PITON Global

EXECUTIVE WHITE PAPER · WP-41

MECHANICAL & PRODUCT DESIGN · PHILIPPINES

The release-readiness standard: the executive economics of mechanical and product design support outsourcing to the Philippines.

An analysis of why models built is a volume vanity metric, how manufacturability and release-readiness — never modeling throughput — decide the true cost of an NPD support operation once tolerance errors, DFM misses, ECO churn, and prototype failures are counted, and the vendor-selection discipline that ships a design that manufactures the first time.

AUTHORS

John Maczynski — Chief Executive Officer
Ralf Ellspermann — Chief Strategy Officer

JULY 2026
MANILA · ESTABLISHED 2001

Contents

ABOUT THIS SERIES — VOLUME 41

Every product-design support proposal is priced per model or per hour, because CAD output is easy to count. This volume is about the only design outcome an NPD program actually depends on: the deliverable modeled to intent, designed to manufacture, and released clean enough that the first prototype fits and the tool cuts without an ECO storm. The gap between building a model and releasing a manufacturable one is where tolerance errors and prototype scrap live. The full series lives at piton-global.com.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors

John Maczynski

Chief Executive Officer

John has bought design-support capacity for three decades, and learned that the fast pod that builds every model and misses manufacturability is the most expensive one a product team can hire — because a tolerance error found at first article is a scrapped prototype, and a DFM miss found at the tool is an ECO storm and a slipped launch. His standing question ignores the model count: of the deliverables you built, how many released clean and manufactured the first time? Applied at PITON-Global, that question is what clients get on every design-support sourcing decision — personally, with no account layer in between.

j.maczynski@piton-global.com
+1 402-598-8740

Ralf Ellspermann

Chief Strategy Officer

Ralf has spent twenty-five years running Philippine engineering-support floors, and he judges a design pod by its release-readiness and first-article-fit rate, not its models-per-week: whether the deliverable was modeled to intent and manufacturable, or thrown over the wall for the client's engineers to fix. The design-support contract in Section 3 codifies what he learned rebuilding pods that modeled fast and released dirty. He now reads design-support providers the way he once ran them — from the ECO log and the first-article-failure rate backward.

r.ellspermann@piton-global.com
Manila, Philippines

PITON Global

ABOUT PITON-GLOBAL

PITON-Global is a buyer-side advisory practice, Manila-based since 2001, that has never operated a delivery floor or taken a placement commission. Its stock-in-trade is knowing the Philippine provider market at the operating level — 1,000+ mapped operations, re-verified continuously, including each design pod's real release-readiness and first-article-fit rate rather than its models-per-week headline — and converting that knowledge into a shortlist of six to ten genuinely qualified partners per engagement. Buyers from Tripadvisor, Garmin, TSYS, HealthPartners, Yetipay, Workrise, and Norman Disney & Young have run their searches through the practice.

SECTION 01

Executive summary

-59%
Fully-loaded cost per release-ready deliverable vs. US onshore

98%+
Release-readiness rate on vetted pods, up from 82–90%

-64%
Reduction in post-release ECO churn on vetted pods

6.3×
First-year ROI, reconstructed case (Section 6)

A design-support engagement is bought on the price per model or hour and paid for on the deliverables that will not manufacture. The pod reports hundreds of models and drawings built at a fraction of the onshore cost; the product team finds the tolerance stack-ups that fail at first article, the DFM misses that surface at the tool, the ECO churn that eats the launch schedule, and its own engineers redoing offshore CAD instead of advancing the program. The pod hit its model count and the program inherited scrapped prototypes, an ECO backlog, and a slipped launch. The gap between building a model and releasing a manufacturable one is where the real cost of cheap design support hides, because a model that will not manufacture is not throughput delivered — it is prototype scrap and a launch at risk.

The Philippines — with a deep, degreed mechanical-engineering talent pool trained on the major MCAD/PLM platforms and Western manufacturing standards — is where design support is delivered with the most release-readiness discipline, because that pool can run the intent modeling, DFM review, and clean release that manufacturability demands. Five findings:

1. **Models built is the industry's most flattering number.** Any pod can push CAD at volume. Vetted pods are measured on outcome: release-readiness, manufacturability, and first-article fit. Weak providers quote a low per-model price the program repays in tolerance errors, ECO churn, and prototype scrap. Section 2 shows the gap.

