
WHITE PAPER

How Far Can a SIP Span?

Why the answer is a drawing, not a number

1. The question

It comes up in the first meeting, nearly every time, and usually in the same words. *How far can a SIP span?*

It is a fair question, and the industry invites it. Every manufacturer publishes span tables. They are freely available, they are stamped by a third-party agency, and they look exactly like the kind of document that answers a question like that.

Open one and there are no distances in it.

There are spans along one edge and panel thicknesses along the other, and what sits in the cells is a load — pounds per square foot. The table is not telling you how far a panel can go. It is telling you that *at a given span*, a given thickness will carry a given load. You arrive with your load and read out how far you are allowed to go.

1.1 Which load?

In the Northeast the conversation usually turns straight to snow, and with some reason — it is the largest single number on most of our roofs, and the one that varies most from town to town. But it is not the load. It is one of five, and they are added together before anything is looked up.

Dead load is the roof's own weight, permanently: the panels, the underlayment, the roofing, the ceiling finish underneath. It is easy to underrate. On a race club we supplied in Rhode Island the dead load came to 35 pounds per square foot against 21 of snow — the heaviest thing that roof carries is itself.

Live load is temporary and movable — the crew and their equipment during construction, and whatever the code requires afterwards. On a children's center in New York it was 20 psf, exactly equal to that roof's dead load.

Snow load is the one everyone means. It starts from a ground snow figure for the site and is then adjusted for exposure, for whether the roof is warm or cold, for how important the building is, and for pitch — and then checked again for drifting, because snow does not lie evenly against a wall or in a valley.

Wind load behaves differently from the others and is easy to forget because it often pushes *up*. Uplift rarely decides how far a panel spans; it usually decides how the panel is fastened down, which is a different question with a different answer — and one §6 returns to.

Concentrated loads are the things bolted on afterwards: a rooftop unit, a solar array, a hanging fixture. The engineer's spreadsheet for that Rhode Island roof carries its own line for rooftop units

and another for photovoltaics. Both were zero on that job. They are on the form because they are frequently not.

Add them up and you have the number to take to the table. Snow is usually the largest term in the Northeast. It is rarely the only one that matters, and on at least one of our buildings it was not even the biggest.

1.2 So the answer moves

Change the ground snow and it moves. Change how much deflection the ceiling finish will tolerate and it moves again. Change the pitch and it moves twice, in opposite directions, for reasons §5 takes up. Add a solar array in year six and it moves once more.

Ask two people the same question about two buildings a hundred miles apart and they should get different numbers, and if they do not, one of them has not done the work.

So there is no single number to give you, and any supplier who offers one has quietly filled in your load for you.

There is, though, a better answer than a shrug — and it is more useful than the number would have been. It is not a distance. It is a drawing: where the structure goes, so that the panels are doing what they are good at.

2. Why a panel spans at all

Lay a sheet of OSB across two sawhorses six feet apart and stand on it. It will bend alarmingly. Glue a second sheet to the first with six inches of foam between them, put it back on the same sawhorses, and it will hold you without complaint.

Nothing about the OSB changed. What changed is where the two sheets are.



The separation is the structure — not the sheet

A structural insulated panel carries load the way an I-beam does, and the resemblance is not a figure of speech. In an I-beam the flanges — the top and bottom — do almost all of the work. Bending pushes the top into compression and pulls the bottom into tension, and the flanges are where those forces are carried. The web in the middle takes comparatively little of them. Its job is to resist shear and, above all, to hold the two flanges apart.

A SIP is the same arrangement in different materials. The two OSB facings are the flanges. The foam core is the web. One facing goes into compression, the other into tension, and the foam holds them at a fixed distance while stopping them sliding past one another.

Watch what happens when it bends. Both facings curve about the same center, so the outer one rides on a longer radius than the inner — and the gap between those two radii is the thickness of the panel. That is what thickness *is*, to a panel in bending.

Because the outer facing is on the longer arc it has further to travel, and must stretch to get there. The inner facing is on the shorter arc and must shorten. Their unwillingness to do either is what resists the bend.

Make the panel thicker and you widen the gap between those radii. Now the outer skin has to stretch further still for the same amount of bend, so it pushes back harder — and the two forces are working across a longer lever besides. The thickness counts twice, once in how hard the facings resist and once in how much leverage they have. Which is why stiffness grows with the square of it: move the facings twice as far apart and the panel is roughly four times as stiff, having gained nothing but foam. It is why a 10¼-inch panel does so much more than an 8¼-inch one, and why the difference costs so little set against what it buys.

It also explains something that surprises people: the foam does not need to be strong. The evaluation reports put the allowable core shear stress of expanded polystyrene at 4.5 pounds per square inch. You could press your thumb into it. Strength was never the core's job. Geometry was.

