THE ROOF COMES FIRST · WHITEPAPER

How much does PV clearance actually solve?

Separating ignition prevention, exposure mitigation and consequence limitation on commercial flat roofs

CENTRAL QUESTION

Increasing the gap between a PV array and the roof changes fire behaviour. How much of the roof's fire risk does that actually solve?

Olivier Langejan
Chief Commercial Officer, AllShield
ol@allshield.com | allshield.com/en/olivier-langejan

Prepared for insurers, risk engineers, fire engineers, PV designers, roofing professionals and asset owners.

Version 1.03 | 1 November 2026

Contents

Author's perspective....3
Executive summary....4
1. Why clearance matters....6
2. Is there a safe clearance?....9
3. What clearance does not change....13
4. The roof does not see clearance in isolation....15
5. What if fire exposure still reaches the roof?....19
6. Three layers, three questions....22
7. A roof-first design sequence for clearance....25
8. What we do not yet know....29
Conclusion....33
References....34
Method and sources....36
REVISION HISTORY....38
THE SERIES....39

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: how much of a commercial roof's fire risk does the clearance between PV modules and the roof actually address? I separate ignition prevention, exposure mitigation and consequence limitation and place clearance within that model. This paper was written as part of my role as Chief Commercial Officer at AllShield. It contains no product case study; the argument in Chapters 1–8 is product-neutral.

Several relationships are relevant to this subject and are stated here. AllShield commissioned roof-PV fire testing at ZAG. In the tested configurations, ZAG concluded that the protective layer prevented direct flame penetration into deeper structural layers, while thermal transfer still produced subsurface damage, including PIR degradation and EPS melting, and localised burn-through of the sheet was recorded in two runs. Grunde Jomaas, who is affiliated with FRISSBE at ZAG, co-authored several of the studies cited here [1–4,19], and the ZAG test method [4] was co-authored by Kirils Simakovs, the responsible investigator for the AllShield-commissioned programme. AllShield BarrierSheet is also FM Approved under FM 4470 and FM 4478 in a defined roof-PV assembly, while this paper discusses FM DS 1-15 as an institutional evidence framework. The documents behind these statements are listed after the references. Where independent published research addresses the same mechanism, I have relied on that research for the argument.

None of the conclusions depends on an AllShield product being selected. Clearance can be a legitimate and substantial exposure-mitigation measure, including as a prescribed route within FM's framework, and roof intervention may be unnecessary where the existing construction, geometry and evidence already provide an acceptable risk position.

Olivier Langejan
Chief Commercial Officer, AllShield

Executive summary

Clearance between rooftop photovoltaic modules and the roof beneath them is increasingly discussed as a way to reduce PV fire risk. The physical basis is sound. A PV module creates a semi-enclosed cavity above the roof and changing that cavity changes flame shape, heat transfer and the conditions under which a local fire can develop into sustained propagation. Experimental research has demonstrated that relatively small changes in gap height can, in specific configurations, produce large changes in fire behaviour. [1,2]

The difficulty begins when an experimentally observed clearance is converted into a general design rule. Published studies do not identify one universally safe distance. Critical regions change with the ignition source, module inclination, frame geometry, roof construction and the definition of the distance being measured. A gap measured to the rear surface of a module is not necessarily the same as the distance to the lower edge of its frame, and neither is necessarily the distance between a connector or cable fault and the roof. Clearance therefore has technical meaning only when the geometry and fire scenario behind the number are understood. [2,3]

Institutional guidance approaches the same issue from another direction. FM DS 1-15 provides a prescriptive route for existing roof covers that will remain in service, in which PV panels are installed at least 0.9 m above the roof “to eliminate the fire hazard”. [7, §2.1.1.2(E), p. 5] Within the FM property-loss-prevention framework, that is a legitimate risk-control route. It remains part of a defined framework with associated wind requirements rather than establishing 0.9 m as a universal physical fire threshold.

This building interaction is fundamental. Clearance is not merely a fire variable: changing PV height and array geometry can also alter wind pressures, support reactions and snow distribution [7,12–16], as well as waterproofing interfaces, maintenance access, material demand and the consequences of future roof replacement. On an existing commercial roof, all of those effects have to fit within a structural and building envelope that already exists. The design objective is therefore not to maximise clearance in isolation but to establish which geometry can be justified across fire, structure, weatherproofing and lifecycle requirements.

The same distinction applies to the fire strategy itself. This paper separates three complementary functions. **Ignition prevention** acts on the probability that an electrical or thermal fault develops into a fire. **Exposure mitigation** acts on the pathway between that source and the roof; clearance belongs primarily in this second layer. **Consequence limitation** addresses what happens when significant thermal exposure nevertheless reaches the roof assembly. The waterproofing, insulation, protective layers, interfaces and supporting construction then determine whether an incident remains local or develops into a larger roof or building loss.

These functions are related but not interchangeable. Better electrical design does not prove that the roof can tolerate fire exposure. A favourable clearance does not prevent every electrical fault and does not alter the materials within the roof. A robust receiving roof does not justify weak electrical installation. The purpose of defence in depth is not to select one of these approaches but to make clear what each measure is expected to accomplish.

This distinction matters most on existing buildings. Before clearance is selected, the actual roof has to be known, including retained insulation, previous membranes and repairs that may be absent from the original drawings. Its condition and remaining life also change the available routes. Where an existing roof is dry, serviceable and capable of remaining in place, a geometry-led route may have the genuine advantage of avoiding intervention in the roof build-up. Where refurbishment or replacement is already required, the project has an opportunity to reconsider consequence limitation within the renewed

assembly. FM DS 1-15 makes a similar distinction between roof covers that remain and roofs that are being replaced. [7, §§2.1.1.1–2.1.1.2, pp. 4–5]

The roof-first gate sequence of this series therefore places clearance after the building has been understood. The actual roof, the consequence that matters, the structural envelope and the condition of the roof come first; mounting geometry and the three defence layers follow. Evidence is then matched to the function being claimed, the geometry is verified as built and the strategy is reassessed when either the PV system or the roof changes.

The evidence also shows where uncertainty remains. Published fire and wind studies generally do not investigate the same progressive clearance configurations. Most critical-clearance fire experiments use reproducible non-electrical ignition sources rather than real DC arcs. There is no common convention for where fire-related PV clearance is measured, and long-term changes such as accumulated debris and displaced cables are not reflected in the experimental data. Current external roof classifications do not, by themselves, reproduce the semi-enclosed fire environment created by PV. [8]

None of these limitations means that clearance should be rejected. They mean that the claim made for clearance should remain proportional to the evidence behind it. An experimental critical gap belongs to the configuration in which it was measured. A prescriptive insurer value belongs to the framework that accepts it. A roof classification belongs to its documented field of application. A complete roof-PV approval provides a different and broader form of evidence.

DESIGN QUESTION

Enough for what fire scenario, on what roof and demonstrated by what evidence?

Clearance can materially reduce the pathway from a PV fire towards roof involvement and should be treated as an important exposure-mitigation measure. But the design decision starts with the roof, not with a preferred distance above it. The objective is a roof-PV system in which ignition prevention, exposure mitigation and consequence limitation are each supported by evidence appropriate to the actual building.

1. Why clearance matters

The distance between a rooftop PV module and the roof beneath it is not a minor installation detail. It changes the geometry in which a fire develops and can therefore influence the heat transferred towards the roof, the direction in which flames travel and the conditions under which local burning develops into sustained flame spread.

Without a PV module above it, a flame on a roof can develop relatively freely. Installing a module creates a semi-enclosed space in which flames and hot combustion gases interact with the module above and the roof below. Radiation from the flame and heated surfaces can increase thermal exposure ahead of the fire while the module can redirect flames and hot gases along the roof surface. The height, width, inclination and internal geometry of this cavity therefore become relevant fire-safety parameters.

Fundamental evidence for the effect of clearance

Kristensen, Faudzi and Jomaas investigated the fundamental effect of gap height using polymethyl methacrylate (PMMA) as a surrogate combustible roof surface and a black stainless-steel plate as a surrogate PV module. Their objective was not to reproduce a complete commercial roof assembly but to isolate the physical mechanisms governing horizontal flame spread beneath an obstruction. For a 300 mm wide PMMA sample, reducing the gap from 200 mm to 170 mm increased the measured flame-spread rate from approximately 0.37 mm/s to 2.41 mm/s. Similar transition behaviour was identified at other tested sample widths. [1]

The significance of this experiment is not that 170 mm or 200 mm can be transferred to a commercial roof as design limits. Its importance is the demonstration that relatively small changes in cavity geometry can produce disproportionately large changes in flame behaviour. Clearance is therefore a genuine fire-dynamics parameter rather than merely a practical mounting dimension.

The simplified experimental arrangement is equally important to understanding what was demonstrated. PMMA was deliberately selected as a reproducible fuel and the steel plate provided a simplified representation of the obstruction above it. The numerical transition values belong to that controlled configuration. The mechanism they reveal is much more general than the numbers themselves.

Moving from a surrogate to a roof construction

Kristensen, Jacobs and Jomaas subsequently moved closer to an actual roof system by conducting experiments on a B_{ROOF}(t4)-compliant flat-roof mock-up. The study varied gap height, module inclination, underlying insulation and whether the obstruction above the roof consisted of a stainless-steel plate or an actual PV module. [2]

For horizontal configurations, the experiments showed a transition between conditions in which flame spread remained within the ignition-source domain and conditions in which sustained propagation developed beyond it. For the horizontal stainless-steel plate, the transition region lay between approximately 100 and 110 mm: no propagation outside the source domain was observed at 120 mm, consistent propagation occurred at 100 mm and the 110 mm tests showed scatter within the transition region. With actual horizontal PV modules, a marked transition occurred between 120 and 110 mm. Inclining the PV modules by 10 degrees increased the critical gap beyond 120 mm. [2] The exact numerical boundary should therefore not be treated as a universal design value, but the experiments clearly demonstrate how sensitive the fire behaviour is to relatively small geometric changes.

The modest difference between the horizontal steel plate and the actual PV module should not be overstated. More revealing was what happened after flame spread began. The geometry of the real module, including its aluminium frame, changed the movement of hot combustion gases and affected the

length and severity of flame spread underneath the module. A nominal gap dimension therefore does not completely describe the effective cavity in which a fire develops. The polymeric backsheet of the tested modules was not found to sustain significant flame spread by itself. [2]

What exactly is being measured?

For practical design discussions, a statement such as “the modules have 300 mm clearance” is incomplete unless the reference points are defined. The relevant distance might be measured to the rear surface of the module, the lower edge of its frame or another part of the mounting structure. With an inclined module the clearance also varies along its length, so the location at which the dimension is specified matters.

