Continuous monitoring narrows part of this gap, but it remains detection-dependent: Sandia research shows that no default insulation-monitoring threshold can both avoid nuisance trips and guarantee a trip before a fire hazard develops. Different detection methods can therefore share one dependency: the fault must first produce an observable consequence.

No universal probability of detection exists for fire-relevant rooftop PV defects, and national fire statistics do not record the inspection history of affected installations. The residual risk after a satisfactory inspection cannot therefore be expressed as a known percentage. For PV above combustible roof construction, that uncertainty is itself relevant to design.

> **Core principle (author's conclusion)**
>
> **A passive roof-level fire protection measure does not replace inspection. It reduces the extent to which fire safety depends on a relevant defect being detectable, detected and acted upon in time.**

A passive layer has its own dependencies. It must be correctly specified and installed, it cannot be inspected periodically beneath the roof covering, and it must be reinstated after every intervention that disturbs it. It is not a reason to inspect less. The paper closes with practical questions for owners, insurers, inspectors and contractors, including the recommendation that

reassessment be linked to events such as inverter replacement and roof work, and not only to the calendar.

1. Why this question matters

Inspection is one of the most important controls available to owners of rooftop photovoltaic systems. It can identify installation defects, damaged components, abnormal operating conditions and deviations from applicable standards, and it can initiate corrective action before those deficiencies develop into larger technical or safety problems. Nothing in this paper argues otherwise.

The question I address here is narrower: what does a satisfactory inspection actually prove about fire risk?

That question matters because an inspection report can acquire a meaning beyond the scope of the inspection itself. In market communication, a system that has passed an electrical or thermographic inspection is sometimes described as safe or as presenting no elevated fire risk. Such language suggests that the relevant fire hazards have been eliminated. That is not what an inspection establishes. It assesses an installation using defined procedures, accessible components and measurements obtained under particular operating and environmental conditions. Where defects are detectable and identified, the process can be highly effective: they are reported and repaired.

> **KEY DISTINCTION**
>
> **The absence of an identified defect is not the same as proof that every relevant defect is absent.**

The distinction is particularly important for rooftop PV because the installation does not stand still. Electrical connections are exposed to thermal cycling, mechanical movement, moisture and ageing. Cables and their supports remain exposed to the roof environment. Inverters are replaced, arrays are extended and roofs are repaired, each time introducing new connections and new workmanship long after commissioning. The installation that is inspected is therefore not necessarily the installation that exists a year later.

Some defects are also easier to identify than others. Thermographic inspection depends on a defect producing a measurable thermal effect under the conditions present during the inspection. Electrical monitoring depends on a fault affecting a monitored parameter enough to cross a detection threshold. Visual inspection depends on the relevant component being accessible and visible. These are not deficiencies in inspection practice. They are characteristics of every control that depends on detection.

A well-designed inspection regime standardises the method, defines competence requirements and creates a process for correcting what it finds. In the Netherlands, SCIOS Scope 12 is one example. Other jurisdictions organise comparable functions through electrical inspection, commissioning, maintenance and condition-assessment regimes. Their terminology differs, but their safety objective is similar: identify relevant deficiencies before they develop into harmful events. Limiting the consequences when a relevant deficiency is not identified in time is a different objective.

This paper concerns rooftop PV installed above combustible roof construction. For that situation, I do not ask whether inspection works. I ask what protects the roof when a relevant fault is not prevented, is not detectable under the conditions present, is not detected, or is not corrected before it becomes an ignition source.

In *From risk models to design requirements* I described that question as part of a Defense in Depth architecture, in which controls with different functions and different dependencies complement one another. For the roof, the final layer is physical rather than diagnostic: what happens when heat, arcing or flame reaches the roof?

> **Core principle (author's conclusion)**
>
> **A passive roof-level fire protection measure does not replace inspection. It reduces the extent to which fire safety depends on a relevant defect being detectable, detected and acted upon in time.**

The difference between those two functions is the subject of this paper.

2. What inspection does well

The value of inspection begins with a simple fact: many PV defects that are relevant to fire safety are preventable or correctable if they are identified in time. Incorrect connector assembly, damaged cables, inadequate cable support, loose terminations, deviations from installation requirements and visible signs of thermal or mechanical damage can all be identified with appropriate methods. Where such deficiencies are found and corrected, the probability that they contribute to a later incident is reduced.

This matters most at commissioning, because installation quality recurs as a cause in incident studies. The BRE investigation into UK PV fire incidents attributed approximately 36% of the incidents identified as caused by the PV system to poor installation practice [1]. Large-scale inspection programmes provide a different form of evidence. An international audit programme covering more than 600 commercial rooftop systems identified major safety concerns in 97% of the systems inspected, including grounding defects, damaged modules, mismatched connectors, poor terminations and hotspots [2]. The audited systems are not a representative sample of all rooftop PV installations, but they demonstrate the range of material deficiencies that structured inspection can uncover. A first inspection can identify errors introduced during design, assembly and installation before years of thermal cycling, weather exposure and later maintenance make the original condition harder to reconstruct.

Periodic inspection addresses a different problem. A compliant installation does not remain unchanged: components age, connections experience temperature cycles and mechanical stress, supports deteriorate and later interventions introduce new workmanship. Periodic condition assessment provides information that commissioning alone cannot provide.

The methods used within an inspection programme complement one another. A visual inspection can identify physical damage that leaves no obvious electrical signature. Thermography can locate abnormal heating that cannot be recognised visually. Electrical testing can reveal insulation or continuity problems without an obvious thermal effect. Between site visits, performance monitoring may identify certain operational deviations; its detection mechanisms and limitations are considered in Chapter 6.