Which leads to the sentence the rest of this paper rests on. **A SIP is a structural member.** Not insulation that happens to be rigid, not sheathing with something useful in the middle — a beam, laid flat, that also happens to be the thermal barrier and the air barrier. It spans on its own account.

The question is how far, and what happens when we ask it to go further than it comfortably can.

3. What holds a panel up

In the ordinary case — the good case, and the one nearly every building we supply is built on — a panel **sits on top of** what holds it up.

The support is below the bottom skin. It might be an exterior wall, an interior bearing wall, a beam, a truss, a purlin, or a timber frame bent. What matters is not which of those it is but where it is: underneath. Load runs down through the panel, into the bearing, and away into the structure. Nothing hangs. Nothing is asked to grip.

It also means nothing passes through the insulation. The panel is continuous across the top of its supports, so the envelope stays whole and the structure stays below it. That is the arrangement a SIP is designed around, and it is worth noticing that it is also the cheapest one — ordinary framing members in ordinary places.

Which reframes the question the paper opened with.

The useful question is not how far a panel can reach. It is **where the supports can go so that the building is better for having them there.**

That is what actually happens when a set of drawings arrives. We do not start by looking up a span. We look at what the building already wants — a ridge, a bearing wall between two rooms, a line of trusses over the shop, a beam where the ceiling changes height — and we ask whether the panels can land on those. Most of the time they can, and most of the time the answer was already in the architect's drawing before anyone thought about panels.

When it is not, the fix is nearly always to add one more line of support rather than to ask more of the panel. The rest of this paper is about why, and what it costs when you go the other way.

4. The vaulted gable

An owner wants a cathedral ceiling. No flat ceiling, no attic, no collar ties across the room — just the underside of the roof, all the way up.

4.1 Nothing spans eave to eave

It is worth clearing away the question people expect to ask, because it was never available. A roof panel cannot span from one eave wall to the other, because there is a break in the structure at the ridge. Two panels meet there; they do not continue through.

Conventional framing has exactly the same problem, and collar ties are the answer to it. Without them the rafters push the eave walls outward and the ridge sags between them — the roof tries to flatten, and the walls go with it. The ties stop the spreading.

Which is precisely what the owner asked us not to build. A vaulted room has nowhere to put a tie.

So for a vaulted gable the ridge has to be a **structural ridge beam** — a beam that holds the two roof planes up rather than a board they lean against. That is not a clever move on our part. It is the entry fee. Every panel then runs from the eave wall up to the ridge beam, bearing on top of a support at each end, and nothing pushes outward anywhere.

You can see the two ends in the details we publish: RB1 at the ridge, where the panels meet on beveled blocking above the beam and are screwed down into it, and RP1 through RP6 at the eave — six ways of landing a roof panel on a wall, depending on what the overhang and soffit are doing.

4.2 So the real question is a shorter one

Not *how far can a panel span across the room*, but: **is eave to ridge short enough?**

Often it is, and there is nothing more to decide. A thirty-two-foot-wide room is sixteen feet from wall to ridge on plan, and a panel of ordinary thickness will cross that without an argument. (It is a little further than sixteen feet in reality, because the panel runs up the slope rather than across the plan — §5 takes that up.)

When it is not short enough, there are two answers, and which one you want is usually an architectural question rather than a structural one.

4.3 One: cut the run with a midspan beam

A **midspan beam** runs parallel to the ridge, partway down the slope between ridge and eave wall. It halves the eave-to-ridge run. A second halves it again.

Each one sits below the panel exactly as the ridge beam does, and the panel simply crosses over the top of it — see RB5, where the panel runs continuous over a beveled glulam and is screwed down into it. There is a choice worth making deliberately at that beam: the panel can run continuous over it, as in RB5, or a panel joint can land on it, as in RB4. Continuous is stiffer. The panel layout decides which you get, which is a good reason to settle the layout and the framing together rather than one after the other.

4.4 Two: turn the panels and use trusses

The other answer changes the direction the panels run.

Set trusses at eight feet on center, or wider, and the panels no longer travel up the slope at all — they span **horizontally, truss to truss**. The trusses can be multiple plies of conventional trusses, or custom beam trusses shaped to whatever the room wants to look like.

This one has a pleasing property: the panel's span is now simply the truss spacing, and the truss spacing is something you choose. Rather than asking how far the panel can go and building to suit, you decide what the panel should span and put a truss there.

Which of the two you use is usually decided by what the ceiling should look like. A midspan beam gives you one line across a continuous slope. Trusses give you a rhythm of them, and a shape you can design.

The way to span further is almost never a stronger panel or a heavier joint. It is one more line of support.

Beams and trusses are cheap. They sit under the insulation rather than running through it. And in a vaulted room they are usually something the design wanted anyway — a ridge, a purlin line, a repeating structural bay.