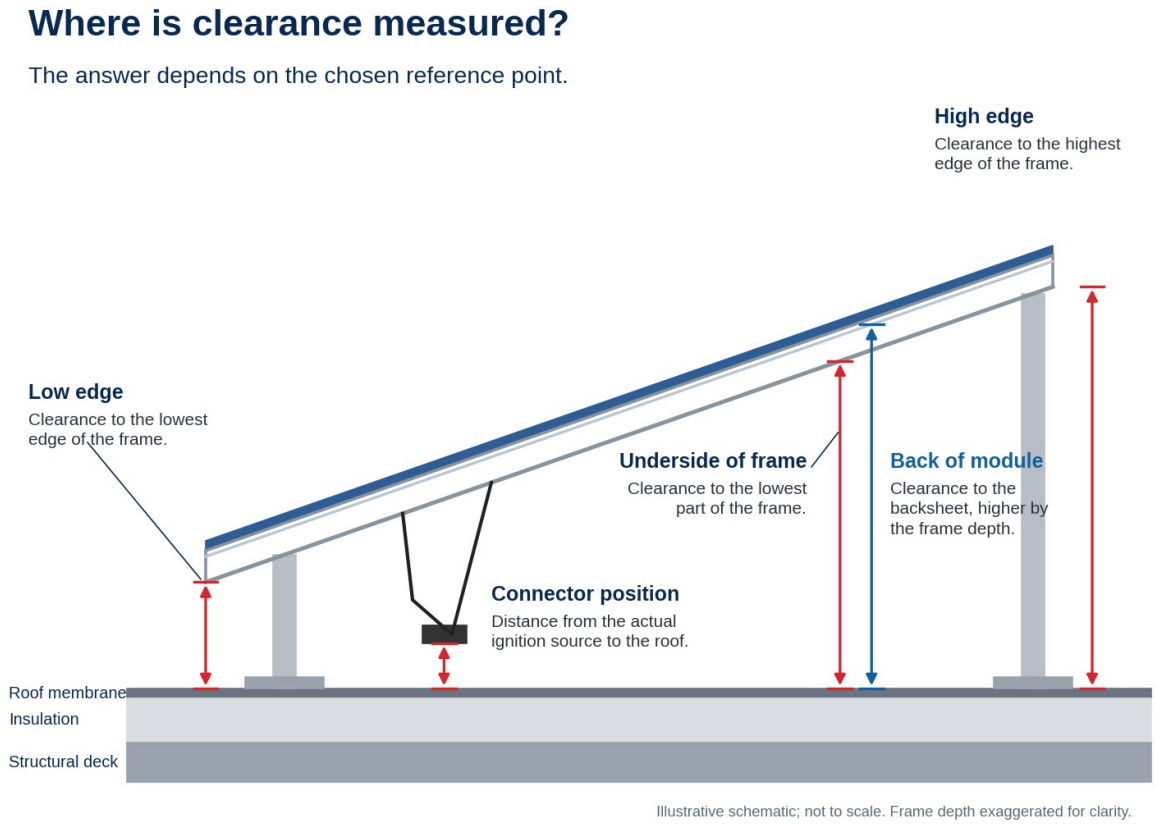


Figure 1. The same installation can be described by different clearances depending on the reference point chosen. Module clearance and source clearance are separate dimensions.

In the Kristensen experiments, the gap associated with actual PV modules was defined in relation to the rear side of the module. The frame itself extended below that surface and changed the cavity flow. This distinction is directly relevant to design specifications because two systems described using the same nominal clearance may create different effective fire geometries. [2]

The experimental roof strip was also limited in width. The authors explain that this restriction mainly affected the subsequent flame-spread pattern because the fire could primarily develop along the length of the specimen. They also state that the ignition region itself was considerably narrower than the roof strip, meaning the limited width did not determine whether propagation initially extended beyond the ignition zone. [2]

Those observations are compatible with the authors' broader caution. In their conclusion they state that the critical gap heights determined for the experimental arrangements cannot be implemented directly in industry because they are not universal and because the limited-width roof mock-up does not represent

an effectively infinite flat roof. [2, p. 2049] The width limitation therefore does not invalidate the observed transition. It is one of the reasons the measured value should not become a universal design dimension.

What Chapter 1 establishes

The experimental evidence supports two propositions at the same time. Clearance can materially influence flame spread and thermal interaction beneath rooftop PV, and increasing it can therefore be an effective mitigation measure. At the same time, the location of the transition depends on the physical configuration in which the measurement was made.

> KEY QUESTION
>
> **Can one clearance demonstrate the fire performance of a roof-PV assembly independently of the fire source, module geometry and roof construction?**

The evidence in the next chapter shows why that question cannot be answered by a single distance.

2. Is there a safe clearance?

Once clearance is accepted as a meaningful fire-mitigation parameter, the obvious design question is how much is enough. Published experiments do identify critical gap regions, but the numerical result changes when the source, geometry, roof construction or objective of the test changes. A critical clearance is therefore not a property of the gap alone. It is an outcome of a defined combination of fire source, geometry and roof-PV configuration.

> **DESIGN QUESTION**
> **Enough for what fire scenario, on what roof and demonstrated by what evidence?**

The fire source forms part of the result

The 2022 Kristensen experiments used a reproducible wood crib as the ignition source rather than a real electrical DC arc. The resulting critical gap values are therefore measurements of how those roof-PV configurations behaved under that controlled source. They are not direct measurements of a safe distance for every possible electrical failure. [2]

This limitation is not a weakness in the research. Reproducibility is essential when the purpose of an experiment is to isolate the effect of geometry. It does mean that the critical clearance must remain connected to the exposure with which it was established. Whether a real DC arc produces a thermal exposure equivalent to a particular wood crib is a separate engineering question, addressed in *DC arcs as a heat source*. For this paper the relevant conclusion is simply that an experimentally determined gap value cannot be separated from the characteristics of the source used to determine it.

Stølen, Fjærestad, Mikalsen and Jomaas provide a useful complementary experiment. Their study concerned a 30 degree sloped roof with a B_{ROOF}(t2)-rated bituminous membrane on a chipboard substrate and is therefore not representative of a commercial flat roof. Its value here is methodological: during part of the programme the researchers held the gap at 120 mm and varied the ignition source. [3]

The outcome changed materially. The study found that both gap distance and the size and characteristics of the ignition source influenced how far the fire propagated. The authors conclude that the complete system of roof composition and PV installation needs to be considered as a whole to ensure an adequate level of fire safety. [3]

The source should not be characterised by peak heat-release rate alone. Different wood-crib configurations varied in geometry, fuel mass, duration and estimated heat-release rate. A lower peak output sustained over a different period can produce a different interaction with the cavity than a shorter source with a higher peak. The relevant variable is therefore the fire exposure as a whole, not one kW figure in isolation.

A test distance is not automatically a design distance

The proposed ZAG test method for flat roofs with PV illustrates a related distinction. It places a wood crib beneath a PV module with approximately 120 mm between the roofing membrane and the rear side of the module at the centre of the ignition source. [4]

Read without context, such a dimension could easily be mistaken for an implicit safe or recommended clearance. The method states something different. Its objective is to assess damage to the roof elements after fire propagation has occurred. If the initial configuration does not result in sufficient ignition and spread, the gap can be reduced or the ignition source enlarged. The method explicitly states that it is not intended as an ignition test. [4]

This demonstrates why a number in a test procedure should never be detached from the purpose of the test. The 120 mm distance is a parameter used to create the fire exposure necessary to evaluate roof response. It is not a conclusion that 120 mm is either safe or unsafe as an installation dimension. The same distinction links the present paper to *What does a fire test actually prove?*: a laboratory geometry only acquires meaning when the question the laboratory intended to answer is understood.

Different scenarios can produce different clearance results

The UL/Solar ABCs programme provides another perspective. Phase 4 investigated whether increasing the separation between a rack-mounted PV surrogate and a low-slope roof could maintain the original Class A roof performance under the UL 790 spread-of-flame exposure. In the tested configurations, flame spread exceeded the applicable Class A criterion even at a separation of 24 inches, approximately 610 mm. [5]

This finding should not be used to claim that 300 or 400 mm is ineffective against a local fault underneath a PV module. The UL exposure was fundamentally different from the wood crib used beneath the modules in the Kristensen programme. It represented a fire approaching the assembly from its edge and interacting with the cavity under a standard roof-fire exposure.

Phase 4 was also not a controlled gap-height series in which all other variables remained constant. The tested roof covering changed across the selected clearance configurations, different distances were associated with different membrane types and the programme used a non-combustible surrogate rather than a functioning PV module. With one experiment per configuration, the results cannot be interpreted as a continuous failure curve solely as a function of gap height. [5] Their relevance is more specific: they demonstrate that a clearance which may be effective for one type of fire exposure cannot automatically be treated as evidence for another.

Earlier experiments in the same programme add another caution. Phase 5 reported earlier measurements in which a 5 inch gap produced higher measured heat flux and roof-surface temperature than the tested 2.5 and 10 inch configurations. Those experiments used a non-combustible roof and were intended to measure thermal exposure rather than flame propagation. They therefore do not prove that 5 inches is generally the most dangerous clearance. They demonstrate that the relationship between cavity height and thermal exposure does not have to be monotonic. [6]

An institutional clearance route also exists

Experimental studies are not the only basis on which clearance is used as a fire-risk measure. FM Property Loss Prevention Data Sheet 1-15, April 2026, provides an important institutional example. Section 2.1.1.2 applies to “existing roof covers that will remain when installing a new PV system above the existing roof cover”. It provides several alternative fire-risk routes, including using components from an FM Approved roof-PV assembly, specified re-cover or retrofit approaches, particular non-combustible roof arrangements and, as option E, elevating the panels. [7, §2.1.1.2, p. 5]

> **FM DS 1-15, §2.1.1.2(E)**
> **“Install PV panels at least 3 ft. (0.9 m) off the roof to eliminate the fire hazard. Ensure adherence with Section 2.1.2 due to increased wind loads.” [7, p. 5]**

Within that framework, increased physical separation is therefore accepted as a prescribed property-loss-prevention route, explicitly linked to the wind provisions of the same Data Sheet.

