

The Missing Middle: Where Structural Verification Standards Run Out for Deployable and Assembled Space Structures

Marcin W. Kasiak — PhD, IWE, PE

PSE-TM-04 · Revision A · 15 October 2026 · postsingularity.engineer · mkasiak@postsingularity.engineer

Independent research, carried out on personal time using public data. Views are the author's own and not those of his employer. Not prepared under, and not offered subject to, a professional seal.

Abstract

Structural verification of spacecraft is governed by a process-and-margin regime rather than by a prescriptive code. Factors of safety, test factors and model correlation targets are specified; every consequential acceptance decision is delegated to a project Technical Authority in the NASA system, or settled by customer-supplier agreement under ECSS. That arrangement works well for hardware that can be tested. It runs out precisely where deployable and assembled structures live.

This review maps the governing documents on both sides, separates what is binding from what is convention, and identifies four gaps that matter for structures that unfold or are built in orbit. First, no standard requires structural verification of the intermediate configurations a structure passes through during deployment; mechanism actuation is covered, the structure at each stage is not. Second, joint nonlinearity — free play, backlash, deadband, latch repeatability — has no required characterisation and no required model form, in a class of structure whose global stiffness is dominated by its joints. Third, when full-scale test in one gravity is impossible, the only quantitative lever is a model uncertainty factor whose every reduction is discretionary, while the credibility framework that exists for models is in practice not applied to structural models at all. Fourth, tolerance propagation to deployed shape is handled by worst-case or root-sum-square convention; statistical propagation is permitted everywhere and required nowhere.

Underlying all four is a difference in accountability. In civil practice a licensed individual seals calculations against a prescriptive code and carries personal liability. In space structures an institutional role exercises judgement against a process. Offshore and nuclear engineering sit between the two, with independent third-party verification as a licensing prerequisite. This paper argues that the same instrument — independent verification by a professionally accountable party — is the natural fit for the analysis-only arguments that deployable and assembled structures increasingly require, and sets out what a requirement would have to say in each of the four cases.

1. Two regimes

A long-span bridge is designed against a prescriptive code. Load factors, resistance factors and acceptance checks are fixed in the document; the engineer applies them, seals the calculations, and carries personal liability for the result. A building official issues a permit on the strength of that seal. The code is the argument, and a named individual stands behind its application.

A spacecraft structure is designed against a different kind of document. Factors of safety are specified, test factors are specified, and model correlation targets are specified — but the decisions that actually determine whether a structure is accepted are delegated. In the NASA system they go to a Technical Authority holding an individual warrant delegated down from the Administrator through the Chief Engineer, deliberately independent of programme and budget, supported where necessary by independent assessment and standing review boards. Under ECSS they are settled between customer and supplier. There is no seal, no permit, and no individual carrying personal liability for the calculation.

Neither arrangement is careless. The space regime is a rational response to hardware that is built in small numbers, flown once, and cannot be inspected afterwards: it substitutes process discipline and institutional judgement for codified rules that could not anticipate every configuration. But the substitution has a consequence. Where a prescriptive code runs out, its silence is visible — there is no clause to cite. Where a judgement-based regime runs out, the gap is invisible, because the process still produces an answer. A Technical Authority approves something; the analysis was performed; the review was held. Nothing signals that no requirement existed.

The argument of this paper is that four such silences matter increasingly, because the structures that expose them — large deployables and structures assembled in orbit — are exactly the ones the industry is now building. Each is identified below with the evidence for its absence, and each is illustrated by a case where the missing requirement would have changed what an analysis had to show.

2. The regime mapped

The documents below are the structural core of the two systems. The column that matters is the last one: what each says when analysis is asked to stand in for test.