A structured inspection regime also does more than generate observations. It turns them into a process. In the Netherlands, SCIOS Scope 12 standardises the inspection method in Technical Document 18, sets competence requirements for inspectors and minimum requirements for measuring equipment, and operates within a certification scheme for inspection companies [3]. An initial inspection is carried out at commissioning or after a substantial modification or extension; a periodic inspection requires the report of that initial inspection, and where no interval is agreed in a policy or contract, the interval is five years. Identified deficiencies must be shown to have been resolved, either through re-inspection or through traceable photographic or documentary evidence supplied by the installer, and an inspection with findings can only be converted into one without findings within twelve months, after which a complete new inspection is required [3, 4]. In this way the chain from observation to corrective action becomes part of the control itself.

For these reasons inspection belongs firmly within a rooftop PV fire-safety strategy. Its effectiveness rests on three conditions:

- the relevant defect has developed sufficiently to become observable;
- the inspection method can detect that type of defect, at that location, under the conditions present;
- the finding leads to timely corrective action.

These are the three later points of failure introduced in Chapter 1, seen from the other side. When all three conditions are satisfied, inspection can interrupt an escalation pathway before a defect becomes an ignition source. The limitations examined in the following chapters begin where one of them is not. A connector may be physically inaccessible. A thermal anomaly may be insufficiently distinct at the available irradiance or operating current. A fault may not yet have developed on the day of inspection, or may arise afterwards. A monitored parameter may remain within its threshold while heat develops locally. None of these possibilities diminishes the value of inspection; together they define its operating envelope.

A sound safety architecture should not require one control to perform a function for which it was never designed. Inspection is designed to identify conditions that can be observed, measured and corrected. It is not a physical barrier between an ignition source and a combustible roof. The question for the next chapter is therefore what can legitimately be concluded from an inspection that finds no relevant defect.

3. What an inspection actually establishes

The three conditions identified in Chapter 2 define the operating envelope of an inspection. For a finding to occur, a defect must have developed sufficiently to be observable, the selected method must be capable of detecting it under the prevailing conditions and the observation must trigger effective corrective action. An inspection report therefore does not assess an installation in the abstract. It produces evidence about an installation within that specific operating envelope.

Infrared thermography provides a particularly clear illustration, because its measurement conditions are explicitly defined. IEC TS 62446-3 specifies outdoor infrared thermography for PV systems in operation. Rather than treating thermography as an open-ended visual scan, the technical specification defines requirements for equipment, environmental and operating conditions, inspection procedures, reporting and personnel competence [5]. For a detailed inspection, the thermographic expert should have a qualification equivalent to Level 2 electro-thermography according to ISO 9712, and at least Level 1 [5, 6]. The method requires interpretation, not simply image capture.

Those requirements reflect the physics of the measurement. Thermography does not measure an electrical defect directly. It measures a surface temperature distribution from which the presence and possible nature of a defect are inferred. For PV module inspection, IEC TS 62446-3 specifies a minimum irradiance of 600 W/m² in the plane of the module. For other electrical components, including cables and connectors, the measured operating current must be at least 30% of the rated system current within the inspected current path, and irradiance above 600 W/m² is recommended [5, Clause 5.3, Table 3]. The specification also limits wind and cloud cover, because they affect the thermal contrast on which interpretation depends.

The importance of operating current can be illustrated by Joule heating. For a resistive contact, electrical dissipation is described by $P = I^2R$. If resistance remains unchanged, operation at 30% of rated current produces $0.30^2 = 0.09$, or 9%, of the resistive dissipation generated at rated current. That does not mean that the observed temperature rise will also be 9%. Surface temperature results from the complete thermal balance, including conduction, convection, geometry, ambient conditions, heat capacity and the thermal state of the component. The calculation demonstrates a narrower point: electrical loading determines the magnitude of the heat source that thermography is trying to observe. Wind can further reduce local temperature rise through convective cooling, changing irradiance can alter the thermal state of the installation during the inspection, and viewing geometry and surrounding materials influence whether an abnormality produces a sufficiently distinct surface signature.

> **MEASUREMENT PRINCIPLE**
>
> **A thermographic inspection does not directly establish whether a relevant electrical defect exists. It establishes whether a thermal abnormality associated with that defect is observable within the inspected scope, under the electrical loading, environmental conditions and viewing geometry present at the time of measurement.**

The same boundary applies to other diagnostic techniques. A visual examination establishes the external condition of components that can be accessed and seen. Electrical measurements establish defined electrical properties under the particular test configuration used. Documentation review establishes whether specified information is present and consistent with the installation assessed. Combining these methods makes an assessment substantially stronger, because each interrogates a different characteristic.

The resulting evidence nevertheless remains bounded by the capabilities of the methods used. IEC TS 62446-3 itself recognises this: a thermal abnormality cannot always be classified unambiguously from thermography alone, in which case further investigation is required [5]. If no abnormality is recorded, the defensible conclusion is therefore that no reportable deviation was identified within the inspected scope and under the recorded conditions. A component may have been inaccessible, the operating condition may not have produced sufficient thermal contrast, a defect may not yet have produced the characteristic being measured, or the component may have been healthy on the day of inspection and deteriorated later. Recording those conditions and limitations is part of what makes an inspection result interpretable, a point I return to in Chapter 8.

The next chapter examines what may remain unseen even when an inspection is competently performed.

4. What may remain unseen

The previous chapter established the measurement boundary in principle. Experimental evidence shows why it matters in practice. Del Prado Santamaría and colleagues examined drone-based infrared thermography on 43 PV modules with deliberately introduced defects, including cell cracks, degraded cell interconnections, potential-induced degradation, glass damage and short-circuited bypass diodes [7]. The modules were operated in strings and imaged at irradiance levels between 200 and 1,000 W/m² and at drone altitudes between 8 and 20 metres.