The scope of the route matters. FM separately addresses new roof-PV systems and projects where the roof is being replaced in conjunction with the PV installation under §2.1.1.1, where the recommended

route is an FM Approved roof-PV system. [7, §§2.1.1.1–2.1.1.2, pp. 4–5] FM also recommends that an existing roof more than ten years old should generally be replaced or re-covered before installing PV expected to remain for 25 to 30 years, and that PV should not be installed over roofs showing deterioration such as leaks, cracks, delamination or blisters. [7, §2.1, p. 4]

Within that framework, the 0.9 m value is a legitimate prescriptive risk-control route. It should not be interpreted as a universal experimentally derived critical clearance applicable independently of the conditions under which FM accepts it. Its structural and aerodynamic implications are considered in Chapter 4.

For context, the highest clearance investigated in the UL Phase 4 configurations was approximately 610 mm, whereas the FM elevated route begins at 900 mm. These values arise from different methodologies and should not be interpreted as points on one continuous fire-response curve. Their juxtaposition simply illustrates how strongly the meaning of a numerical clearance depends on the framework in which it is used.

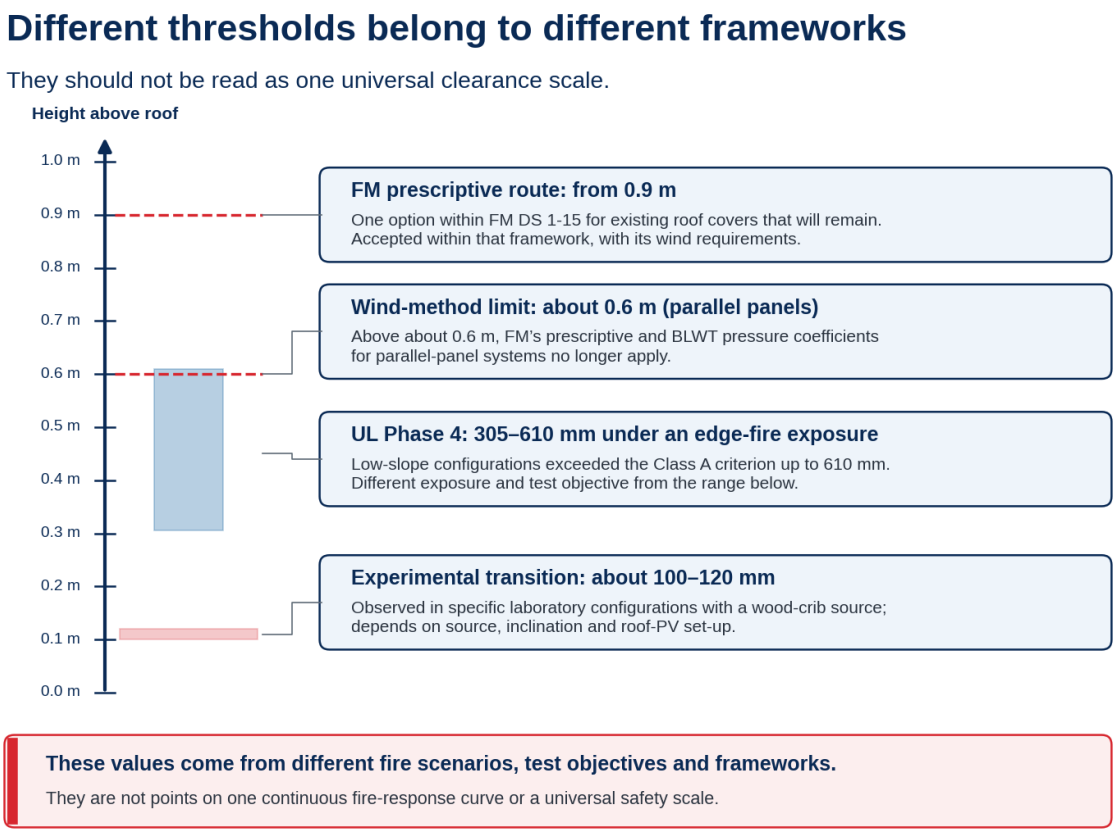


Figure 2. Values from different fire scenarios, test objectives and frameworks. They are not points on one universal clearance scale.

Geometry remains part of the fire scenario

Gap height interacts with several other geometric variables. Module inclination, frames, rails, array edges and the position of the ignition source can alter both the direction of hot-gas flow and the amount of radiative and convective exposure delivered to the roof.

Kristensen and colleagues found different critical conditions for horizontal and inclined modules and showed that the frame of the real PV module affected gas movement inside the cavity. The same nominal gap can therefore represent a different fire environment when another part of the geometry changes. Their experiments also showed that adding PV above a roof that had already demonstrated external-fire performance can create fire dynamics that are not represented by testing the roof in isolation. [2]

The independent review reaches the same boundary

The 2026 UK literature review prepared for the Building Safety Regulator treats critical gap height as a potentially useful mitigation measure while identifying reported values as case-specific and cautioning against turning them directly into universal installation requirements without further validation. [8]

The review does not reject clearance. It recognises clearance as one legitimate design parameter while separately discussing measures relating to roof construction, including non-combustible roof coverings, exterior-to-interior fire resistance and PV-layout measures such as gap height and array segmentation. [8] That separation is important. A geometrical measure can reduce the probability or severity of exposure while the construction underneath determines what happens if significant exposure nevertheless reaches the building envelope.

The result is a narrower but stronger conclusion. Clearance acquires technical meaning within a defined fire scenario, PV geometry and roof construction, and only to the extent that the evidence used was designed to demonstrate that scenario.

3. What clearance does not change

Understanding what clearance does not change is as important as understanding what it does. Raising or lowering the PV array primarily modifies the cavity between the source and the roof. It can change flame geometry, heat transfer and the likelihood that a local incident develops into sustained roof involvement. It does not remove the initiating fault and it does not alter the composition of the roof below it.

This distinction separates three engineering functions: ignition prevention, exposure mitigation and consequence limitation. Clearance belongs primarily to exposure mitigation.

Clearance does not remove the initiating fault

Changing the module height does not prevent defective connectors, damaged cables, junction-box failures or other electrical faults from occurring. Component selection, installation quality, cable routing, inspection, monitoring and fault detection remain separate lines of defence.

Clearance serves another function. By changing the physical relationship between a local source and the roofing system it can reduce the thermal exposure delivered to the construction. This can substantially influence the probability of escalation, but it should not be confused with removing the initiating event. The fire experiments discussed above do not establish the relationship between an actual DC arc and every laboratory ignition source; that question is reserved for *DC arcs as a heat source*. Keeping the distinction explicit prevents clearance from being assigned a preventive function that it does not perform.

Module clearance is not source clearance

Another practical distinction becomes important when raised PV is discussed. Module-to-roof clearance and source-to-roof separation are not necessarily the same dimension. A module can be installed several hundred millimetres above a roof while connectors, cables or other electrical components remain much closer to the membrane. Raising the module does not automatically raise every credible ignition source by the same amount.

The companion paper *DC arcs at the roof-PV interface* discusses this problem at component level. In the ARC002 unprotected reference configuration, a connector positioned approximately 50 mm above EPDM or bituminous roofing carried a 35 A arc that was deliberately maintained for five minutes, resulting in membrane ignition and subsequent flame spread. [9, §2.1] The relevant distance in that experiment was source-to-roof separation rather than module-to-roof clearance.

That distinction matters in design reviews. If clearance is being credited as a fire-mitigation measure, the actual location of connectors, cable loops, junctions and other possible heat sources should be established rather than inferred from the nominal module height. A raised array can improve the cavity geometry while still containing local electrical components close to the roof. The geometry of the source and the geometry of the module should therefore be documented separately.

Clearance does not change the roof build-up

Changing the distance above the roof leaves the underlying construction unchanged. The membrane remains the same membrane, the insulation retains the same composition and any adhesives, separation layers, vapour-control layers, existing coverings and structural deck remain part of the same assembly.

Clearance can reduce the thermal demand reaching those materials in a defined scenario. It does not change their behaviour once a particular exposure reaches them. The resistance of the roof assembly therefore remains a separate design question from the geometry used to reduce its exposure. The Building Safety Regulator review makes the same separation in its recommendations: minimum gap

heights and array configuration are treated as layout measures, while non-combustible roof coverings and appropriate exterior-to-interior fire resistance are identified separately as means of limiting fire spread and roof penetration. [8]

> **KEY DISTINCTION**
> **Clearance acts on the cavity. It does not change the roof build-up beneath it.**

This distinction does not diminish clearance. It defines its place within a layered risk strategy.

Residual exposure remains possible

No cavity geometry can be assumed to reduce exposure to zero under every possible event. Radiation and convection can continue to transfer heat to the roof and some module constructions may release hot or burning material when involved in fire.

Kristensen and colleagues observed burning droplets from one of the tested PV-module types. They also concluded that the droplets did not materially increase the exposed area in their flat-roof test arrangement. [2] That finding supports the possibility of falling burning material without supporting a broader claim that increased clearance automatically spreads such debris over a larger roof area.

The more important question is what happens if the residual thermal exposure is still sufficient to involve the roofing system. At that point the design problem moves from exposure mitigation to consequence limitation.

Different measures answer different questions

The distinction developed in this chapter can be stated simply. Electrical integrity and monitoring act primarily on the probability of ignition. Clearance and cavity geometry influence the pathway by which a local incident exposes the roof. The roof assembly determines how the construction responds if significant exposure nevertheless occurs. These functions are complementary rather than interchangeable. Chapter 6 brings them together in a three-layer defence-in-depth model.

4. The roof does not see clearance in isolation

Clearance is simultaneously a fire-dynamics parameter and a building-design parameter. Changing the distance between a PV module and the roof can influence ventilation, module temperature, rear-side irradiance, wind loading, support forces, waterproofing details, snow accumulation, future roof renewal and access for inspection. The relevant design problem is therefore not whether more or less clearance is inherently desirable. It is how a particular clearance performs against several requirements that may not share the same optimum.

Clearance can provide benefits beyond fire

A larger space beneath a PV module can improve natural ventilation in some configurations. Increased airflow can reduce module operating temperature, although the magnitude of that effect depends on the panel arrangement, roof geometry and environmental conditions. It should therefore not be translated into a universal efficiency or energy-yield percentage. [10]

Clearance can also influence bifacial performance. Rear-side irradiance depends on surface reflectance, row arrangement, shading and the geometric relationship between the module and the surface beneath it. Greater height can consequently form part of a bifacial optimisation strategy without being inherently optimal in every installation. [11]

Depending on the mounting arrangement, additional physical space may also improve access for visual inspection, wiring checks and roof maintenance. Support frames, bracing and cable runs can equally restrict that access, so maintainability should be considered as an actual layout issue rather than inferred from clearance alone. These advantages matter because the engineering question should not be framed as a list of penalties associated with raising modules. Clearance is a legitimate multi-objective design variable.