DOCUMENT	SCOPE	WHAT IT SAYS ABOUT ANALYSIS IN PLACE OF TEST
NASA-STD-5001	Design and test factors of safety	Hardware is normally verified by both analysis and test. A no-test option exists at substantially higher factors, with project-specific criteria submitted for approval by the responsible Center.
NASA-STD-5002	Load analyses of spacecraft and payloads	Correlation thresholds are binding requirements. Model uncertainty factors are prescribed by analysis cycle, and every reduction requires approval by the delegated Technical Authority.
NASA-STD-5017	Mechanisms	Torque and force margin at all points of travel. Verification of margin by analysis alone requires prior approval of the analytical approach and the margin value by the technical authority.
NASA-STD-7002	Payload test requirements	A test matrix marks where analysis may be employed in lieu of testing. Deployment appears as a functional test requirement. Tailoring requires delegated Technical Authority approval.
NASA-STD-7009	Models and simulations	Defines a credibility assessment scale across eight factors and requires results used in critical decisions to carry an uncertainty estimate and a credibility score.
GSFC-STD-7000 (GEVS)	Environmental verification	Verification by test and/or analysis, with modal survey and orthogonality goals for dynamic model verification. Heavily tailored in practice.
ECSS-E-ST-32C	Structural general requirements	Requires a dimensional stability budget in which contributors are identified, analysed and allocated. Stability is to be verified by test; sub-scale representativity must itself be verified.
ECSS-E-ST-32-10C	Structural factors of safety	Deterministic factors, plus a model factor covering uncertainty in mathematical models when predicting dynamic response, loads and load paths. Probabilistic approaches permitted only against a demonstrated reliability objective.
ECSS-E-ST-32-11C	Modal survey assessment	Correlation quality criteria, most of them stated as non-normative examples.
ECSS-E-ST-10-02C	Verification	Defines verification by test, analysis, review of design and inspection, and permits analysis where test is impractical by customer agreement.
ECSS-E-ST-33-01C	Mechanisms	Motorisation factor and uncertainty factors on resisting torque sources.

One defence document is worth naming separately. The US Space Force standard for independent structural loads analysis requires an independent check of models, forcing functions and response calculations, with independently verified mass properties. It is the closest thing in the Western system to institutionalised independent checking of a structural analysis — and it exists on the defence side, not the civil space side.

The pattern across the table is consistent. Analysis in place of test is permitted everywhere, conditioned everywhere on approval, and quantified almost nowhere. The condition is procedural: somebody with authority agrees. What that person should require before agreeing is, with one exception discussed below, left open.

3. What is binding, and what is convention

The distinction between a requirement and a convention is the whole subject of this paper, and it is not always legible from the documents. Three examples set the pattern.

Correlation thresholds differ in status between the two systems. On the NASA side the load analysis standard states them as binding requirements: frequency agreement between test and analysis within five percent for target and significant modes, cross-orthogonality with diagonal terms at or above 0.9 and all other

terms below 0.1. On the European side the equivalent table — mode assurance criterion above 0.9 with frequency deviation under three percent for fundamental bending, relaxing for less significant modes — is explicitly labelled as non-normative, offered as an example of satisfactory correlation, though cross-references make parts of it binding in the narrower context of a reduced test-analysis model. The same numbers therefore mean different things depending on which system a programme is run under.

Effective mass capture is a goal on both sides. Capturing above ninety percent of rigid-body mass in the retained modes is quoted widely and required nowhere as a general condition. For a stiff, compact spacecraft this rarely matters. For a large flexible deployable, where modal density is high and the choice of target modes is genuinely ambiguous, it decides whether the model is complete enough to be correlated at all.

The model uncertainty factor is quantified as guidance and controlled as a requirement. Conventional values run from about 1.4 to 1.5 for an uncorrelated model early in design, through roughly 1.2 at final design, to 1.1 once the structure is built and the model is test-verified. Those numbers appear as guidance. What is binding is the procedure around them: the factor is approved per analysis cycle by the delegated Technical Authority, reduced only gradually as the model matures, and held above the floor if verification is judged inadequate. The standard is explicit that the factor is not a substitute for model verification.

That last point deserves emphasis, because it is the hinge of everything that follows. The uncertainty factor is the system's answer to the question "how much should we distrust this model?" — and the answer is calibrated against how much test correlation the model has received. For a structure that cannot be tested at full scale, the calibration has no anchor. The standard says the factor cannot replace verification. It does not say what to do when verification is impossible.

4. Gap one: the configurations between stowed and deployed

A deployable structure exists in three kinds of state. Stowed, where it survives launch. Deployed, where it works. And a continuum of intermediate configurations in between, each of which the structure must pass through intact.

The standards cover the first two and the mechanism that moves between them. Torque and force margin is required at all points of travel, with prescribed factors on resisting torque sources and a motorisation factor in the European system. A functional deployment test is required, nominally and at high and low energy. What is not required anywhere is verification of the *structure* at each intermediate configuration: that every partial state is stable, that member loads and load paths through the swing are within limits, that the transient at latch is survivable, and that the abort cases — a hinge that binds, a line that releases late — leave the structure in a state it can tolerate.