4.5 The same logic off the end of the roof

Overhangs work the same way. A panel will cantilever past the wall on its own for a foot or two; ask for more and the industry's advice is to reinforce the joint between panels with heavier material.

On a recent commercial job we did the other thing. Rather than escalate the panel joints to carry a long rake overhang, we drew forty LVL outriggers — twenty-six of them ten feet long, the longest a fourteen-foot two-ply — and let ordinary members carry the cantilever. Same principle as the midspan beam, turned on its side: give the panel something to sit on rather than asking it to do more.

5. Slope, which surprises people twice

Two things about pitched roofs catch people out, and they pull in opposite directions.

5.1 The panel is longer than the room it covers

Snow load is defined per square foot of **plan area** — the ground the roof covers — for the sensible reason that this is where the snow falls. A square foot of flat ground collects a square foot of snow whether the roof above it is flat or steep.

The panel, though, runs up the slope. It is longer than the distance it crosses on plan, and the steeper the roof the greater the difference. At 6/12 the panel is about 12% longer than its run. At 12/12 it is 41% longer.

Manufacturers publish correction factors for this. R-Control's are 1.05 at 4/12, 1.12 at 6/12, 1.20 at 8/12 and 1.41 at 12/12 — which is exactly the ratio of the sloped length to the flat one at each pitch. You multiply the horizontal run by the factor to get the panel, and divide the load by it to get what the panel actually carries per square foot of its own surface.

So a vaulted roof does not just look longer than the plan suggests. It *is* longer, and it is carrying over a longer distance.

5.2 Steep roofs shed snow, but later than anyone expects

The other direction. ASCE 7 reduces snow load as the pitch increases, on the reasonable grounds that snow slides off steep roofs. Everybody knows this, and most people assume it starts early.

It does not. For a warm roof that is not slippery — asphalt shingles, which is most of what gets built — the reduction does not begin until **30 degrees**, and a 6/12 pitch is 26.6 degrees. It gets nothing at all. Slippery roofs, standing seam among them, start earning credit at 5 degrees. And snow guards, which plenty of New England houses need over an entry, cancel the reduction entirely: snow that cannot slide is snow you must carry.

Two of our own jobs make the point better than the standard does. A children's center in New York has a 7/12 roof — 30.3 degrees, having only just crossed the threshold — and its snow calculation carries a slope factor of 1.000. A race club in Rhode Island has a 9/12, steep enough to earn a real reduction, and its engineer took 1.000 anyway.

Two steep roofs. Neither earned a thing for it.

5.3 Which is why the beam usually wins

Put the two together and the arithmetic is not what people expect. Take a room and vault it at 6/12, and the panel gets 12% longer while the snow load stays exactly where it was. The vaulted version of a given room is **harder** on the panel than the flat version, not easier.

That is one of the concrete reasons the ridge beam and the midspan beam in §4 tend to be the right answer. The slope has already spent some of the panel's capacity before anyone starts looking up spans.

6. The structural spline

A spline is the thing that joins one panel to the next.

It sits **between the skins**. A 2x spline drops into a slot cut in the foam along the panel edge and is nailed through both faces into the lumber; a surface spline is a strip of OSB in a shallower slot; a block spline is a narrow piece of panel doing the same job. All of them do one thing: hold two panels together in line so the roof behaves as one surface.

Then there is what the industry calls a **structural spline**, and it is a different animal — because it is not joining panels any more, it is holding them up.

6.1 What changes when it starts carrying

When a spline becomes structural, the load path turns through ninety degrees.

In §3 the panel sat on top of its support and the load ran straight down. Now the panel is beside its support rather than on it, and the only thing connecting the two is the panel's own skin — nailed to the side of a piece of lumber. The panel hangs, and it hangs from 7/16 of an inch of OSB glued to foam.

Recall from §2 what bending is already doing to that bond. The two facings ride on different radii, which means bending is permanently trying to change the distance between them, and the foam and the glue line are what refuse. Now add a demand to carry the panel's whole share of the roof into a member at its edge, through the same skin.

You can see how much this matters in an engineer's own numbers. On one of our jobs the fastener capacities are worked out both ways: a screw pulling straight out of the framing is good for 506 pounds, while the same screw's head pulling **through the OSB facing** is good for 237 — and the calculation notes that the second figure is what controls. Less than half. **The skin gives out before the fastener does.**

The evaluation reports say the same thing in their own language. They publish a *face peeling factor* — a number applied when a panel bears on its facing without something solid behind it — and its value is **0.4**. Bear that way and you keep forty percent of the capacity. The other sixty is the price of hanging rather than sitting.

6.2 So what does "spans up to 24 feet" mean?

It is the number everyone quotes, and it is not wrong. It is just describing something other than what most people hear.

Read the tables carefully and three things are always true of that 24 feet:

- The panels are about **four feet wide**, not eight. SIPA's own design guide says so.
- The structural splines sit at **four feet on center**. All three of the manufacturers
- The panel therefore spans **four feet**, between splines.