The results did not show that infrared thermography is ineffective. They showed something more useful for this paper: detectability depended strongly on the type and severity of the defect and on the conditions under which it was observed. Short-circuited bypass diodes and heavy soiling produced pronounced hotspots across a wide range of conditions. Cell cracks were difficult to distinguish consistently from healthy cells, and degraded interconnections produced signatures whose strength depended on their severity. Within the tested range, irradiance had substantially more influence on detectability than drone altitude. Spatial resolution matters only once a component produces a thermal signature to resolve; higher resolution cannot create a signature that the component is not producing.

The study also included modules affected by potential-induced degradation (PID), which did not produce a pronounced thermal abnormality. That finding should be interpreted narrowly. PID is not used here as evidence of a fire-relevant defect, and the induced cases should not be treated as representative of severe field cases. Its relevance is methodological: a defect can be present and independently confirmed without producing a thermal signature that allows reliable identification by thermography alone.

> **DEFECT VERSUS SIGNATURE**
>
> **The physical presence of a defect and the detectability of its thermal signature are two different conditions.**

That distinction becomes particularly important when the component of interest is not the module surface. Connectors, cable joints, junction boxes and sections of DC cabling are often located underneath modules, close to mounting rails or within cable management. Where a component is outside the direct field of view of an infrared camera, its surface temperature cannot be measured. Detection then depends on the defect producing an indirect signature on a visible surface, or on the component being assessed in another way.

Connector faults are not hypothetical. DiGregorio, Burnham and King analysed 6,276 connectors harvested from residential rooftop PV systems in the United States [8]. Among connectors displaying thermal damage, the median measured resistance was 164 mΩ, and 90% exceeded the study's 50 mΩ threshold for critical failure. These figures apply to the thermally damaged subgroup, not to the whole population. The connectors were removed by a single installer during a proactive replacement programme after connector issues had been noticed; the authors note that the systems therefore had known issues and do not represent the industry as a whole, and the study concerns residential rather than commercial roofs. Two

further observations from the same study bear directly on inspection. Two connectors showed complete separation of pin and socket while appearing ordinary from the outside. And the authors note that researchers with site access depend on visual inspection and thermal imaging to assess connectors, which is time-consuming and, depending on cable management, sometimes not feasible.

Two independent boundaries therefore determine whether a defect can be found. The first is **physical accessibility**: can the relevant component itself be seen or measured? The second is **signature formation**: under the conditions present, does the defect produce a characteristic that the chosen method can distinguish from normal variation? Increasing camera resolution does not address the first boundary at all: a component outside the field of view remains outside it. It addresses only part of the second, because resolution can help distinguish a signature but cannot create one.

Inspection practice recognises the accessibility boundary explicitly. Under Scope 12, where cable connections and cable management beneath the modules cannot be seen during the visual inspection, this must be documented in the report with photographs, supplied in advance, that are unambiguous and traceable to the location. The report records why aids such as drones, platforms or endoscopes could not be used, and cable routes that become inaccessible after installation must be inspected during installation [3, 4]. This is a sensible practical response. It also means that, for those components, assurance rests on records made at installation rather than on observation at the time of the periodic inspection. Chapter 7 returns to that same form of assurance for passive roof-level protection.

Different diagnostic methods therefore remain complementary. A defect that produces little thermal contrast may be revealed by an electrical or visual technique; a component hidden from an aerial camera may be accessible during a closer physical examination; and a clearly visible thermal abnormality may still require further testing before its cause is known. The relevant question is not whether a drone, a thermographic camera or an electrical inspection can find defects; each can. It is whether a fire-safety strategy may assume that every defect capable of becoming an ignition source will produce a detectable signature, in an accessible location, under the conditions present, before escalation occurs. The evidence reviewed here does not support that assumption.

No robust universal probability of detection for fire-relevant connector faults has been identified in the literature reviewed for this paper. The residual risk can therefore not be represented as a known percentage of faults that inspection will find. For a roof built with combustible materials, every relevant defect that remains unseen is also a potential ignition source above that roof.

The next question is temporal. Even a defect that would be detectable under suitable conditions must first exist when the inspection is performed.

5. The inspection gap

The previous chapter considered defects that may remain outside the operating envelope of an inspection. A second boundary cannot be removed by better equipment or greater competence: time. A periodic inspection assesses an installation at intervals, while the installation continues to change between them.

It changes in two ways. The first is ageing. Connections experience thermal cycling, moisture and mechanical movement; cable jackets and supports remain exposed to UV and abrasion; and power electronics are often replaced during the life of the modules and the roof. These processes do not run on one clock, and they do not necessarily progress at a constant rate. Field-aged connectors, for example, have shown low and stable contact resistance during earlier years of service, followed in some samples from severe climates by markedly non-linear resistance increases after longer exposure [9].

> **DIFFERENT CLOCKS**
>
> Connectors can remain electrically stable for years and later show non-linear resistance growth, depending on exposure [9]. PV cable jackets degrade under combinations of UV, temperature and mechanical stress, and system life may exceed cable life [10]. A substantial share of inverters fails within fifteen years and is replaced during the life of the installation [11]. These observations are not component lifetimes. They illustrate why the condition of an installation does not evolve at the uniform rate of an inspection calendar.

The second source of change is intervention. Inverters are replaced, arrays extended, modules exchanged and roofs repaired. Each intervention means new work: cables rerouted, connectors disconnected and remated, new terminations made. Laboratory testing has shown that repeated mating cycles alone can measurably increase connector contact resistance, even without environmental ageing [12]. A system may therefore be fifteen years old while a connection made during an inverter replacement is six months old. This distinction between **component age** and **workmanship age** means that a PV installation no longer has one meaningful age, and that the condition assessed at a previous inspection may cease to describe the installation after later intervention.