2. **The release-ready, manufacturable deliverable, not the fast model, is the product.** A design that ships comes from modeling to design intent, reviewing for manufacture, and releasing to standard — not from clearing a CAD queue. Pods that release clean manufacture the first time; pods that model loosely throw tolerance errors and DFM misses over the wall that cost more than the model.

3. **A dirty release cuts both ways — scrap and delay.** A tolerance error scraps a prototype and its tooling; a DFM miss triggers an ECO cascade and slips the launch. Vetted pods hold release-readiness above 98% precisely because both tails cost — the part that will not fit and the schedule that will not hold.

4 **The ECO churn and the scrap, not the model, are what cost.** A deliverable released dirty costs the client-engineer rework, the scrapped prototype and tool, and the launch delay — many times the per-model saving. Pods that release clean cut ECO churn and hit first-article fit, and that protected schedule — not the offshore rate alone — is where the case study's return is built.

5 **Contract on the release-ready deliverable, not the built model.** The contracting failure of design support is paying per model and hoping it manufactures. Vetted deals price against release-readiness floors, DFM-review gates, and first-article-fit targets, making the manufacturable deliverable the pod's problem, which is where it belongs.

The intended reader is a VP of engineering, NPD program lead, or COO buying design support per model and paying for it in tolerance errors, ECO churn, and prototype scrap. Sections 2–4 separate the pods that release ready from the pods that build models; Sections 5–7, the method and the proof.

SECTION 02

The volume mirage: models built versus release-ready deliverables

Figure 1 reads five deliverable classes as release-gate bullet bars: the vetted pod's release-ready result (bar) against the readiness target every class must clear to advance (tick). The models-built number is identical in spirit to a chatbot's containment rate — a figure engineered to look like output while tolerance errors, DFM misses, and ECO churn accumulate underneath.

FIG. 1 — RELEASE-GATE READINESS BY DELIVERABLE CLASS: RESULT VS. TARGET

Bar = vetted-pod release-ready result · vertical tick = readiness target to advance the gate · faint bar = sheet-count vendor

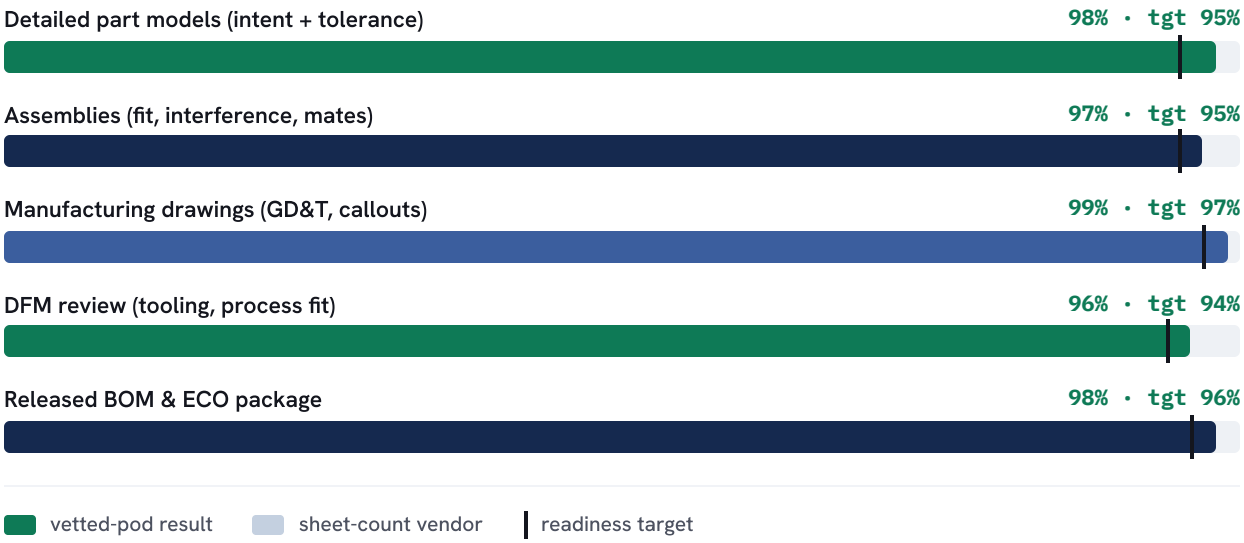

Bars: PITON-Global engagement instrumentation, 2025–26, release logs audited. Sheet-count vendor bars are a composite of pre-engagement reporting reviewed in diligence.