Wind loading is also clearance-dependent

The same geometric change affects aerodynamics. Alrawashdeh and Stathopoulos specifically investigated the effect of underneath-array clearance on multipanel PV systems mounted above a flat roof. Their wind-tunnel programme compared full-scale equivalent clearances of 0, 200 and 400 mm and measured pressures on both upper and lower panel surfaces. [12]

The results do not support a simple rule that either higher or lower arrays are always preferable for wind. The effect depended strongly on wind direction and on the position of an individual panel within the array. In this study, the lowest clearance produced the more severe uplift-related pressure conditions while reducing downward net pressure in other load cases. [12]

Fire and wind therefore do not necessarily demand opposite geometric changes. The fire evidence is itself non-monotonic at small gaps and neither body of evidence establishes a general optimum; the direction of each effect has to be established for the actual configuration. Research on low-profile arrays on large, nearly flat roofs supports the wider point that wind behaviour depends on building and array geometry. Kopp found that loads varied with building dimensions and array configuration and that the roof zones relevant to PV arrays differed from those used for the bare roof. [13]

The engineering conclusion is therefore not that fire engineering wants higher arrays while wind engineering wants lower ones. It is that a geometry selected from fire evidence is not automatically the geometry that produces the most favourable structural wind response.

FM illustrates risk displacement in practice

FM DS 1-15 provides an unusually clear example of why a geometrical fire measure cannot be separated from other engineering disciplines. Its fire section allows panels to be installed at least 0.9 m above an existing roof cover as one of the available routes under §2.1.1.2. The same Data Sheet states in §3.4.14 that pressure coefficients from wind-tunnel tests remain valid only within the geometric limits of the tested configuration. For parallel-panel systems raised more than approximately 0.6 m above the roof, FM states that the prescriptive and boundary-layer wind-tunnel pressure coefficients are no longer valid. [7, §3.4.14, pp. 22–23]

The significance is not that the 0.9 m fire route is structurally problematic. It is that the prescribed fire-separation route crosses an aerodynamic evidence boundary before its prescribed fire clearance is reached. The designer therefore has to change the wind-assessment basis rather than assume that evidence applicable to a conventional low-profile array remains valid. This is a concrete example of risk displacement: solving one problem through geometry can move the project into a different engineering domain in which additional evidence is required.

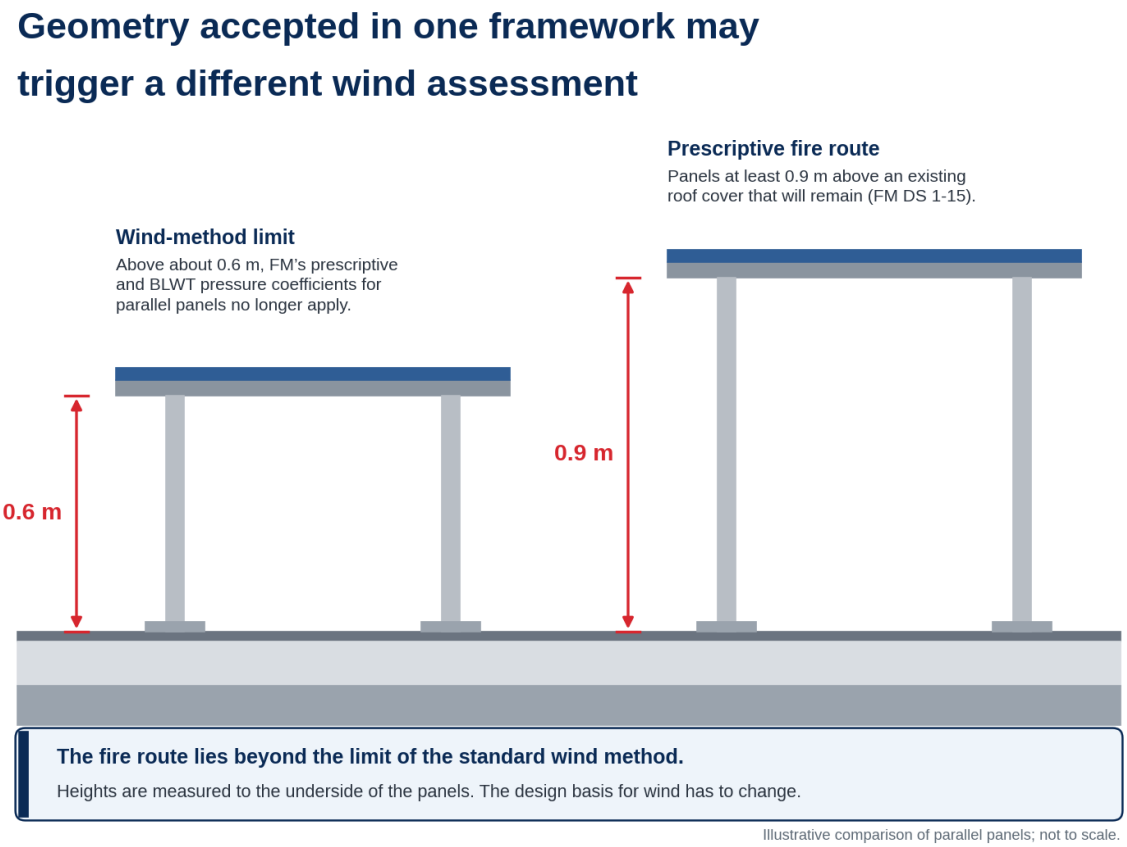


Figure 3. A geometry accepted within one fire framework can move the design outside the limits of the standard wind-assessment method.

The load path remains a building question

Wind pressure on a PV module does not end at the module. The resulting forces must pass through the mounting system and ultimately into the building. A ballasted installation may resist part of those actions through self-weight, ballast and friction. A mechanically attached installation transfers forces through fasteners or supports towards the structural deck or frame. Hybrid systems combine these mechanisms.

Changing the geometry can alter forces and moments within this load path. Whether this results in additional ballast, different fixing requirements or a change in support spacing depends on the actual

system. There is therefore no defensible generic conversion from millimetres of clearance to kilograms of ballast. For an existing commercial roof, the critical question is where those forces are transferred and whether the existing deck, purlins and primary structure have sufficient capacity to resist them.

Beyond structural actions

PV geometry also interacts with the waterproofing, service life and material strategy of the roof. Any support or anchorage that penetrates or locally interrupts the waterproofing layer becomes a detail that must remain watertight while accommodating structural and thermal movement. This is not unique to elevated PV, and an elevated system does not necessarily create an unacceptable waterproofing risk. The relevant issue is the actual support concept, the number and complexity of interfaces and whether they remain inspectable and repairable throughout the service life.

The same need for symmetry applies at roof replacement. A complete renewal of the waterproofing will normally require substantial removal of **any** rooftop PV installation, including modules, low-profile racking and ballast. An elevated support structure does not create that basic requirement. It can, however, add another layer of work where structural members, fixed supports or penetrations also have to be partly dismantled, inspected, reinstated and re-detailed before a new continuous waterproof layer can be completed. The relevant lifecycle question is therefore not whether one PV configuration must be removed and another does not, but how much intervention each geometry adds to the inevitable future interaction between the PV system and the roof beneath it.

Material use can be assessed in the same way. Increasing structural height may change the quantity of racking, steel, ballast, anchorage or other supporting material required, and on a large commercial roof even small changes per square metre become substantial at building scale. This does not establish that an elevated or low-profile system necessarily has a lower environmental impact; that requires a project-specific comparison of quantities, environmental product data, service lives, replacement cycles and potential reuse. Waterproofing, future roof renewal and material use should therefore be considered alongside wind and snow when alternative clearance strategies are compared.

Snow is another geometry-dependent action

PV also changes the geometry presented to snow. Tilted or elevated panels create obstructions around which snow can accumulate, erode or redistribute differently from an unobstructed flat roof. The second-generation EN 1991-1-3:2025 introduces a dedicated case for flat roofs with parallel rows of tilted panels, explicitly including solar and photovoltaic panels. [14,15] The model relates snow-load distribution to panel geometry rather than assuming that the roof behaves as an unobstructed flat surface.

The second-generation Eurocodes are in a transition period. Their formal applicability depends on national adoption and the relevant National Annex. The Joint Research Centre timetable provides for national implementation by 30 September 2027 and withdrawal of conflicting first-generation national standards by 30 March 2028. [16] The applicable edition therefore has to be established for the jurisdiction and project being assessed.

The important principle for this paper does not depend on a particular national coefficient. A geometrical change introduced partly for fire reasons can alter the snow-load case and has to be checked structurally. Generic statements such as “400 mm doubles the snow load” should not be used without the actual panel geometry, climatic snow load and applicable design rules.

The existing roof defines the design space

The distinction between a new building and an existing commercial roof is fundamental. In a new building the roof structure and PV installation can be designed together: deck, purlin spacing, primary structure

and drainage can all be selected with the final PV configuration in mind. A retrofit starts with those decisions already made. The deck, purlins, primary frame, waterproofing, insulation and drainage exist before the PV designer arrives, and their condition and remaining capacity establish the design space within which the array has to fit.

For that reason, remaining structural capacity should not be represented by a generic number such as 15 or 20 kg/m². Buildings that look similar externally can have different spans, deck profiles, original design assumptions, modifications, corrosion histories and imposed-load reserves. The same roof-first review also needs to consider drainage: supports or ballast can interact with existing drainage routes while local deflection can influence ponding. These are not inevitable consequences of a higher array, but they are examples of building conditions that cannot be understood from a PV layout in isolation.

Clearance as a constrained design variable

The combined evidence suggests a different way of treating clearance. It is neither a simple safety distance nor a parameter that should automatically be maximised. It is a constrained design variable whose value has consequences across several engineering disciplines.

Fire engineering can establish how a particular geometry influences flame spread and thermal exposure. Wind engineering can determine the aerodynamic response of the actual array and building. Structural assessment can establish whether the resulting actions can be transferred safely through the existing roof. PV design can address electrical integrity, energy performance and maintainability, while the roof assessment adds waterproofing, drainage, condition and remaining service life. The purpose of a roof-first sequence is to make these requirements meet before a single variable is optimised in isolation.

This brings the first half of the paper back to two connected but distinct questions. The first is whether the chosen design sufficiently reduces the probability that a local PV incident develops into significant roof exposure. The second is whether the complete roof-PV assembly is sufficiently resilient if significant exposure nevertheless occurs. A favourable answer to the first question does not automatically answer the second.