> **THE INSPECTION GAP**
>
> **A satisfactory inspection describes an assessed condition. It does not preserve that condition until the next inspection.**

None of this represents an inspection failure. A defect that develops after an inspection did not exist when the inspection was performed. Continuous monitoring can narrow part of this gap where a developing fault changes a monitored variable; Chapter 6 examines how far that reaches. Inspection linked to events can address another part. Commissioning, substantial extension, inverter replacement, significant electrical repair and major roof work all create a new condition that deserves reassessment rather than waiting for the next calendar date, a distinction I return to in Chapter 8.

Events of this kind also matter for the roof itself. Where the fire strategy includes a passive protective layer, significant roof or PV work affects two systems at once: the electrical installation may need reassessment, and the protective layer must be reinstated and its continuity recorded. A passive layer does not depend on detecting an electrical fault, but it does depend on correct installation and on remaining intact after later work.

The supported conclusion is narrow. The available evidence does not show that older PV systems are inherently unsafe, or that a particular inspection interval is too long. It shows that the state of a PV installation changes on several timescales while inspection samples that state at discrete moments. On a combustible roof, whatever develops in between does so above fuel that no inspection report protects.

6. Detection between inspections

The inspection gap raises an obvious question: if periodic inspection samples the installation only at discrete moments, can continuous monitoring remove the uncertainty between them? It can remove part of it. Inverters supervise electrical parameters, some systems compare string performance, insulation monitoring can detect certain deterioration towards earth, and arc-fault detection can recognise specified signatures associated with arcing. These controls do not wait for the next site visit. Their effectiveness nevertheless depends on the developing fault altering a monitored quantity sufficiently to satisfy the relevant detection logic.

Detection requires a measurable consequence

A developing fault does not announce itself simply because it is hazardous; it must create an electrical, thermal or other measurable characteristic. This is particularly relevant for local resistive faults. A deteriorating connector can dissipate significant heat at a very small location while the corresponding power loss remains small relative to the power of the complete string. Using the median resistance reported for thermally damaged connectors in the DiGregorio study as an illustration, a resistance of 164 mΩ at 10 A dissipates approximately 16 W locally. That calculation does not establish an ignition threshold and should not be read as one. It shows that meaningful local heating can represent only a small fraction of the output of a multi-kilowatt string. Whether such a condition is visible in inverter or string-level monitoring depends on measurement architecture, resolution, thresholds and comparison method, and no probability that conventional performance monitoring will detect a single degrading connector has been identified for this paper.

Thresholds are part of the safety function

Insulation monitoring provides a better documented example. PV inverters can supervise insulation resistance between the DC circuit and earth, so that deteriorating cable insulation or another ground-related fault produces a quantity that is monitored continuously. Research by Sandia National Laboratories shows that the trip threshold involves a genuine engineering trade-off [13]. It must sit above the normal leakage of the array to avoid unwanted tripping, yet low enough to detect hazardous faults. The authors concluded that without specific knowledge of the array it is impossible to define a threshold that both provides fire safety and avoids unwanted tripping: a default setting will either cause nuisance trips in leaky systems or fail to trip before a fire hazard in well-isolated ones. In their field tests, faults located inside the array produced readings above the trip threshold and could have persisted. Detection is therefore not binary: the threshold is part of the safety function.

Detection and interruption are different controls

Arc-fault circuit interrupters go one step further: they are designed to recognise specified signatures associated with arcing and to interrupt the circuit automatically. The regulatory position differs between markets; in the United States the National Electrical Code requires DC arc-fault protection for defined PV applications [14], while equivalent protection is not universally required in Europe. The DC-arc papers in this series examine arc detection in more detail. For

the present argument one distinction is sufficient: AFCI protection removes the human-response step for the faults it detects, but it still depends on the event producing a signature within the detection capability of the device.

The dependency chain

Inspection and monitoring can now be placed in one framework, which expands the three conditions of Chapter 2. For periodic inspection to prevent escalation, a defect must exist, become observable, fall within the inspected scope, be recognised, be correctly interpreted, be reported and be corrected in time. For continuous monitoring the sequence is shorter but still conditional: the fault must develop, the monitored variable must change, the threshold or algorithm must be satisfied, and an alarm or automatic response must follow. Where the response is not automatic, the chain also includes people. An alarm may reach an owner, a remote monitoring service or an O&M contractor, and still requires timely interpretation and action. This is one of the ways in which operational risk can drift after commissioning, as discussed in *The economics of the rooftop PV decision*.

Multiple controls are not automatically independent merely because they use different technologies. Thermography, visual inspection, performance monitoring and automatic fault detection may all depend, in different ways, on the same developing defect first producing an observable consequence. A hidden or weakly expressed fault can therefore challenge several detection-based controls at once. The purpose here is not to calculate a formal LOPA credit for inspection or monitoring; the evidence does not support that level of quantification. It is to identify where these controls derive their protective function and where their dependencies are shared.

> **DEPENDENCY PRINCIPLE**
>
> **Detection-based controls differ in method, speed and response, but they share one fundamental dependency: the relevant fault must produce a detectable condition before escalation.**

A passive roof-level barrier has a different dependency structure, as described in *From risk models to design requirements*. It does not detect an abnormal temperature, resistance, insulation condition or arc signature, and it does not wait for a threshold to be exceeded. Its function begins later in the sequence, when heat, arcing or flame has already reached the roof. That does not make it superior to inspection, monitoring or automatic interruption. It must still be correctly specified, installed and kept intact through later work, and those are real failure modes with their own quality assurance. But once present and intact, its fire-safety function does not depend on the initiating fault first becoming observable. The earlier controls seek to prevent or interrupt the event; the passive layer addresses what happens when they do not.