The economics follow the release, not the model count. The vetted pod clears the readiness target on every deliverable class, so parts fit at first article and the tool cuts without an ECO storm; the sheet-count vendor lands well short on every bar — models delivered, but not release-ready — and every gap is a tolerance error, a DFM miss, or an ECO waiting downstream. A low per-model price that releases dirty is not a saving; it is prototype scrap and a launch slip, deferred and compounding. The release-readiness audit — first-article fit, DFM-review pass, ECO churn — is the cheapest due-diligence instrument in this paper.

SECTION 03

The design-support contract: model to intent, design for manufacture, release it clean

Figure 2 runs the NPD deliverable through four stage-gates — the machinery a floor read verifies before any per-model price is believed. A deliverable advances only when it passes the gate; the pod that skips a gate is the one that ships prototype scrap.

FIG. 2 — THE FOUR-GATE DESIGN-SUPPORT CONTRACT, AS OPERATED ON VETTED PODS

G1 — Model to intent

Geometry is built to the design intent and the requirement set, fully parametric and rebuild-clean, with tolerances applied to function — not eyeballed. Gate passes when the model matches intent and the feature tree survives a rebuild without errors.

G2 — Verify fit & tolerance

The assembly is checked for interference, mate integrity, and tolerance stack-up, and GD&T is applied so the drawing is unambiguous. Gate passes when the stack-up closes and nothing interferes across the motion envelope.

G3 — Design for manufacture

The design is reviewed against the target process — moldability, machinability, draft, wall, and tooling — and revised so it can actually be made at cost. Gate passes when DFM review clears with no open manufacturability risk.

G4 — Release it clean

The BOM, revision, and ECO package are released to the PLM standard with a clean audit trail — with a release-readiness report the program can read. Gate passes when the package releases with no open redlines and manufactures the first time.

THE RELEASE-READY DELIVERABLE

98%+ RELEASE-READY | DFM-CLEARED | FITS FIRST ARTICLE

A deliverable that passes all four gates is modeled to intent, tolerance-closed, manufacturable, and released clean — so prototypes fit, tools cut, and ECO churn falls. The operation gets cheaper per release-ready deliverable and faster to launch at the same time, which is the entire promise of outsourcing design support to people who measure release-readiness, not models built.

Operating parameters from PITON-Global engineering-support engagements, 2025–26. Design support runs under the client's engineer of record; nothing here substitutes for the client's design authority.

In diligence, the contract is tested at the gates: hand the pod a representative part and assembly and check whether the model rebuilds clean, the stack-up closes, and the DFM review catches the moldability risk; audit the ECO log and the first-article-failure history. Perhaps one pod in ten survives it cleanly. Steps 3 and 5 of the framework exist to find that one before the contract does — because everything in Table 3 of Section 6 depends on release-readiness, not models built.

SECTION 04

Cost and performance: cost-per-release-ready-deliverable economics and benchmarks

TABLE 1 — FULLY-LOADED COST PER RELEASE-READY DELIVERABLE: THREE OPERATING MODELS (2026)

OPERATING MODEL	COST / DELIVERABLE	VS. US BASELINE
US onshore in-house design engineer	$430	baseline
Low-cost offshore CAD mill (per-model)	$285	−34%
Vetted Manila pod — release-ready, DFM-cleared	$176	−59%

Basis: PITON-Global engagement pricing 2025–26; cost per release-ready deliverable = fully-loaded operating cost divided by deliverables released clean and manufacturable (dirty and reworked deliverables excluded); the CAD mill looks cheap per model but ECO churn, prototype scrap, and client-engineer rework lift its true cost per release-ready deliverable; US comparator fully-loaded in-house engineer cost.

TABLE 2 — DESIGN-SUPPORT BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
Release-readiness rate	98%+	82–90%	The only design number launch follows
First-article fit rate	96%+	78–88%	Fits the first time, or scrapped
Post-release ECO churn vs. baseline	−64%	baseline	"Design for manufacture" — held or not
Design-cycle time vs. baseline	−42%	baseline	To launch faster, or ECO-bound
Client-engineer rework reclaimed	+31%	baseline	Time on the roadmap, or on fixes

Source: PITON-Global engineering-support engagement data 2025–26; baseline = typical pre-engagement operations billed on models built.