5. What if fire exposure still reaches the roof?

The preceding chapters establish that clearance can materially influence fire exposure but cannot demonstrate that significant roof exposure will never occur. For an asset owner, insurer or designer, that leaves a second decision that should be made explicitly: what residual exposure is being accepted, what outcome from the roof is considered acceptable and what evidence demonstrates that outcome? The question is no longer only how effectively the PV geometry limits the transfer of heat. It becomes what the roof is expected to do if that mitigation is exceeded.

A roof classification answers a defined test question

External-fire classifications provide important evidence, but only for the roof construction, field of application and exposure represented by the applicable test method. A B_{ROOF}(t4) classification, for example, is derived from a prescribed external-fire exposure incorporating burning brands, wind and supplementary radiant heat. [17,18] That evidence remains valuable when PV is proposed, but it should not be extended beyond the question the test was designed to answer. PV creates a semi-enclosed geometry in which flames, hot gases and radiation can interact differently with the roof. The test exposure also represents fire arriving at the roof from outside, such as burning brands from a neighbouring fire. It does not address an ignition source located permanently on the roof itself, such as the electrical components of a PV installation. Even where the PV geometry no longer materially changes the cavity exposure, that on-roof source remains outside the scenario the classification was designed to assess.

FM makes this boundary particularly explicit within its own framework: “Testing has shown that unless specifically FM Approved as a roof-PV system, an FM Approved Class 1A roof assembly (without PV) will have the FM Approved Class 1A rating negated once PV panels are placed above it.” [7, §2.0, p. 4] Within that framework, PV is treated as a change to the system being evaluated rather than as an accessory added above an otherwise unchanged classification.

The same principle underlies the Building Safety Regulator review, which treats PV-layout measures and roof-construction measures separately. [8] A roof classification remains relevant evidence. It should not be expected to demonstrate a fire geometry outside the test from which the classification was derived.

Once exposure reaches the roof, the assembly matters

If sufficient thermal energy reaches the roof, clearance is no longer the only parameter governing the outcome. The waterproofing, protective layers, insulation, adhesives, vapour-control layers and supporting construction become part of the event.

Independent large-scale experiments by Kristensen and Jomaas provide a useful illustration. Four flat-roof mock-ups combined a B_{ROOF}(t2) PVC roofing system and 150 mm EPS insulation with either a 40 mm PIR or a 30 mm mineral-wool mitigation layer beneath six PV modules. The configurations were not like-for-like material comparisons and should not be used as a generic ranking of PIR against mineral wool. The project was partially sponsored by IKEA Services AB. [19]

In all four experiments, fire spread across the area beneath the array but no self-sustained fire continued beyond the PV-array area. In the wider PIR experiment, the PIR retained its protective function for approximately 63 minutes before its mechanical properties failed, the layer was penetrated and the underlying EPS ignited. Until that penetration, the maximum measured temperature below the PIR was approximately 175 °C. The maximum measured below the mineral wool was approximately 243 °C. Both values were more than 100 °C below the cited 366–405 °C range for piloted ignition of flame-retarded EPS. [19]

The finding demonstrates why measured temperature, direct penetration and deeper material involvement are different performance questions. In the PIR configuration, thermal degradation and eventual loss of mechanical integrity changed the protection mechanism despite the earlier temperature measurements. In the mineral-wool configurations the mitigation layer remained mechanically stable and EPS ignition was not observed. [19] The study is cited here not to select an insulation material; *Beyond flame spread* examines the underlying mechanisms in more detail. It is used only to establish that limiting direct flame penetration and limiting deeper thermal involvement are not necessarily the same performance objective.

The actual roof must first be known

This requirement creates a practical problem on existing commercial buildings: the complete roof assembly is often not reliably documented. An older roof may contain several refurbishment generations, such as a new membrane installed over an earlier covering, insulation retained from previous work, locally replaced boards, or adhesives and vapour-control layers that no longer correspond with the original specification. The construction shown on a drawing or in an old tender document may not be the construction present on the building today.

An acceptance criterion cannot meaningfully be checked against an assembly that has not been identified. Before evidence from a tested roof-PV configuration can be applied, the actual roof should therefore be established with sufficient confidence. Depending on the building and the quality of the existing information, that may require document review, visual inspection, local opening-up, core sampling, moisture assessment or a combination of these methods. This is more than a documentation exercise: the existing assembly determines both the fire response that has to be assessed and the structural boundary conditions discussed in Chapter 4.

What outcome is acceptable?

Consequence limitation is most useful when expressed in terms of outcomes rather than products. Depending on the building, occupancy, property value, business-continuity requirements and insurer expectations, an acceptable outcome may allow local damage while requiring that the incident remains confined and does not escalate into a wider property loss.

Possible performance objectives include limiting sustained flame propagation across or through the roof assembly, preventing direct flame penetration into deeper layers or the space below and limiting thermal involvement of combustible insulation to an acceptable extent. Structural performance may also become relevant where prolonged heating could compromise the deck or supporting construction.

These outcomes should not be treated as interchangeable. Preventing direct penetration does not demonstrate that combustible insulation remains thermally unaffected. Preventing fire spread outside a PV array does not demonstrate that every material beneath the array remains undamaged. Conversely, local damage does not necessarily mean that the consequence-limiting function has failed. What constitutes an acceptable outcome is also an asset decision rather than a fire-physics constant. As discussed in *The economics of the rooftop PV decision*, potential business interruption, damage to stored goods, production loss and recovery time can materially change the level of residual damage an owner or insurer is willing to accept.

Evidence should match the accepted outcome

A project review can connect the required outcome directly to the evidence used to support it.

Intended performance outcome	Primary defence layer	Evidence that can address the question
Limit flame propagation over the roof-PV interface	Exposure mitigation and consequence limitation	Representative roof-PV assembly fire test with relevant geometry and exposure
Prevent direct flame penetration into deeper roof layers	Consequence limitation	Test in which those layers are present and inspected after exposure
Limit thermal involvement of combustible insulation	Consequence limitation	Temperature measurements within the assembly combined with post-test inspection of insulation and interfaces
Demonstrate performance under a defined standard external-fire exposure	Consequence limitation	Applicable external-fire classification within its documented field of application

Table 1. Intended performance outcomes and the evidence that can address them.

The purpose of this matrix is not to prescribe one mandatory test. It illustrates the underlying rule: evidence should be selected because it addresses the required function, not simply because a certificate or test report already exists. A module certification does not establish the response of the roof. A clearance experiment does not establish the resilience of every underlying roof construction. An external-fire classification obtained without PV does not necessarily reproduce the semi-enclosed exposure created by the proposed array.

This follows the same evidence logic developed in *What does a fire test actually prove?* and *DC arcs at the roof-PV interface*: the claim, source, geometry, duration, receiving system and failure criterion have to correspond to the function being credited. The design sequence should therefore run in the opposite direction from a certificate-led approach. First establish the actual roof. Then define the unacceptable consequence, identify which defence layer is expected to prevent it and select evidence that addresses that specific function. That is the practical meaning of consequence-limiting design.

6. Three layers, three questions

The three-layer model used here is a refinement of the incident chain developed in *From risk models to design requirements*. That broader model follows the development from initiating fault through PV fire, roof involvement and potential building escalation. The present paper focuses on the transition between a local PV fire and involvement of the roof, and clearance acts primarily on that pathway.

The evidence reviewed in this paper can therefore be organised into three complementary layers of defence: ignition prevention, exposure mitigation and consequence limitation. Separating them prevents a successful measure at one stage in the event sequence from being treated as evidence for another stage. The model is not intended to prescribe three products or imply that every project requires identical measures. Its purpose is to clarify which engineering question each measure answers and where its evidence belongs.

Mapping the three layers to the wider series

This paper	Functional term used in the series	Relevant part of the incident chain
Ignition prevention	Source	Fault → ignition / local PV fire
Exposure mitigation	Pathway	PV fire → thermal transfer → roof involvement
Consequence limitation	Receptor / receiving system	Roof involvement → propagation / penetration → building loss

Table 2. The three layers of this paper mapped to the terminology and incident chain used elsewhere in the series.

The mapping matters because a measure belongs in the layer where it acts physically, not where its commercial description suggests it belongs.

Layer 1: Ignition prevention

The first layer acts on the probability that a damaging thermal source occurs. Appropriate components, compatible connectors, controlled terminations, cable management, inspection, maintenance, monitoring and fault detection can all reduce the probability that an electrical defect develops into ignition.

Real-world evidence reinforces the importance of this layer while illustrating why it cannot be treated as infallible. The Building Safety Regulator review reports Clean Energy Associates findings that safety issues linked to ignition hazards were identified in 97 percent of more than 600 audited rooftop PV installations. The review also refers to FM observations that fires can still occur at installations designed in accordance with its recommendations. [8] The conclusion is not that prevention is ineffective. It is that reducing ignition probability and managing residual consequences remain separate tasks.

> **LAYER 1 QUESTION**
>
> **How effectively does the design reduce the probability that a damaging thermal source occurs?**

Evidence for this layer should primarily address the electrical installation, its components, workmanship, inspection and detection strategy. It does not establish how the roof will respond after ignition.

Layer 2: Exposure mitigation

The second layer becomes relevant once a thermal source exists. Its function is to reduce the probability that this source produces sufficient thermal exposure to involve the roof or propagate across the roof-PV interface.

Clearance belongs primarily in this layer. Chapters 1 and 2 show that cavity height, module inclination, frames, source characteristics and other geometric parameters can materially alter flame behaviour and heat transfer. Array segmentation and separation from vulnerable roof details can act on the same pathway.

The companion paper *DC arcs at the roof-PV interface* distinguishes local interposition from measures acting over a wider part of the array. Connector housings, local shields and plates positioned between a specific source and the roof act on the pathway and therefore belong in Layer 2. Module clearance also modifies the pathway, but over the wider roof-PV cavity. Neither should automatically be treated as a property of the receiving roof assembly.

Their effectiveness remains scenario-dependent. A clearance has technical meaning only within the source, geometry and roof-PV configuration for which supporting evidence exists. Any optimisation within this layer also remains constrained by the wind, snow, load-path, waterproofing, lifecycle and existing-roof considerations set out in Chapter 4.

> **LAYER 2 QUESTION**
>
> **If an ignition source occurs, how effectively does the design reduce the thermal exposure transferred to the roof and limit the pathway for propagation?**

A validated clearance can make a substantial contribution to this layer. It should not be expected to answer the questions assigned to the layers before and after it.