7. Consequence-limiting design at roof level

The preceding chapters considered controls that depend, in different ways, on identifying a developing fault before it becomes an ignition source. A passive roof-level fire-protection measure performs a different function. It does not attempt to determine why heat, arcing or flame is present. Its purpose is to change what happens when that thermal exposure reaches the roof construction.

The physical basis for that function is addressed in *Beyond flame spread*, including the effect of PV geometry on heat transfer and the role of the roof build-up beneath the membrane. What a particular fire test does and does not demonstrate is the subject of *What does a fire test actually prove?* Neither is repeated here. For the present paper one point is sufficient: consequence limitation does not depend on successful diagnosis of the initiating electrical fault. A connector can remain inaccessible to inspection, a developing fault can fail to produce a distinct thermal signature, monitoring can remain below its threshold and a defect can arise after the most recent inspection. Those conditions affect the probability that the electrical problem is identified in time, but they do not change the intended function of a correctly specified and installed passive layer.

The evidence required for such a layer is therefore different from the evidence required for inspection. An inspection method is assessed by what it can observe, under which conditions and with what process for corrective action. A passive fire-protection measure must demonstrate performance under relevant thermal and fire exposure, and that evidence remains construction-specific: a successful test does not automatically establish equivalent performance for every roof build-up, detail, exposure duration or installation condition.

Assurance moves from detection to installation

A passive layer removes some dependencies from the electrical fault sequence but introduces others. Its performance depends on correct specification, continuity, correct installation and compatibility with the roof construction. Penetrations, later repairs or removal during roof work can compromise that continuity. Quality assurance is therefore concentrated at installation and whenever the roof is subsequently disturbed. Electrical inspection asks whether a relevant defect can be identified and corrected; passive protection asks whether the intended protection is still present and continuous when an ignition source occurs. Both questions remain relevant.

This is particularly important after intervention. As Chapter 5 showed, inverter replacement, array modification, module replacement and roof repair introduce new electrical workmanship, and some of the same events can disturb a roof-level protective layer. Event-driven assurance should therefore consider both: whether the electrical installation has changed, and whether the physical protection of the roof has been affected. The location, extent and detailing of a passive layer should be documented at installation, and any later work affecting it should end with documented reinstatement.

This assurance has a limit that should be stated plainly. A layer beneath the roof covering cannot be inspected periodically without opening the roof. If later work disturbs it and the

reinstatement is not recorded, its condition cannot readily be verified. Its dependability therefore rests on records and change control rather than on recurring observation, which places responsibility with whoever manages work on the roof. In that respect it is assured in the same way as connectors concealed beneath modules are under Scope 12: through records made at installation, not through later observation.

The scope of the protection must also remain clear. A roof-level layer protects the area and construction in which it is correctly installed. Penetrations, edges, expansion joints, rooflights and other discontinuities can form separate pathways and require their own treatment, as discussed in *From risk models to design requirements*.

The false-security argument

There is a legitimate objection to passive protection. If owners, installers or insurers interpret it as a reason to reduce inspection, maintenance or fault detection, the additional layer could create behavioural risk rather than reduce it. That interpretation would contradict the Defense in Depth principle, because replacing one layer with another reduces depth rather than increasing it. A passive layer does not make defective connectors acceptable, damaged cables irrelevant or competent installation unnecessary. Electrical faults can still damage equipment, reduce performance, create electrical hazards and cause fire elsewhere in the installation.

> **Consequence-limiting principle (author's conclusion)**
>
> **Passive roof-level protection is not a substitute for inspection. It addresses the residual case in which prevention, detection or interruption does not occur in time, through a function that does not depend on detecting the fault.**

The strongest safety architecture therefore retains both functions. Inspection and monitoring reduce the probability that developing electrical defects reach ignition; passive roof-level protection limits the consequence if that pathway is nevertheless completed. Chapter 8 translates this into practical questions.

8. What this means in practice

The preceding chapters distinguish three forms of assurance. Inspection establishes information about the condition that could be assessed with the methods, access and operating conditions available at the time. Monitoring extends observation between inspections but remains dependent on measurable signals and thresholds. Passive roof-level protection addresses a later point in the sequence and requires assurance of specification, installation, continuity and later intervention. The practical task is not to choose between these controls, but to prevent one control from being credited with the function of another.

Calendar-driven and event-driven assessment

A calendar-driven inspection asks whether enough time has passed since the previous assessment. An event-driven assessment asks whether something has happened that changes the condition previously assessed. Both are needed. Under Scope 12, commissioning and substantial modification or extension already require an initial inspection, and cable routes that will become inaccessible must be inspected during installation [3, 4], which corresponds to the hold points before concealment recommended in *Rooftop PV on new commercial buildings*. Inverter replacement, significant electrical repair and major roof work introduce new connections and workmanship in the same way. In the sources reviewed for this paper it is not defined whether such events count as a substantial modification; that should be made explicit rather than assumed. The same events are also the moments at which a passive roof-level layer may be disturbed and must be reinstated.

Where a roof relies on a passive separation layer, the inspection scope can usefully include the parts of that layer that remain visible or accessible: joints and overlaps at edges and openings, mechanical fixings, cable entries and other penetrations, local damage after maintenance or repairs, and the actual distance between connectors and the roof covering. Because such a layer lies beneath the membrane over most of the roof, most of its assurance still has to come from installation records and change control rather than from later inspection. These are attention points for the inspection scope, not requirements set by a product manufacturer.