The strategic read: the CAD mill's −34% per model is a false economy — ECO churn, prototype scrap, and client-engineer rework lift its true cost per release-ready deliverable far above the per-model gap, and a launch slipped by a DFM miss is a cost no rate recovers. The vetted Manila pod's −59% is real because it is

measured where the value is: the deliverable released clean, DFM-cleared, that fits at first article. Buy on models built and you buy the short bars of Figure 1; buy on cost-per-release-ready-deliverable and you buy release-readiness — the same discipline every volume in this series keeps arriving at from a different door.

SECTION 05

Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to product design

Design-support diligence adds one hard requirement to every step: a per-model price is real only when the release-readiness behind it survives a gate test on a representative part and assembly. The funnel below is representative of a 30–70-seat NPD support search.

FIG. 3 — THE 7-STEP FRAMEWORK AS AN NPD SUPPORT FUNNEL

	STEP	WHAT IT ELIMINATES IN A DESIGN-SUPPORT SEARCH	FIELD REMAINING
1	**Requirements definition**	Model-count scoping: fixes the release-readiness floor, the first-article-fit target, the DFM-review gate, and the ECO-churn ceiling the operation must meet by deliverable class	1,000+
2	**Market mapping**	CAD mills: pods that bill model volume and treat tolerance errors and DFM misses as the client's problem, and generalist shops with no degreed-engineer or DFM depth	~44
3	**Capability screening**	Pods without the machinery: degreed mechanical engineers, MCAD/PLM proficiency, a DFM-review and tolerance-analysis workflow, and coverage across your processes and materials	~13
4	**Compliance & security audit**	Weak IP governance: design-IP and trade-secret protection, PLM access controls, export-control alignment, secure-facility discipline, SOC 2 scope covering the design stack	~7
5	**Operational diligence**	The gate test: hand over a representative part and assembly and check rebuild-clean modeling, tolerance closure, and DFM catch; audit the ECO and first-article-failure logs; interview on a manufacturability edge case	4–6
6	**Competitive RFP & negotiation**	Per-model pricing: finalists bid against release-readiness floors, first-article-fit targets, and DFM-review gates — all carrying credits for deliverables that come back dirty or fail first article	2–3
7	**Transition & launch governance**	Cold start: the pod calibrates against your CAD/PLM standards, tolerancing conventions, and target processes on a pilot part before it owns production, with the release-readiness floor certified before steady state, PITON-Global embedded buyer-side	1

Field-remaining counts represent a 30–70-seat NPD support search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no design floors, builds no models, and takes no placement economics, so the shortlist reflects audited release-readiness performance, not inventory; **right-sizing** — pods for whom your program is a flagship account. Six to ten

genuinely qualified providers, delivered within 48–72 hours, free to the buyer, competing through a fully transparent, fair, comprehensive, and highly competitive RFP.

SECTION 06 · ANONYMIZED CASE STUDY

Re-basing the operation on release-readiness: a 44-seat NPD support rebuild, reconstructed

Engagement MD-069 is a US industrial-products manufacturer — running roughly 21,000 design deliverables a year across parts, assemblies, drawings, and ECO packages — that had offshored design support to a per-model CAD mill. The model price was excellent; the release-readiness was not: release-readiness sat near 86%, first-article failures were common, ECO churn was eating the launch calendar, and the client's own engineers were redoing offshore CAD. The VP of engineering wanted the arbitrage without the scrap and the ECO storm. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–8, GATE TESTS)

Requirements fixed a 98% release-readiness floor, a 96% first-article-fit target, and a DFM-review gate, with a mandatory tolerance-analysis and PLM-release workflow. Mapping produced 44 claimants; the machinery screen — degreed engineers, MCAD/PLM proficiency, DFM review, process coverage — left 13; the IP-and-security audit left 7. Gate tests on a representative part and assembly shortlisted 5. A Manila pod won priced against release-readiness and first-article fit, with a gate re-test written into the reporting spec.