The omission is not obviously wrong. For a simple two-position mechanism, checking the endpoints and demonstrating the actuator has margin is a reasonable proxy. It stops being reasonable as soon as the structure is long, flexible, multi-stage, or dependent on several hinge lines moving in some relationship to each other — which describes most large deployables.

A worked example makes the size of the omission concrete. In a staged analysis of a hinged multi-panel array, carried through twenty configurations with inertia loads taken from an independent multibody solution, two results appear that no endpoint check could produce. The first is that the peak member loads occur near the end of travel, not at release, and that their magnitude is governed almost entirely by a single modelling choice — whether the damper is represented as a constant torque or as a torque proportional to rate. The constant

representation is correct for the torque margin calculation, which asks what resists at worst case; it is wrong for the dynamics, where the damper exists to cap arrival speed. Carrying the margin-budget value into the deployment analysis oversized the peak member bending by a factor of twenty.

The second is that the worst deployment sequence changes identity depending on that same choice. Undamped, the fastest sequence is worst, because inertia loads scale with the square of rate. With a rate damper fitted, the fastest sequence becomes the least stressed and a sequence in which one line completes early becomes the worst — because a line that latches stops contributing damping for the rest of the deployment. Both statements are true of their own model. Only one is about the structure.

Neither result is reachable from a requirement set that asks for torque margin and a functional test. Both are ordinary questions in staged construction analysis of a terrestrial structure, where every erection stage is checked as a matter of course because a half-built structure is where collapses happen.

5. Gap two: the joints

In a jointed structure the joints set the answer. Their rotational and axial stiffness sets the modes; their free play and clearance set how the structure behaves at low amplitude; their preload and its repeatability set how much of that behaviour returns after the structure has been taken apart and rebuilt. None of these is a material property. All of them are properties of an interface that was made by hand or by a mechanism, and that re-expresses itself every time the structure is loaded.

No standard reviewed here requires that joint stiffness be characterised, that free play or backlash be bounded, that latch repeatability be measured, or that any particular model form be used to represent a nonlinear joint. The European factor-of-safety standard comes closest to acknowledging the difficulty, observing for nonlinear cases that no commonly agreed value can be provided by the space community. Practice fills the gap with linearised joints and a model factor, or with bounding cases at the extremes of an assumed range.

The cost of that convention is measurable. In the author's tolerance study of a flown jointed truss, reported separately in this series, the reassembly scatter measured across repeated disassembly and rebuild was about seventeen times larger than everything dimensional tolerance could account for — and the gap survived the measurement's own confidence interval by an order of magnitude. A tolerance budget built from the drawing package would have predicted six percent of what the hardware actually did. The residue is carried by joint behaviour under load, and the most plausible single candidate is a threaded locking sleeve that the test documentation records as being set by finger pressure, with no wrench, at every reassembly.

That is the shape of the problem in one sentence. The parameter that governed the measured behaviour was not on a drawing, was not toleranced, was not measured, and was not in anyone's model — and no requirement asked for it.

There is a second-order effect worth stating, because it defeats the obvious remedy of modelling joints in more detail. Tolerance sources divide into two classes that propagate differently. Some are set once at assembly and translate directly into geometry: member length, latch position, node machining. Others only manifest when load pushes the joint through its clearance: free play, backlash, deadband, preload. In the study above, seventeen parameters fell in the first class and five in the second — and not one of the five carried a documented value. The class that governs the behaviour is the class nobody writes down.

6. Gap three: credibility when the test cannot be run

The hardest case is the one the industry is moving toward. A large aperture, a long boom, an assembled truss: too big for any chamber, too flexible to hold its shape in one gravity, and existing in its working configuration only after it has flown. The verification argument has to be carried by analysis, and the framework offers two instruments.

The first is the model uncertainty factor, which is calibrated against test correlation and therefore has no anchor when correlation is impossible. The second is the credibility assessment framework defined for models and simulations, which scores a model across eight factors — verification, validation, input pedigree, uncertainty in results, robustness, use history, management and people — and requires results used in critical decisions to be reported with an uncertainty estimate and a credibility score.