Practical questions

For	Question	Why it matters
Owner	What was actually assessed, under which conditions, and what was inaccessible?	A satisfactory report applies to the assessed scope and conditions, not automatically to every component on the roof.
Owner	What has changed since the inspection: inverter replacement, repair, extension or roof work?	Later intervention can change the assessed condition and can disturb passive protection.
Owner	Who maintains the change record for both the PV installation and the roof?	Event-driven assurance depends on work being recorded and followed up.
Insurer or risk engineer	Is the report treated as evidence of an assessed condition, or implicitly as a continuing guarantee of fire safety?	Those are different claims.
Insurer or risk engineer	Are several detection-based measures credited as if they were independent?	Different technologies can share a dependency on the same fault first becoming

For	Question	Why it matters
		observable.
Insurer or risk engineer	If passive protection forms part of the strategy, how are installation, coverage and reinstatement documented?	Its assurance depends on specification and change control rather than on detection.
Inspector	Are measurement conditions and inaccessible areas recorded in the report?	They define the evidence boundary of the inspection.
Inspector	Does the wording distinguish “no abnormality detected” from “no relevant defect exists”?	“Not observed” should not silently become “not present”.
Inspector	Are changes since the previous inspection identified?	New workmanship creates a new condition to assess.
Roof or PV contractor	Has the work introduced new connectors, terminations or cable routes?	The affected parts may require reassessment.
Roof or PV contractor	Has the protective layer been penetrated or disturbed, and has its reinstatement been recorded?	A concealed layer cannot readily be verified afterwards.
Roof or PV contractor	Are penetrations, rooflights, edges and joints treated separately?	Protection of the main roof area does not control every route through the roof.

Table 1. Practical questions for those who rely on, perform or affect inspection and roof-level protection.

The practical framework can be reduced to four questions:

1. **What was assessed?** Which components, locations and conditions were actually examined?
2. **What has changed since then?** Has ageing, repair, modification or roof work altered the assessed condition?
3. **What depends on detection?** Which parts of the strategy require a developing fault to become observable and to trigger a response?
4. **What limits the consequence if detection does not occur in time?** What protects the combustible roof when the fault is missed, develops later or falls outside the detection envelope?

Those four questions keep the functions separate without placing them in competition. The objective is not less inspection. It is to use inspection for what it proves, monitoring for what it can detect, and passive protection for the consequence-limiting function it is intended to perform.

9. What we still do not know

The evidence reviewed in this paper supports a clear distinction between detecting a developing fault and limiting its consequences. It does not support precise claims about how often detection succeeds, how quickly every relevant defect develops, or how much residual fire risk remains after a satisfactory inspection. Those uncertainties should remain visible.

The probability of detection is not known

No robust universal probability of detection has been identified for fire-relevant defects in rooftop PV systems. Studies demonstrate that particular defects are detectable under defined conditions, and some report high agreement between techniques for selected module-level abnormalities. That is not the same as knowing the probability that a developing connector, cable or other fire-relevant fault will be detected before it becomes an ignition source. Chapters 3 and 4 show why a single percentage would be difficult to establish: detectability varies with defect type, severity, accessibility, electrical loading, environmental conditions, method and the stage at which the defect is examined. The correct status for a general probability of detection is therefore unknown, not low and not high.

Fire statistics are not linked to inspection history

No incident dataset identified for this paper systematically links a PV fire to the inspection history of the installation: whether it had been inspected, by which method, when, with what recorded limitations, and whether it had since been modified. That missing link prevents two opposite claims from being tested. The available statistics cannot show how often an installation with a satisfactory inspection later develops a fire, but neither can they show that this is exceptionally rare. A recent comparison of national PV fire statistics in the United Kingdom, Italy, Sweden and Slovenia found substantial differences in reporting methods and could not establish the age of the affected systems from the available data [15]. Incident reporting would be considerably more useful if it recorded system age, the date and result of the most recent inspection, later modifications, and whether monitoring or automatic protection had operated.

The field progression rate of connector faults is uncertain

Field-recovered connector studies demonstrate that severe resistance increases and thermal damage occur in operating installations, and accelerated tests demonstrate mechanisms through which connector condition deteriorates. What they do not yet provide is a reliable clock. A laboratory observation that a configuration can deteriorate rapidly under defined loading, movement or ageing should not be converted into a prediction for field connectors, because current profile, temperature cycling, product design, workmanship, moisture and mechanical loading differ substantially between laboratories and operating roofs. This paper therefore does not assign a general period from initial degradation to hazardous heating or arcing. That interval remains an important research need, because it determines how a given inspection frequency relates to different failure mechanisms.

A possible masking effect after thermal damage

In the study of harvested rooftop connectors, connectors showing early thermal damage had median resistances of approximately 196–199 mΩ, while fully melted connectors had a lower median of 63 mΩ, some below 0.2 mΩ. The authors suggest that arcing and local melting may sometimes fuse pin and socket together, re-establishing a low-resistance connection; they also note that connector resistance is not static and can fall when higher current is applied [8]. Whether such a connector would present a less conspicuous signature in a later inspection has not been tested. A useful programme would follow connectors through degradation, thermal event, cooling and subsequent electrical and thermographic assessment.

The limits of continuous monitoring are not quantified

Chapter 6 showed why monitoring has a temporal advantage over periodic inspection and why threshold-based detection involves trade-offs. What remains unknown is the probability that inverter, string-level or other monitoring will identify a single developing high-resistance connection before hazardous local heating occurs. That local heat loss is a small proportion of string power does not establish that monitoring will miss the fault; it shows why system-level performance and local thermal severity need not scale together. Defect-specific detection studies are required before performance monitoring can be assigned a quantitative risk-reduction factor for connector-related fire hazards.