Layer 3: Consequence limitation

The third layer becomes relevant when significant thermal exposure reaches the roof despite the preceding measures. Its purpose is to control the response of the receiving roof assembly and prevent a local incident from developing into a substantially larger roof or building loss. Layer 3 therefore contains measures that alter the receptor itself. These may include changes to the exposed roof surface, protective layers incorporated within the roof build-up or an inherently non-combustible or otherwise demonstrably robust roof construction.

A loose shield or plate positioned above the membrane solely to intercept a local source does not belong in this layer. Following the terminology used in *DC arcs at the roof-PV interface*, it is an interposition measure in the exposure pathway and therefore belongs in Layer 2. By contrast, a protective layer that forms part of the roof assembly and changes how that assembly responds after exposure belongs in Layer 3. The Building Safety Regulator review reflects the same functional separation by considering measures affecting PV layout independently from measures affecting roof construction. [8]

> **LAYER 3 QUESTION**
>
> **If significant thermal exposure reaches the roof, how effectively does the receiving assembly prevent the incident from escalating into a larger roof or building fire?**

The acceptable outcome remains project-specific. A high-value distribution centre, a food-production facility and a simple storage building containing little combustible or business-critical content may justify

different performance requirements. What matters is that the required outcome is stated explicitly and supported by evidence appropriate to the actual roof construction.

The layers are complementary, not interchangeable

The value of the model lies in separating functions. Layer 1 acts mainly on the probability of the initiating source. Layer 2 acts mainly on the pathway between that source and the roof. Layer 3 acts on the consequences once significant exposure has reached the construction. Clearance therefore sits primarily in Layer 2. Demonstrating that it performs effectively in a relevant scenario does not demonstrate that electrical faults cannot occur, and it does not establish how every possible underlying roof construction will behave once significantly exposed.

The logic also works in the opposite direction. A robust roof assembly does not justify poor connector practice, inadequate installation or weak inspection. Consequence limitation should not be used as an excuse to weaken prevention. Equally, strong prevention does not remove the need to understand residual consequences where they are material to the building risk. This is defence in depth: the three layers are not alternative answers competing for priority but address different stages in the same escalation pathway.

From three layers to a roof-first decision sequence

The model also explains why specialists can reach apparently different but individually valid conclusions. An electrical engineer may focus on reducing faults. A fire researcher may demonstrate that a particular clearance substantially reduces flame spread. A roof specialist may identify that the existing insulation or roof build-up remains vulnerable under significant thermal exposure. The error occurs only when the answer from one discipline is extended beyond the question that discipline actually examined.

For a commercial roof, the design process therefore has to resolve three questions in sequence: how is ignition being prevented, how is exposure being mitigated if ignition nevertheless occurs and how are consequences being limited if significant exposure still reaches the roof? The first practical step in answering them is not to select a clearance or a fire-protection measure. It is to establish what roof is actually present, what condition it is in and what constraints it imposes.

That does not require a new workflow. *Rooftop PV on existing commercial buildings* already defines the broader roof-first gate sequence. Chapter 7 therefore follows that sequence and asks one narrower question at each gate: what does clearance change?

7. A roof-first design sequence for clearance

The preceding chapters show that clearance is a meaningful fire-mitigation parameter, but not one that can be selected independently from the building beneath it. The practical question is therefore not simply how much clearance to provide. It is when clearance should enter the decision process, what has to be known before it is selected and what evidence is required before it can be credited as part of the fire-safety strategy.

This chapter does not introduce a new methodology. It follows the roof-first gate sequence developed in *Rooftop PV on existing commercial buildings* and asks a narrower question at each stage: what does clearance change at this gate? The sequence matters because clearance is not the starting point. It is one design variable within a decision that begins with the existing asset.

Baseline: establish what roof is actually present

The baseline gate is factual. Before any mitigation strategy is selected, the actual roof construction needs to be established with sufficient confidence. Drawings and specifications may provide a starting point, but previous refurbishments, retained insulation, local repairs and overlaid waterproofing can mean that the construction present today differs materially from the original documents.

The assessment should therefore identify the relevant waterproofing, insulation, protective layers, supporting deck and interfaces. Where records are incomplete, document review may need to be supplemented by inspection, local opening-up, core sampling or moisture assessment. The purpose at this stage is not to decide whether the roof should remain or how it should be protected. It is simply to establish what the receiving construction actually is, because neither clearance evidence nor roof-response evidence can be meaningfully applied to an unknown assembly.

Asset: define the consequence that matters

The asset gate establishes what loss the design is intended to prevent. The same local PV incident can have very different consequences depending on occupancy, stored goods, production processes, recovery time and business-continuity exposure. A degree of local roof damage that may be tolerable for one asset may be unacceptable for another.

The relevant question is therefore not only whether fire can reach the roof but what must not happen afterwards. Unacceptable outcomes may include wider propagation, penetration into the building, involvement of combustible roof layers, loss of weatherproofing, damage to valuable contents or prolonged operational interruption. As discussed in *The economics of the rooftop PV decision*, acceptable residual damage is an asset and risk decision rather than a fire-physics constant. Clearance can influence the exposure but cannot determine what consequence an owner or insurer should accept.

Structure: establish the geometrical design envelope

The structural gate establishes the physical range within which the PV geometry can be designed. The existing deck, purlins, primary structure, spans, previous modifications and available load reserve determine what additional actions can be accommodated. Changing array height can alter wind response, forces in the mounting system, support reactions and snow distribution, so a fire-driven clearance cannot be selected independently from those constraints.

As Chapter 4 showed, FM's elevated fire route lies beyond the height at which FM's normal low-profile pressure coefficients remain directly applicable, so the wind-assessment basis has to change with it. The purpose of the structural gate is therefore not to derive one universal maximum or minimum clearance. It is to define a defensible geometrical design envelope for the actual building.

Roof: decide whether the existing roof should remain or change

The roof gate addresses condition, remaining service life and the need for intervention. This is deliberately separate from the baseline gate. Baseline establishes what roof is present; the roof gate asks whether that construction is capable of remaining in service for the intended PV lifetime and whether refurbishment or replacement is already required.

FM DS 1-15 makes essentially the same distinction. Section 2.1.1.2 applies where existing roof covers will remain and provides a series of alternative fire-risk routes, including option E with PV at least 0.9 m above the roof. Section 2.1.1.1 applies instead to new roof-PV installations or projects where the roof is being replaced in conjunction with the PV installation, and directs those projects towards an FM Approved roof-PV system. [7, §§2.1.1.1–2.1.1.2, pp. 4–5] FM also advises replacement or re-covering of older roofs before installing long-life PV and states that PV should not be installed over roofs showing deterioration such as leaks, cracks, delamination or blisters. [7, §2.1, p. 4]

This supports a symmetrical roof-first decision rule. Where an existing roof is dry, serviceable and has sufficient remaining life, avoiding intervention in that roof can be a genuine benefit of a geometry-led strategy. Where the roof already requires refurbishment or replacement, that benefit largely disappears and the project gains an opportunity to reconsider the receiving roof assembly itself.

> **ROOF-FIRST RULE**
>
> **The value of clearance depends partly on whether it preserves a roof that could otherwise remain untouched. Neither route is inherently preferable; the condition of the existing roof decides which comparison is relevant.**

Mounting: define the geometry that can actually be built

The mounting gate converts the structural design envelope into a real PV configuration. At this stage, clearance should no longer be described by one nominal height alone. The design should define the reference points used for measurement, module inclination, frame depth, supports, row spacing and other geometric features that materially affect the cavity.

The mounting concept also determines how the system interacts with the waterproofing and how accessible the roof remains for inspection and future work; those consequences are addressed in Chapter 4. For the clearance decision, the key requirement is that the geometry credited in the fire strategy is the geometry the mounting concept is actually capable of delivering.

Risk: define the role of clearance within the defence strategy

The risk gate determines what the selected geometry is expected to achieve. The three layers developed in Chapter 6 provide the framework: Layer 1 reduces the probability of ignition, Layer 2 reduces the pathway by which a local source exposes the roof and Layer 3 limits escalation once significant exposure reaches the receiving construction. Clearance belongs primarily in Layer 2.

The roof response is considered at this gate because it determines how much reliance can reasonably be placed on Layer 2. A predominantly non-combustible roof presents a different consequence pathway from an assembly containing combustible insulation or uncertain legacy layers. *Beyond flame spread* examines the mechanisms of surface propagation, penetration and deeper thermal involvement in detail. The question here is narrower: after the residual exposure from Layers 1 and 2 has been defined, can the actual receiving roof deliver the consequence-limiting performance required at the asset gate?

Clearance is therefore neither automatically sufficient nor automatically insufficient. Its importance follows from the combination of the credible ignition source, the exposure pathway, the receiving roof and the consequence that has been judged unacceptable.

PV: distinguish module geometry from source geometry

The PV gate establishes where the credible electrical sources actually sit within the selected array geometry. Module-to-roof clearance and source-to-roof separation are different dimensions. Raising a module does not necessarily raise connectors, cable loops, junctions or other possible local heat sources by the same amount.

This connects directly with *DC arcs at the roof-PV interface*. That paper addresses the local pathway between an electrical source and the receiving surface, while the present paper addresses the wider cavity created by the PV array. A design that credits module clearance as an exposure-mitigation measure therefore also needs to record the position of relevant electrical components within that cavity. Neither distance should be inferred from the other.

Evidence: demonstrate the function being credited

The evidence gate asks whether each claimed function is supported by evidence that actually addresses it. Where clearance is credited with reducing cavity exposure, the evidence should be relevant to the source, geometry, clearance definition and receiving construction being claimed. Where the roof itself is credited with limiting consequences, the evidence should address the receiving assembly and the failure criterion that matters.

A prescriptive institutional route can also constitute evidence within its own framework. A project complying with FM DS 1-15 §2.1.1.2 option E can rely on that route within the FM property-loss-prevention framework, provided the associated conditions, including the applicable wind requirements, are also satisfied. That is a legitimate basis for acceptance within that framework. It does not turn 0.9 m into a universal physical threshold or automatically establish equivalence for another insurer, regulatory system or engineering assessment. [7]

This is the same principle developed in *What does a fire test actually prove?*: evidence has meaning in relation to the claim and the framework it is intended to support. The project dossier should therefore be built around the functions being credited rather than around whichever certificates happen to be available.

Approval: make the residual risk explicit

Approval should bring the preceding decisions together rather than reduce them to confirmation of one clearance dimension. The owner should understand which measures act on ignition, which act on exposure and which limit consequences. The insurer or risk engineer should understand the boundaries of the supporting evidence, while the roof, structural and PV designers should confirm that the final configuration remains consistent with their own design assumptions. Where the actual roof is uncertain, the structural geometry unresolved or the consequence criterion unsupported by evidence, achieving a nominal clearance on the PV drawing does not close the decision.