That second instrument is, on paper, exactly the right tool. In practice it is not applied to structural models. Its pilot applications were elsewhere, it is widely regarded as burdensome, and structural model credibility continues to be judged by the correlation numbers discussed above — which is to say, by a test the structure cannot undergo. The credibility scale and the structural acceptance workflow have never been joined.

The precedent everyone reaches for is the large infrared observatory whose sunshield could not be tested as a whole in any vacuum chamber. The solution was a series of engineering models at ambient temperature, a third-scale model that fitted inside a chamber, subsystem testing, integrated optical and thermal modelling, and incremental verification as the observatory was assembled, culminating in an end-to-end cryogenic optical test of what could be tested. It worked. But it was approved through a programme's own technical authority and review chain, not against a codified standard for accepting an analysis-only argument. The precedent is a demonstration that such arguments can be made and trusted. It is not a rule for how.

That missing rule is where a verification practice has something specific to offer, because the question is answerable. An analysis-only argument can state what was calibrated and against what data; what was predicted without that data and how well it matched; which mechanisms were tested and eliminated with quantified reasons; what the two independent implementations do and do not both contain; and where the reference itself is uncertain. A study of that form, reported separately in this series, was carried out by the author against a jointed truss flown with measurements taken in both environments: three quantities were predicted after the model was frozen, the strongest to within a tenth of the experiment's own repeatability, and seven candidate mechanisms for an unexplained discrepancy were eliminated with selectivity and driver quantified for each. None of that was required by any standard. All of it is what a credibility argument would consist of if one were required.

7. Gap four: how tolerances are combined

The European structural standard requires a dimensional stability budget: contributors identified, analysed and allocated, with stability verified by test. It is a real requirement and a good one. What it does not do is prescribe how the contributors combine. Neither does anything on the NASA side, where the equivalent obligations are distributed across mission-specific specifications rather than gathered into one clause.

In the absence of a rule, practice uses worst-case stack-up or root-sum-square. Both are defensible for a stack of small, independent, roughly linear contributors. Neither is appropriate when the map from tolerance to deployed shape is nonlinear, when contributors interact, or — most importantly — when one contributor

turns out to dominate through a mechanism nobody anticipated. Statistical propagation is permitted everywhere and required nowhere.

What that costs is best shown by example. In the author's Monte Carlo study over the tolerances of that same jointed truss, five independent channels were sampled separately and together. Cable tension setting carried eighty-seven percent of the variance in the fundamental frequency. All dimensional tolerance together carried thirteen percent. The single dimension in the entire source corpus that carried a drawn tolerance — the only number a draughtsman had written down — carried three tenths of one percent, and halving it moved the answer by a tenth of a percent.

The mechanism behind that result is the transferable part. Seven of the eight bays had no tension adjuster; tension was set by the geometry of an over-centre lock, so the cable's cut length turned directly into tension, amplified by the ratio of axial stiffness to working tension — a factor of four hundred and fifty. A cut tolerance of the same class as the one documented callout produces a tension scatter of four and a half percent. The same tolerance acting through two different load paths differs in effect by three orders of magnitude.

No worst-case or root-sum-square stack-up would have surfaced that, because the amplification is not in the tolerance list; it is in the structure. And the conclusion is the kind that changes a manufacturing decision: tightening the tolerated dimension is worthless, and controlling the cut length of a cable that nobody thought to tolerance is everything.

8. Who is accountable

The four gaps share a root. A prescriptive code makes its silences visible; a judgement-based process does not. Which raises the question of who exercises the judgement, and on what basis.

	LONG-SPAN CIVIL	SPACE STRUCTURES	OFFSHORE AND NUCLEAR
Accountable party	Named licensed engineer	Institutional role: delegated Technical Authority, or customer and design authority	Operator, plus an independent verification body
Basis of acceptance	Compliance with a prescriptive code	Compliance with process and margins, plus project-tailored criteria	Compliance with code, verified by a third party
Personal liability	Yes, licence at risk	No individual seal	Corporate, with the verifier's accreditation at risk
Independent check	Design check or owner's engineer, often statutory	Independent loads analysis on defence programmes; independent assessment bodies and review boards otherwise	Certified verification agent or third-party inspection agency, as a licensing prerequisite
External gate	Building official issues a permit	No equivalent; launch licensing covers public safety, not payload adequacy	Regulator will not license without verification

The middle column is not a deficiency to be corrected by importing the left one. Sealing calculations against a prescriptive code requires a code, and no code could anticipate the configurations spacecraft structures take. The Technical Authority model is a rational alternative: an individual warrant, deliberately insulated from programme pressure, with a formal route for dissent.