Passive protection also has open questions

The absence of a detection dependency does not remove uncertainty from passive protection. Its performance must be established for the relevant roof construction, exposure and detailing; its continuity after later work depends on records and change control; and penetrations, edges, rooflights and other discontinuities remain separate pathways. Long-term field evidence is also more limited than short-duration fire-test evidence, and, as noted in *Beyond flame spread*, ageing remains largely outside the available fire-test evidence. These uncertainties differ from those of inspection and monitoring, but they deserve the same evidence discipline. A passive measure should be credited only for the performance supported by its tested construction, installation details and lifecycle assurance.

Conclusion

Inspection works by creating knowledge about the condition of a PV installation and using that knowledge to initiate corrective action. Thermography, visual inspection, electrical testing, monitoring and automatic protection each extend that capability in different ways. The evidence also defines their boundary. A relevant defect must exist, become observable, fall within the capability and scope of the method and lead to timely action before the escalation pathway is interrupted. Some of these dependencies can be reduced through better inspection, continuous monitoring and automatic protection, but they cannot all be removed by adding more detection methods, because those methods can share the same starting condition.

The installation also continues to change. Components age, later work introduces new connections and workmanship, and faults can develop after a satisfactory inspection. No universal probability of detection for fire-relevant rooftop PV defects has been established, and fire statistics do not connect later incidents to the inspection history of the installation. That uncertainty should not be turned into an argument against inspection. It should determine which conclusions an inspection report is allowed to support.

A passive roof-level fire-protection measure addresses a different point in the sequence. Its intended function does not require the initiating defect to be identified first. Its own dependability rests instead on specification, tested performance, correct installation, documentation and change control. Neither function replaces the other.

> **FINAL STATEMENT**
>
> **For rooftop PV above combustible roof construction, the question is not only whether the installation has been inspected. It is whether the safety strategy still has a functioning layer when a relevant fault is not prevented, not detected or not interrupted in time.**

References

[1] BRE National Solar Centre (2018). *Fire and Solar PV Systems – Investigations and Evidence*. Report P100874-1004, Issue 2.9.

[2] Intertek Clean Energy Associates (2024). *Resolving safety concerns and keeping people safe at solar installations*. February 2024.

[3] SCIOS (current versions). *Technisch Document 18 (TD18): Inspectie van zonnestroominstallaties* (SCIOS Scope 12) and *SCIOS Certificatieregeling*.

[4] RJ Inspecties B.V. (2025). *Scios Scope 12* (versie 2025). Description of the Scope 12 inspection procedure under TD18, including initial and periodic inspection, inspection interval, non-visible installation parts, repair evidence and registration.

[5] IEC (2017). *IEC TS 62446-3:2017 Photovoltaic (PV) systems – Requirements for testing, documentation and maintenance – Part 3: Photovoltaic modules and plants – Outdoor infrared thermography*. Geneva: IEC.

[6] ISO. *ISO 9712: Non-destructive testing – Qualification and certification of NDT personnel*. Geneva: ISO.

[7] Del Prado Santamaría, R., Alves dos Reis Benatto, G., Dhimish, M., Koc, T., Friansyah, R., Kari, T., Mahmood, A., Poulsen, P. B. & Spataru, S. V. (2025). Influence of irradiance and drone altitude in infrared thermography inspections of photovoltaic plants. *EPJ Photovoltaics*, 16, 31. doi:10.1051/epjpv/2025020.

[8] DiGregorio, S. J., Burnham, L. & King, B. H. (2025). Rapid characterization and failure analysis of 6276 rooftop-harvested photovoltaic connectors. *Solar Energy*, 301, 113916. doi:10.1016/j.solener.2025.113916.

[9] Kapoor, K., Muthusamy, E., Goyal, D. & Nikhil, P. G. (2025). Study on reliability of field-aged photovoltaic connectors. *Eksploatacja i Niezawodność – Maintenance and Reliability*, 27(3), 199496. doi:10.17531/ein/199496.

[10] Miller, D., Arnold, R., Hacke, P. et al. (2025). Photovoltaic cable jackets – A comparison of representative products using combined-accelerated stress testing. *Solar Energy Materials and Solar Cells*, 282, 113372. doi:10.1016/j.solmat.2024.113372.

[11] Bern University of Applied Sciences (BFH) (2022). Life expectancy of PV inverters and optimizers in residential PV systems. *8th World Conference on Photovoltaic Energy Conversion (WCPEC-8)*, paper 3DV.1.46.

[12] Rogowski, S. et al. (2026). Degradation of photovoltaic DC connectors under real switching conditions. *Energies*, 19(10), 2332.

[13] Flicker, J., Johnson, J., Albers, M. & Ball, G. (2015). *Recommendations for Isolation Monitor Ground Fault Detectors on Residential and Utility-Scale PV Systems*. SAND2015-4667C. Albuquerque: Sandia National Laboratories.

[14] NFPA. *NFPA 70: National Electrical Code*, Article 690.11, Arc-Fault Circuit Protection (Direct Current). Quincy, MA: NFPA.

[15] Rus, N. et al. (2026). Fire incidents involving photovoltaic systems – An analysis of different national statistics. *Fire Safety Journal*.

Documents referred to in the author's perspective

ZAG, Slovenian National Building and Civil Engineering Institute (2025). *Report 448/25-560-1-EN: About performance comparison of lightweight non-combustible sheet in rooftop PV-related fires*. Responsible investigator: Kirils Simakovs. 1 August 2025. Orderer: AllShield Holding B.V.

Other papers in the series cited in this paper: *From risk models to design requirements*; *Beyond flame spread*; *What does a fire test actually prove?*; *The economics of the rooftop PV decision*; *Rooftop PV on new commercial buildings*.

Method and sources

This is a targeted inventory and synthesis of identified knowledge, updated through the checks documented for 1 November 2026. It does not claim exhaustive coverage of the latest literature. Source findings, author analysis and practical design proposals remain distinct. The author endorses the conclusions within these evidence boundaries.