Investment: compare complete design routes

The investment gate compares the resulting technical routes rather than the cost of clearance in isolation. Where one route changes structural material, roof intervention, maintenance, future replacement or operational disruption, those consequences belong in the same comparison. The economic methodology is addressed in *The economics of the rooftop PV decision* and is not repeated here. The lowest initial-cost

geometry is not necessarily the lowest-cost asset decision, just as the largest clearance is not automatically the best fire-safety decision.

Delivery: verify the geometry as built

Clearance is particularly vulnerable to a difference between design intent and delivered construction because it is a physical geometry assembled on site. A clearance credited during design therefore has to be verified as built, using the same reference points by which the requirement was originally defined.

The same applies to source geometry. Connector positions, cable loops and cable management can change during installation even when the module layout follows the drawing. Substituting another module, frame or mounting component can also alter the effective cavity while leaving the nominal array layout largely unchanged. Commissioning should therefore verify both the relevant module geometry and the source positions on which the fire strategy depends. Where the installed configuration differs materially from the assessed configuration, the evidence and risk decision need to be revisited before handover.

> **DELIVERY**
> **A clearance credited in design has to be verified as built.**

Lifecycle: revalidate the strategy when either side of the interface changes

Clearance is not a permanent property of the building. It belongs to a particular relationship between a PV system and a particular roof assembly. During the service life of the building, modules may be replaced, racking may change, source locations may move and a second generation of PV may be installed. The roof itself may also be re-covered, refurbished or reconstructed.

A change on either side of that interface can affect the original strategy. A different mounting system may alter clearance and wind behaviour, while a different roof build-up may change the consequence-limiting role originally assigned to the receiving construction. Where cavity geometry is credited as part of the fire strategy, inspection should also establish whether accumulated combustible debris, displaced cables or other changes have materially altered the conditions on which the strategy was based.

Any clearance relied upon as part of the fire-safety strategy should therefore be revalidated when the PV geometry, mounting system, relevant source locations or roof build-up change. The lifecycle gate brings the process back to the questions used for initial approval: what roof is now present, what geometry exists above it, what function is clearance expected to perform and what evidence still applies?

What the sequence changes

The roof-first sequence changes clearance from a number selected at the start of a PV project into a variable that emerges from the building assessment. Its value depends on the actual roof, the asset consequences, the structural envelope and the combination of defence layers chosen for the project.

Different buildings can therefore rationally reach different conclusions. A serviceable existing roof may justify greater reliance on geometry to avoid unnecessary intervention. A roof already requiring renewal may create an opportunity to strengthen consequence limitation within the renewed assembly. Another project may combine both approaches. The objective is not to identify one universally preferred route, but to ensure that the selected clearance follows from the actual building, the required outcome and the available evidence.

8. What we do not yet know

The evidence reviewed in this paper is sufficient to establish that clearance matters. It is also sufficient to show that its effect depends on the fire source, PV geometry and receiving roof construction. What it does not provide is one universally validated clearance that can be transferred between those configurations. The remaining uncertainties therefore do not undermine clearance as a mitigation measure. They identify the boundaries within which experimental evidence, institutional requirements and engineering judgement can currently be applied.

Critical clearance has not been mapped across comparable roof assemblies

Existing experiments demonstrate critical transition behaviour, but the published studies do not form a matched dataset in which clearance is varied across otherwise identical realistic commercial roof assemblies, PV modules and ignition sources. Kristensen and colleagues isolated the effect of gap height effectively within specific experimental configurations, while the UL, Stølen and other programmes addressed different sources, materials, geometries or test objectives. [1–6]

In the sources reviewed for this paper, no programme was identified that systematically repeats the same clearance series across several representative commercial roof constructions while holding the remaining variables constant. The evidence therefore cannot yet show how much a measured critical region shifts solely because the receiving roof construction changes. A useful next step in research would not necessarily be the search for one universal clearance, but a matched programme showing how the same geometrical variable behaves when the roof, module or source is changed one parameter at a time.

The relation between real DC faults and module clearance remains poorly characterised

Most published clearance experiments use wood cribs or other controlled ignition sources because reproducibility is necessary when geometry is being compared. That is scientifically appropriate, but it leaves a translation problem between those test exposures and a real electrical fault.

The research series contains one source-specific data point. In the unprotected reference test reported in the manufacturer's technical note ARC002, a connector positioned approximately 50 mm above EPDM or bituminous roofing carried a 35 A arc deliberately maintained for five minutes, and the membrane ignited with subsequent flame spread. [9, §2.1] That observation is directly relevant to source-to-roof separation, but it does not establish the required module-to-roof clearance.

The evidence gap is therefore precise. There is an adverse DC-arc-specific observation at a small source-to-roof distance, while the critical module-clearance literature largely uses reproducible non-electrical sources. In the sources reviewed here, no systematic programme was identified in which actual DC arcing was combined with progressively varied module-to-roof clearance on realistic commercial roof assemblies. Critical gaps established with wood cribs should consequently remain linked to those exposures rather than being presented as validated separation distances from every credible DC fault.

Fire and wind evidence have not generally been generated for the same geometries

Fire behaviour and structural wind behaviour are both strongly affected by PV geometry, but published research typically examines them in separate experimental programmes. Fire studies select configurations to investigate cavity heat transfer and flame propagation, while wind-tunnel studies select configurations to quantify aerodynamic pressures. The resulting datasets therefore do not generally contain the same progressive clearance geometries tested against both phenomena. [1,2,12,13]

This should not be read as meaning that integrated approval frameworks do not exist. FM provides an important counterexample. Within its approval system, roof assemblies and roof-mounted rigid PV systems are assessed using standards including FM 4470 and FM 4478, while DS 1-15 connects the accepted roof-PV system to wind and other property-loss-prevention requirements. [7,20,21] FM therefore demonstrates that fire and structural resilience can be brought together within one approval framework, even though they are not established through one single experiment.

The narrower research gap is that publicly available fire and wind studies generally do not test the same progressive clearance configurations, which makes it difficult to place experimentally observed fire behaviour and wind response alongside each other for one identical array geometry. At project level, designers normally have to combine those separate evidence streams or work within an approval framework that already defines how they are brought together. The same applies increasingly to snow. The second-generation Eurocode gives a more explicit basis for snow around rows of tilted panels, but experience applying those provisions to the geometries studied in fire experiments is still developing. [14–16] There is therefore no basis for assuming that a geometry favourable in one discipline represents a multi-disciplinary optimum.

Clearance itself is not yet described by one common measurement convention

The experimental literature also exposes a more basic problem: clearance is not always the same geometrical quantity. It may refer to the distance to the rear module surface, the lower edge of a frame or another reference point. For an inclined panel the distance changes along the module, while framing and rails can extend into the nominal cavity. [2]

In the sources reviewed for this paper, no harmonised convention was identified that defines one reference geometry for reporting fire-related PV clearance across research, standards and installation guidance. A statement such as “300 mm clearance” can therefore remain ambiguous unless its measurement points are specified.

The experimental evidence is also geometrically narrower than much of the commercial market. Much of the critical-gap work discussed in this paper concerns horizontal or lightly inclined modules, individual test modules or relatively narrow experimental roof strips. Commercial flat roofs may use multi-row and east-west configurations in which the cavity contains additional edges, ridges, frames and interacting flow paths. A future test or installation framework should therefore define not only a required distance but also where and how that distance is measured.

The cavity changes during the life of the installation

Laboratory experiments generally start with a clean and controlled cavity. Real roof-PV systems do not remain in that condition indefinitely. Leaves, dust, vegetation debris, bird nests and other combustible material can collect below arrays, while cables may move and fixing arrangements can change during maintenance. FM identifies debris such as leaves and bird nests as possible combustible material beneath PV panels. [7]

What is not well characterised is how accumulation of such material changes the relationship between module clearance and fire exposure over long periods. An intermediate fuel package can create a source geometry very different from the clean membrane conditions used in a laboratory experiment. In the sources reviewed for this paper, long-term cavity contamination has not been incorporated into the critical-clearance datasets on which the present discussion relies. The design response belongs in the lifecycle gate described in Chapter 7; the research gap is the absence of quantified evidence showing how much ageing and contamination alter the fire benefit attributed to a given cavity geometry.

Roof classifications and PV fire exposure are not yet joined by one harmonised test framework

A wider standards gap remains between conventional external-fire classification and the modified fire environment created by rooftop PV. The 2026 literature review prepared for the Building Safety Regulator concludes that traditional roof-classification testing does not adequately characterise the altered fire dynamics created by PV, and identifies the need for more representative testing or for combinations of mitigation measures. [8]

This does not mean that no system-oriented assessment exists. FM Approval provides one such framework, and proposed research methods such as the ZAG procedure address specific aspects of complete roof-PV behaviour. [4,20,21] The gap is the absence of a broadly harmonised international test and classification framework that connects realistic PV geometry, roof construction, source exposure and clearly defined failure criteria in a form that can be applied consistently across the market. A related question is whether there is a height above which an open array no longer materially changes the roof exposure, so that the original roof classification becomes representative again. Physically, such decoupling is plausible, particularly for open arrays well above the roof. In the sources reviewed for this paper, however, no validated threshold was identified above which the cavity effect may be disregarded for the purpose of an external-fire classification. FM sets such a boundary within its own framework at 0.9 m; no equivalent was identified for B_{ROOF}. For inclined rows such a threshold would also have to be defined at the lowest relevant edge, and it would not address ignition sources that remain close to the roof surface.

Until such a framework exists, a B_{ROOF} classification, a laboratory clearance threshold and a system approval should not be treated as interchangeable evidence. Each remains meaningful within the scenario and decision framework it was designed to address.

The derivation of some prescriptive values is not publicly traceable

FM's 0.9 m route is particularly relevant because it is a clear example of an institutional clearance requirement. Within the FM property-loss-prevention framework, compliance with that route can be a valid basis for acceptance when its scope and associated requirements are met. [7]

The technical derivation by which the 0.9 m value was selected was not identified in the publicly available sources reviewed for this paper. That does not establish that the value lacks an engineering or experimental basis. It means only that the derivation is not published in the sources reviewed here. The distinction matters because prescriptive insurer guidance, experimental critical-gap measurements and regulatory classifications represent different forms of evidence. Each may be legitimate for its intended purpose without being evidence of the same underlying physical threshold.