But it has a structural weakness that matters here, and it is a weakness of reach rather than of rigour. Technical Authority is internal to an agency. It does not travel to a commercial constellation operator, a small deployables vendor, or a startup assembling structures in orbit — none of whom has one, and most of whom

tailor the standards heavily on the reasonable grounds that requirements written for flagship missions do not fit a satellite built in quantity. What binds them instead is a launch provider's interface document and an insurer's judgement, neither of which asks whether an analysis-only argument is sound.

The right column is the instructive one. Offshore and nuclear engineering face the same problem — structures too large or too consequential to test as built, operated by commercial entities — and they solved it with independent third-party verification made a condition of licensing. The verifier is not the designer, is professionally accountable, and has standing that does not depend on being inside the operator's organisation.

That is the instrument the space sector lacks and the one that fits its trajectory. As deployable and assembled structures move from flagship programmes into commercial fleets, the internal-authority model covers a shrinking fraction of what flies, while the analysis-only arguments get harder. The gap is not in the rigour of the agencies. It is that most of the industry now sits outside them.

9. What a requirement would say

Identifying a gap is cheap. The test of whether one is real is whether a requirement can be written for it that a programme could comply with and an authority could check. Four are proposed below, each deliberately modest.

On deployment configurations. That structural verification be demonstrated at a defined set of intermediate configurations spanning the deployment, not only at the stowed and deployed states; that the set be justified against the kinematics rather than chosen for convenience; that each configuration be shown stable and within limits; and that the credible abort cases — a line that binds, a line that releases late, a latch that does not engage — be shown to leave the structure in a survivable state. This is staged construction analysis, and its terrestrial form has been routine for decades.

On joints. That every joint parameter which materially affects the verified quantities be either measured, bounded by test, or declared as an assumption with its sensitivity reported. Not a required model form — that would age badly — but a requirement that the parameters be named and their influence quantified, so that a reviewer can see which of them the answer rests on.

On analysis-only arguments. That where test cannot close the verification, the argument state explicitly what was calibrated and against what; what was predicted without that data and how it compared; what was independently cross-checked, and specifically what the independent implementations do and do not both contain; and what the analysis does not support. The existing credibility framework provides most of the structure; what is missing is its connection to structural acceptance.

On tolerance propagation. That where the map from tolerance to a verified quantity is nonlinear or where contributors interact, propagation be statistical rather than by stack-up; and that the result be reported as a variance budget identifying which contributors carry the outcome. The budget is the useful artefact: it tells a programme where tightening is worth paying for and where it is not.

The venues for such proposals exist and are active. The European system has a formal change-request route. The American standards ecosystem has consensus committees owning the mechanisms standards. And the in-space servicing and assembly community has, through its consortium, its national implementation plan

and a recent government accountability review, openly identified standards as one of the field's principal gaps — which is an invitation, not merely an observation.

10. Limitations of this review

Requirement and guidance are not always separable from the text. Where normative status could be confirmed from primary text it is stated; where a threshold is quoted without that confirmation it should be treated as convention. The same numerical criterion is binding in one system and an illustrative example in the other, and a reader should check which system governs before relying on either.

Several documents were not read in full. Parts of the consensus-standards and defence-standards sets sit behind paywalls, and characterisations of those rest on scope statements, catalogue entries and cross-references rather than on the clause text. Revisions move; every document cited should be checked against its issuing authority before it is relied upon.

The central claims are absence claims, and absence is harder to prove than presence. "No standard requires X" is supported here by the absence of such a requirement in the documents reviewed, together with the fact that the relevant communities are themselves actively identifying standards gaps. That is reasonable evidence. It is not proof, and a single clause in a document not reviewed would refute any of them.

The illustrative cases are studies on public data, not flight programmes. They were built from published sources, used one article or one declared design each, and stated their own limits. They show that the missing requirement would have changed what an analysis had to demonstrate. They do not establish how often that matters across a fleet.

The proposals in section nine are untested. None has been through a standards committee, where a requirement meets objections about cost, scope and verifiability that a paper does not. The claim here is only that each is specific enough to be argued about, which is the minimum condition for a gap to be real.