Authorship and contributions

Olivier Langejan is the lead author, supported by a small team. His professional focus is building and roof design. The papers synthesise the cited research, standards and test evidence; specialist fire-science and electrical findings are attributed to those sources rather than presented as expertise personally established across every discipline. The author endorses the synthesis and its stated limits. AI assistance is described separately below.

Key source boundaries

FM's public Power & Renewables release of 5 October 2026 supports early engineering input while site, technology and design decisions can still be changed. Its survey concerns the wider energy sector and does not establish rooftop-PV incident or performance data. Source: FM, "With Power and Renewables Risks Evolving, Engineering Must Move Upstream", 5 October 2026, https://www.fm.com/newsroom/news/2026/fm-power-and-renewables-press-release.

Scope and selection

This paper synthesises the detectability and time boundaries of satisfactory rooftop PV inspection. Inspection standards, experimental detection studies, field observations and incident-reporting evidence are included where relevant to residual fire risk. Detailed roof fire-performance comparisons are addressed in companion papers.

Sources and dates

The evidence base is the research, standards, approval documents, technical guidance and incident material identified in the references. This is a targeted narrative synthesis, not a systematic review. English-language sources and relevant Dutch or German material are included where cited. The original search dates and a complete search log are not available in the retained publication record, so no reconstructed search period or exhaustive publication cut-off is asserted. Sources retain their individual publication dates. The present editorial and targeted source checks were completed on 1 November 2026.

Publication range

Publication dates and editions are identified for the individual references. This targeted synthesis does not claim exhaustive chronological coverage or a reconstructed complete search period.

Verification

This revision checked the requested wording against the source papers and carried out targeted primary-source checks for the additions and source corrections identified in its revision history. It is not a new claim-by-claim audit of the entire paper. References identify primary documents where available; secondary accounts and access limitations are identified where material. Author interpretations and design proposals do not acquire test or approval status.

Known gaps

No universal probability of detecting all fire-relevant defects is established. The complete current TD18 text was not newly obtained for this revision; details attributed to a commercial inspector remain secondary-source descriptions.

AI-assisted research and writing

Artificial intelligence was used to help identify and locate technical literature, compare findings across sources, analyse evidence boundaries, organise the argument and assist with drafting and editing. AI-generated summaries were not treated as technical evidence. The selection of sources, interpretation of the evidence, conclusions and final wording remain the responsibility of the author.

Updates

New research, test results, incidents and regulation may lead to new versions. Editorial corrections and technical additions are identified separately in the revision history. Missing sources and evidence that challenges a conclusion are welcome.

REVISION HISTORY

Version	Date	Changes
1.0	21 October 2026	First published version. Added the continuity of a passive separation layer as an attention point for the inspection scope. Revision history moved to the back of the paper.
1.01	1 November 2026	Editorial corrections: reframed as a roof-first knowledge synthesis; shortened the author perspective, added method and source limitations, labelled author interpretations and corrected the requested wording and transparency statements. Technical additions and source checks: No new inspection-performance data. Secondary descriptions of inaccessible Scope 12 documentation remain identified as an access limitation.
1.02	1 November 2026	Method and sources: clarified targeted synthesis; FM release of 5 October 2026 used only for early engineering, not rooftop-PV survey or handover performance claims. Existing inspection uncertainty and Scope 12 source-access limits retained.
1.03	1 November 2026	Authorship and contributions added; source boundaries and expertise scope clarified. Existing technical conclusions and the outstanding TD18 access limitation retained.

THE SERIES

The roof comes first is a synthesis of knowledge on fire safety and rooftop PV from a roof-first perspective, by Olivier Langejan. Each paper stands on its own; together they follow one line of thought.

At the time of this version, the series consisted of:

Paper	For	In one line
Rooftop PV on existing commercial buildings	Asset owners, developers, project teams	A decision framework for resolving roof constraints before the PV investment is committed.
Rooftop PV on new commercial buildings	Developers, architects, project teams	Designing the roof and PV as one system from the start.
The economics of the rooftop PV decision	Asset owners, CFOs, risk managers	Probability, consequence and lifecycle cost from the asset owner's perspective, including Business Interruption.
From risk models to design requirements	Risk engineers, insurers, designers	An application of five established risk frameworks to consequence-limiting design for roofs with PV.
Beyond flame spread	Risk engineers, insurers, fire engineers	How far a rooftop PV fire can travel into the roof before it becomes a building loss, and what a separation layer changes.
What does a fire test actually prove?	Insurers, risk engineers, specifiers	What a fire test result does and does not demonstrate about a real roof with PV.
DC arcs as a heat source	Insurers, risk engineers, PV designers	How DC arc faults start, persist and end in rooftop PV.
DC arcs at the roof-PV interface	Insurers, risk engineers, PV and roof designers	What determines roof exposure to a sustained DC arc and what current evidence shows about measures to limit it.
How much does PV clearance actually solve?	Insurers, risk engineers, PV designers, asset owners	What module-to-roof clearance does and does not demonstrate about the roof beneath it.
What if the PV panel could not burn?	Asset owners, risk engineers, insurers, PV designers	What a lower-combustibility PV module does and does not change for the roof beneath it.
What does an inspection actually prove? (this paper)	Insurers, risk engineers, asset owners, inspectors	What inspection can establish, what it cannot and what still protects the roof when a fault is missed.

All papers and their latest versions: allshield.com/en/olivier-langejan

Comments, corrections, missing sources, evidence that contradicts a conclusion and field experience are welcome at ol@allshield.com. Relevant input is used in future versions. To be notified of new papers and versions, send a short email to the same address.

This paper may be shared freely, unchanged and with attribution.