Which leaves the obvious question of what is spanning the other twenty. The spline is. A triple LVL spline at four-foot centers will carry a long roof for the same reason a triple LVL beam will, and those published "spline span tables" are, in substance, beam tables for the member with a panel attached.

There is a tell. In one manufacturer's charts a double 2x spline outperforms an I-joist spline. That looks strange if you think you are comparing panels, and it is not strange at all once you realize you are comparing beams.

The code reports draw the line themselves. Surface, block and I-joist splines are covered by the evaluation report. Dimensional lumber and engineered structural splines "must be designed in accordance with accepted engineering practice" — which is the polite way of saying you have left the document and are now on an engineer's desk.

6.3 Both strategies, one building

We can show the whole argument with two sheets from one job.

On a steel-framed children's center, the roof panel layout for one set of planes shows eight-foot panels with block splines between them — a joint, nothing more, with the skins nailed at six inches on center. The panels are doing the spanning.

The layout for the other planes shows narrow panels with double and triple LVL splines between them, and the skins nailed at **four** inches on center. Fifty percent more fasteners into the same skin, because the same skin now has fifty percent more work to do.

The engineer's calculations for that roof settle which member is which. The reinforcement — the spline — comes out at **100 percent** of its shear capacity: fully spent, the governing check in the whole package. The panel spanning between splines sits at 48 percent. And the fasteners tying one to the other get a line of their own, at 70 percent.

The spline is the structure. The panel is a passenger, and the nails are what stand between them.

7. When a structural spline is right, and when it is not

Nothing above is an argument that structural splines are bad. They are right in one class of building and almost never right in another, and the reason is economic rather than structural.

7.1 In houses: almost never

We hardly ever use a structural spline on a house, and not out of principle. It is that the alternative from \$4 — one more beam, one more line of trusses — is cheaper, keeps the insulation layer whole, and lands where the architecture wanted a line in any case.

If you find yourself reaching for a spline on a residential roof, the question worth asking first is where a beam could go instead. Nine times in ten there is somewhere obvious, and it costs less than the LVL would have.

7.2 In steel buildings: often, and for a good reason

Structural steel spans further than a SIP does, so a steel frame is naturally spaced more widely than a panel would like. That leaves a choice: add another line of steel, or add structural splines in the roof plane and let the panels bridge between what is already there.

Steel is expensive. Splines are not. The spline wins, and it wins on cost rather than on any structural virtue.

So the honest rule is not *avoid structural splines*. It is this: **a structural spline is a trade against the cost of another support, and which way that trade goes depends on what the supports are made of**. In wood, another beam is cheap. In steel, another frame line is not.

7.3 What it costs either way

Two things, one obvious and one that matters more.

The obvious one is **thermal bridging**. A structural member sitting in the plane of the insulation is a hole in the envelope. It is a small hole, and it is not a reason to panic — but everything the panel is for gets slightly worse, and you have spent money on a panel selected for how well it does exactly that.

The one that matters more is harder to see on a drawing. A structural spline means the panel is being **carried** rather than **carrying**. The structure is doing the job the panel was bought to do, and you are paying for both. In a steel building that trade is worth making, because the alternative costs more. In a house, where a beam costs very little, it rarely is.

Which is the sentence we end up saying to clients, in one form or another, on nearly every job:

You only want to use SIPs where you are maximizing the structural, insulating and air-infiltration properties of the SIP. The benefit of using SIPs is reduced when you are not taking advantage of all of its inherent strengths.

A panel that is spanning, insulating and sealing at once is doing three jobs for one price. A panel hung off a beam in the plane of the insulation is doing one and a half, and being charged for as though it were doing three.

8. Which is why we ask for the drawings early

Span is not a specification you look up. It is an outcome — of the snow at that address, the pitch of that roof, the finish on that ceiling, and above all of where the structure goes.

Those things are cheap to move while the drawings are still soft and expensive to move afterwards. A beam that appears on a plan in week two is a line on a page. The same beam in week twenty is a change order, a revised truss package and a conversation nobody enjoys.

So when a set of drawings arrives here, we do not begin by looking up a span. We look for what the building already wants — the ridge, the bearing wall between two rooms, the line of trusses over the shop, the beam where the ceiling steps — and we ask whether the panels can land on those. Usually they can. Usually the answer was in the architect's drawing before anyone mentioned panels at all.

When they cannot, we say so early, while adding one more line of support is still an easy thing to draw.

Which is the real answer to the question this paper is named after. **Not how far a SIP can span — but where the structure should go, so the panels do all three of the things you bought them for.**

That is a drawing, not a number. Send us your plans and we will show you ours.

This white paper is provided for informational and educational purposes. It does not constitute legal, engineering, or code compliance advice for any specific project.