

The threshold is a tax question. The answer is a sourcing one.

Domestic content and the material-assistance rules are cost accounting on the surface. Underneath they reach into who you buy from, what you make yourself, where you build it and what your supply base is allowed to look like — and none of that can be answered by anyone who does not know the product.

The adviser's memo arrives and it is a good memo. It sets out the thresholds, the formula, the safe harbours, the documentation standard and what your current number appears to be. Then somebody asks the only question that matters — so what do we change? — and the memo has nothing to say, because answering it requires knowing which part in your bill of materials could plausibly come from somewhere else, what the alternative costs, whether it can be qualified in time, whether the supplier can deliver at project quantity, and what moving it does to the rest of your supply base.

That is not a tax question. It is a product, sourcing and footprint question wearing a tax costume, and it is the reason these exercises stall for months in companies that treat them as a compliance filing.

The rule tells you the number. Only the operation can tell you which part to change.

What the test actually reaches into

Two things about the mechanics matter, because they determine where the lever is.

First, it is assessed at the component level and weighted by cost. Where the finished product was assembled is close to irrelevant on its own; "made in the USA" and "final assembly USA" are worth nothing without a cost-weighted breakdown beneath them, and the certificate of origin most suppliers reach for first is not the document the rule asks for.

Second, the production and assembly cost of a manufactured product counts toward the domestic numerator only if *all* of that product's components are domestically sourced. One foreign component disqualifies the whole product's production cost. Which means the lever is not gradual. You do not creep toward a threshold by improving things generally. You move a specific part, or you bring a specific scope in-house, and the number steps.

Identifying which part, out of several hundred, is worth moving is an engineering and sourcing judgement made against a costed bill of materials. There is no version of this where the answer arrives from outside the operation.

First, the test that can disqualify you

Since the One Big Beautiful Bill Act of July 2025 there are two overlapping sourcing tests on the same project. The material-assistance restriction measures the share of cost that does *not* come from a prohibited foreign entity, and failing it disqualifies the project from the credit **entirely**. The domestic content bonus measures the share that is US-origin, and failing it costs **only the bonus**. Same cost basis, broadly the same arithmetic, wildly different downside — so run the disqualifying test first, every time, and treat the bonus as the second question.

What it does to sourcing

This is where the operational weight sits, and it changes four things at once.

Who is on the list. An approved vendor list qualified on capability, quality and price is no longer sufficient; suppliers now also have to be qualified on content and on ownership. Second sources have to clear both bars, which means the qualification programme you had is not the qualification programme you need.

Where the content actually originates. Origin attaches to components, not to your supplier's address. A domestic supplier assembling imported subassemblies moves your number by very little, and finding that out requires a breakdown they may never have been asked for. Expect to send documents back.

What it costs, permanently. The adder between a standard build and a qualifying one is not a one-off compliance cost. It is a change to landed cost that persists for as long as you buy that configuration, and it belongs in the should-cost model rather than in a project memo.

And a diligence activity that did not previously exist. The prohibited-entity tests turn on ownership, control and debt — not on manufacturing location. A supplier can be caught while manufacturing nowhere near the country in question. So procurement now has to know who owns its suppliers, and in places who owns theirs, and has to keep knowing it, because ownership changes and a change can move a supplier across the line mid-programme. Most procurement functions have never done this work and have nobody in post who can. It has to be built: a screening standard, a record, a re-check cadence, and contract terms that oblige suppliers to disclose ownership changes, cooperate with audit, and flow the same obligations down to their own sub-tier.

The part that moves your number most is usually the part with the thinnest domestic supply base.

Which produces the trap worth naming out loud. Qualification and resilience pull in opposite directions. The component that shifts your percentage furthest tends to be a high-value one with a short domestic supplier list, sometimes constrained by an upstream material with a single domestic producer and allocation to match. Qualifying it concentrates your supply at precisely the point where you would otherwise be building a second source. That trade is real, it has a price, and somebody who understands the supply market has to put a number on it before the credit is offered to a customer.

What it does to manufacturing scope and location

The same mechanics reach into the footprint, and they change the shape of the question.

Make-or-buy is now a tax-affected decision. Production cost you incur yourself can count toward the numerator, but only under the all-components rule — so bringing one subassembly in-house is not a marginal improvement, it is a step change that can flip an entire product's production cost into the numerator, or fail to move anything at all if one bought-in part remains foreign. That asymmetry rewards a specific, deliberate scope choice and punishes a general intention to localise.

So the footprint question stops being "where is it cheapest to build" and becomes "what scope, brought into a domestic plant, moves the number by enough to justify the capital and the time." That is a manufacturing strategy question with a tax number attached, and it should be answered inside the location model — as a cost line and as a gate — rather than alongside it.

Two constraints bound the answer. Capacity: an optimised source that cannot deliver at project quantity is not an option, so scope decisions get made against qualification-constrained supply rather than against theoretical availability. And

timing: the thresholds step up by the year construction begins, so the same bill of materials can pass or fail on schedule alone. A slipped construction start is a compliance event, not merely a programme one, and almost nobody treats it as one until it has cost them.

The discipline that makes any of it hold

None of the above survives a weak number. Two conventions do most of the work. Anything without a signed, cost-weighted attestation behind it enters the model at zero per cent — which produces a figure a reviewer cannot improve on, and makes the gap list generate itself, because every zero is either a genuine zero or an outstanding document whose value you can now rank. And the output is never pass or fail. It is the margin above the threshold. A calculation that clears by a tenth of a point is not a position you can hold once a supplier revises a breakdown, a quote becomes a purchase order at a different price, or scope moves.

The cost side comes from a dual quote: standard build and US-content-optimised build of the same product, at the same quantity, priced separately, with capacity and lead time confirmed on the optimised option. A blended number hides the only figure the decision turns on, and an optimised source without capacity is a hope with a price on it.

And the answer might be no

Net benefit is the incremental credit minus the sourcing cost adder, and it is negative more often than people expect. That is a finding, not a failure, provided it lands before the credit has been offered to a customer — which is the argument for running this as a bounded feasibility-and-pricing exercise rather than as the first phase of an execution plan that has already assumed its own conclusion.

A word on where I sit in this

The tax opinion, the credit calculation and the filing position belong to a tax specialist, and any engagement should say so in writing. What I own is everything the memo cannot reach: the bill of materials resolved to component level with costs against it, which parts are substitutable and what qualification would cost, the make-or-buy and scope options and what they cost to stand up, the dual quotes and the capacity behind them, the supplier and ownership diligence, and a defensible net-benefit number in front of the people who have to decide.

The rules are still moving — the material-assistance guidance is interim, its safe harbours expire shortly after the proposed regulations land, and the definitions of which entities are prohibited have been expressly deferred, which is live exposure for anyone with a joint venture, a tax equity structure or a long-term supply agreement in the picture. Check the position at the time; do not carry a conclusion between projects. But the operational work is stable, because it is the same work whichever way the thresholds move: know the product, know the bill of materials, know the supply market, know what you could make yourself, and know what each of those is worth.

That is the work I do.

WHERE THIS GOES NEXT

Where you build it — and how you can defend it →

The manufacturing strategy and footprint method — product, process, make-or-buy, flow, capacity, ownership, geography and organisation, in that order. Content origin belongs inside that model, not beside it. Or start a conversation: rtackconsulting.com/contact