What these gaps mean for design today

The remaining uncertainties do not justify waiting for a future universal clearance before installing rooftop PV. They justify greater discipline about what current evidence is being asked to prove. Clearance can be credited as an exposure-mitigation measure when its geometry, scenario and evidence basis are explicit. Institutional routes can be followed within the frameworks that define them. Roof-response measures can address credible residual exposure where the receiving construction requires additional consequence limitation. What the current evidence does not justify is turning one experimental threshold, one insurer dimension or one roof classification into a universal statement about the resilience of every commercial roof-PV system.

The evidence therefore continues to point towards a roof-first decision: establish the actual roof and the consequence that matters, define the structural and environmental design space, reduce ignition

probability, optimise the exposure pathway within those constraints and determine what the receiving roof has to do when meaningful exposure remains.

Conclusion

The engineering question is not whether clearance works. It does. Experimental evidence shows that changing the cavity between a PV module and a roof can materially change heat transfer, flame behaviour and the transition to sustained fire propagation. In some configurations, relatively small changes in gap height have produced large changes in fire behaviour. Within defined institutional frameworks, clearance is also used explicitly as a property-loss-prevention measure.

The harder question is whether optimising clearance is sufficient to demonstrate that the complete commercial roof remains resilient when ignition occurs. The evidence reviewed in this paper indicates that, in general, it is not. Within a framework that explicitly accepts a prescribed geometry for defined conditions, such as the relevant FM route, clearance can itself form an accepted part of the risk-control strategy. Outside such a framework, clearance alone does not demonstrate the resilience of the roof beneath it.

The effect of clearance depends on the source, the way the distance is measured, the module and mounting geometry and the construction receiving the residual exposure. The same geometrical decision also interacts with wind, snow, structural capacity, waterproofing and lifecycle requirements. This does not diminish the value of clearance. It locates it correctly. Within the defence-in-depth model developed in this paper, clearance acts primarily on the exposure pathway between an ignition source and the roof. It is therefore an important Layer 2 measure. It does not prevent every initiating fault and it does not, by itself, determine what happens inside the roof assembly if significant thermal exposure still reaches it.

That distinction is particularly important on existing commercial roofs. The first question should not be how high to mount the PV panels, but what roof is actually present, whether it can remain in service and what consequence the asset owner is trying to prevent. A serviceable roof that would otherwise remain untouched can make a geometry-led strategy attractive. A roof already requiring renewal creates a different opportunity, because consequence limitation can be reconsidered within the new roof assembly itself.

Evidence should then follow the function being claimed. A critical gap derived from one experiment belongs to that source and geometry. An insurer's prescribed clearance is valid within that insurer's framework and conditions. A roof classification applies within its documented field of application. A complete roof-PV system approval provides another form of evidence. None should be extended into a universal claim that it was not designed to support.

> **DESIGN QUESTION**
> **Enough for what fire scenario, on what roof and demonstrated by what evidence?**

That is a more useful design question than asking for one universally safe clearance. A roof-first approach does not reject clearance and it does not prescribe one alternative solution. It places clearance in the sequence where it can be assessed properly: after the existing building is understood and before residual risk is accepted. The objective is not the largest possible gap. It is a roof-PV system in which ignition prevention, exposure mitigation and consequence limitation are each addressed by measures and evidence appropriate to the actual asset.

> **CORE PRINCIPLE**
> **Clearance can reduce the pathway to roof involvement. It cannot, on its own, define the resilience of the roof that remains beneath it.**

References

[1] Kristensen, J. S., Faudzi, F. B. M. & Jomaas, G. (2021). Experimental study of flame spread underneath photovoltaic (PV) modules. *Fire Safety Journal*, 120, 103027. doi:10.1016/j.firesaf.2020.103027.

[2] Kristensen, J. S., Jacobs, B. & Jomaas, G. (2022). Experimental Study of the Fire Dynamics in a Semi-enclosure Formed by Photovoltaic (PV) Installations on Flat Roof Constructions. *Fire Technology*, 58, 2017–2054. doi:10.1007/s10694-022-01228-z.

[3] Stølen, R., Fjærestad, J. S., Mikalsen, R. F. & Jomaas, G. (2024). Experimental study of fire propagation on sloped roof with building applied photovoltaics. *Journal of Physics: Conference Series*, 2885, 012047. doi:10.1088/1742-6596/2885/1/012047.

[4] Rus, N., Simakovs, K. & Jomaas, G. (2025). *Fire test method for flat roofs with photovoltaic (PV) modules*. Version 3, January 2025. Ljubljana: Slovenian National Building and Civil Engineering Institute (ZAG) / FRISSBE. doi:10.13140/RG.2.2.24246.15688.

[5] Backstrom, B. & Sloan, D. (2012). *Report of Experiments of Minimum Gap and Flashing for Rack Mounted Photovoltaic Modules – Phase 4*. Project 11CA43479, File IN15911. Northbrook, IL: UL LLC / Solar ABCs, 29 March 2012.

[6] Backstrom, B. & Sloan, D. (2012). *Considerations of Module Position on Roof Deck During Spread of Flame Tests – Phase 5*. Project 11CA43479, File IN15911. Northbrook, IL: UL LLC / Solar ABCs, 24 July 2012.

[7] FM (2026). *Roof-Mounted Solar Photovoltaic Panels*. FM Property Loss Prevention Data Sheet 1-15, April 2026, full revision.

[8] OFR Consultants Limited (2026). *Thermal exposure to roofs from fires involving photovoltaic panels*. Project OX24065, Revision R02, 7 January 2026. Prepared for the Building Safety Regulator.

[9] Viridian Solar (2026). *ARC002 Functional Performance Testing of ArcBox Solar Fire Safety Enclosure*. Technical Note, Version 1.0, 12 March 2026.

[10] Naghavi, M. S., Esmaeilzadeh, A., Singh, B., Ang, B. C., Yoon, T. M. & Ong, K. S. (2021). Experimental and numerical assessments of underlying natural air movement on PV modules temperature. *Solar Energy*, 216, 610–622. doi:10.1016/j.solener.2021.01.007.

[11] Deline, C., MacAlpine, S., Marion, W., Toor, F., Asgharzadeh, A. & Stein, J. S. (2017). Assessment of Bifacial Photovoltaic Module Power Rating Methodologies: Inside and Out. *IEEE Journal of Photovoltaics*, 7(2), 575–580. doi:10.1109/JPHOTOV.2017.2650565.

[12] Alrawashdeh, H. & Stathopoulos, T. (2022). Experimental Investigation of the Wind Loading on Solar Panels: Effects of Clearance off Flat Roofs. *Journal of Structural Engineering*, 148(12), 04022202. doi:10.1061/JSENDH.STENG-10957.

[13] Kopp, G. A. (2014). Wind Loads on Low-Profile, Tilted, Solar Arrays Placed on Large, Flat, Low-Rise Building Roofs. *Journal of Structural Engineering*, 140(2), 04013057. doi:10.1061/(ASCE)ST.1943-541X.0000825.

[14] CEN (2025). *EN 1991-1-3:2025 Eurocode 1 – Actions on structures – Part 1-3: Snow loads*. Brussels: CEN.

[15] European Commission, Joint Research Centre (2025). *Snow load: EN 1991-1-3*. The second generation Eurocodes: key changes and benefits through design examples. Online workshop, 3–5 June 2025.

[16] European Commission, Joint Research Centre (2026). *Second Generation of the Eurocodes*. Eurocodes: Building the future. Accessed 4 October 2026.

[17] CEN (2012). *CEN/TS 1187:2012 Test methods for external fire exposure to roofs*. Brussels: CEN.

[18] CEN (2016). *EN 13501-5:2016 Fire classification of construction products and building elements – Part 5: Classification using data from external fire exposure to roofs tests*. Brussels: CEN.

[19] Kristensen, J. S. & Jomaas, G. (2018). Experimental Study of the Fire Behaviour on Flat Roof Constructions with Multiple Photovoltaic (PV) Panels. *Fire Technology*, 54(6), 1807–1828. doi:10.1007/s10694-018-0772-5. Partly sponsored by IKEA Services AB.

[20] FM Approvals (2022). *Examination Standard 4470: Single-Ply, Polymer-Modified Bitumen Sheet, Built-Up Roof (BUR) and Liquid Applied Roof Assemblies for Use in Class 1 and Noncombustible Roof Deck Construction*. April 2022.

[21] FM Approvals (2025). *Examination Standard 4478: Roof-Mounted Rigid Photovoltaic Modules*. March 2025.

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.

FM Approvals (2025). *Approval Report PR472014: AllShield BarrierSheet*. FM 4470 (April 2022) and FM 4478 (March 2025). Date of approval: 11 December 2025. RoofNav assembly 568563-0-0.

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

Approval Report PR472014 (11 December 2025), §§1.3, 3.2 and 9, identifies the applicable standards and evaluated assemblies and refers to RoofNav for the listing. The programme uses previous test data and waivers, including the 4478 external-fire requirement based on 4470 evidence. It does not document a new fire test with PV. Approval of a listed assembly and project-specific engineering acceptance are distinct decisions.
ZAG Report 448/25-560-1-EN, summary p. 3 and conclusion p. 46, reports prevention of direct flame penetration into deeper structural layers in the tested configurations during the documented test periods. Heat transfer, local gaps or burn-through and damage to underlying insulation were also recorded. These observations do not mean that all underlying layers remained undamaged. The programme has no unprotected reference from which a quantified reduction in risk can be calculated.

Scope and selection

This paper synthesises what module-to-roof clearance changes in fire behaviour and in the wider commercial roof design. Fire experiments, source-to-roof evidence and wind, snow and lifecycle guidance are included where geometry affects the decision. Connector-to-roof distance is kept distinct from module-to-roof clearance.

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

The dated references reviewed here run from 2012 to 2026. Publication years are taken from the references, not DOI, ISSN or page-range digits. This is not a claim of exhaustive coverage.

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 clearance is established across membranes, source conditions, array geometries and structural designs. Matched progressive fire-and-wind datasets and long-term contamination data remain limited.

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. Attributed the ZAG conclusion to ZAG and added the localised burn-through; corrected the author description.
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 fire-performance claims or tests.
1.02	1 November 2026	Method and sources: added PR472014 §3.2 and ZAG pp. 3, 46 boundaries and clarified the targeted synthesis. Existing FM DS 1-15 §§2.1, 2.1.1.2 routes retained; no universal clearance threshold introduced.
1.03	1 November 2026	Method and sources: publication range corrected to 2012–2026 from the references. Chs. 3/8: ARC002 identified locally as a manufacturer technical note. Third-party source attribution and authorship clarified.

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? (this paper)	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?	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.