TRANSITION (MONTHS 2–7, CALIBRATED FIRST)

The pod calibrated against the manufacturer's CAD/PLM standards, tolerancing conventions, and target processes on pilot parts before owning production; the release-readiness floor was certified in month 3. The gate-and-DFM loop fed calibration from day one. By month 7 release-readiness reached 98.3%, first-article fit hit 96%, ECO churn fell 64%, and design-cycle time dropped 42%. PITON-Global chaired the release-readiness review board to steady state.

TABLE 3 — MD-069 FIRST-YEAR VALUE BUILD-UP (US$)		
VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 44-seat design operation	$1.58M	82.5k productive hrs × ($29.10 − $9.90)
Prototype scrap & tooling rework avoided	$0.66M	First-article failures removed by higher readiness
Client-engineer rework reclaimed	$0.30M	In-house engineer hours returned from CAD fixes
Launch-schedule & ECO value	$0.20M	Faster launch and lower ECO-handling load
Gross value created	**$2.74M**	
Less: engagement & transition costs	($0.44M)	Calibration, pilot parts, CAD/PLM integration; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	**6.3×**	$2.74M ÷ $0.44M

Figures rounded; readiness, first-article, and ECO effects audit-verified. Method in Methodology & notes.

Fifty-eight percent arbitrage, forty-two percent from avoided scrap, reclaimed engineer time, and faster launch — a typical split for the series, earned because the operation was re-based on release-readiness rather than model volume. Every line traces to a diligence step: the readiness floor to Step 1, the readiness-based pricing to Step 6, the gate workflow to Step 3's contract. The VP of engineering's line at the year-one review: "The model price barely moved. What moved is that parts fit at first article now — and the ECO storm is gone."

SECTION 07

The executive playbook: governance for the first 180 days

A design-support transition succeeds when the pod is measured and paid on release-readiness and first-article fit — client in command throughout.

DAYS 0–30

Define readiness, contract on it

Set the release-readiness floor, the first-article-fit target, and the DFM-review gate and ECO-churn ceiling with engineering leadership before launch, and document the tolerancing conventions and PLM-release standard by deliverable class. Contract against release-readiness and first-article fit — never per model alone — with floors and gates carrying credits. Stand up the gate re-test as the program's arbiter of truth.

DAYS 30–90

Calibrate the pod, pilot before production

Calibrate the pod against your CAD/PLM standards and target processes on pilot parts before it owns production; certify the readiness floor and the four-gate loop. Switch on the gate-and-DFM loop from day one. Publish the design dashboard: release-readiness, first-article fit, ECO churn, design-cycle time, engineer-rework reclaimed.

DAYS 90–180

Own production, certify steady state

Hand production design to the pod only after it holds the readiness floor and first-article-fit target. Review the release board monthly with the ECO log and first-article-failure trend on the table. Certify steady state on two months at floor including a full launch cycle; hand over on documented criteria; keep the quarterly gate re-test standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Paying per model while deliverables will not manufacture · modeling fast instead of to design intent · skipping DFM review until the tool · releasing dirty and letting ECOs pile up · owning production before release-readiness is proven. All five are settled at selection and contract.

The design-support argument, in one sentence: a built model is CAD output and a release-ready deliverable is a part that fits at first article and a tool that cuts clean, and a vetted Philippine pod delivers −59% cost per release-ready deliverable with 98%+ release-readiness and ECO churn cut by nearly two-thirds — for buyers who gate-test the readiness instead of counting models built, and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. Anyone can build a model. We make sure you buy the pod whose design manufactures the first time.

METHODOLOGY & NOTES

How the numbers were built

Readiness and ECO figures (Fig. 1, Table 2) derive from engagement instrumentation in PITON-Global engineering-support engagements, 2025–26. Release-readiness and first-article fit are measured by auditing release logs and first-article results against the deliverable class; ECO churn, design-cycle time, and engineer-rework reclaimed are

measured against the client's pre-engagement baseline. The sheet-count-vendor bars are a composite of pre-engagement reporting reviewed in diligence, shown for contrast.

Cost-per-release-ready-deliverable figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, MCAD/PLM licensing and technology, management and QA overhead including DFM reviewers and checkers, attrition-driven recruiting — divided by deliverables released clean and manufacturable, with dirty or reworked deliverables excluded from the denominator; US comparators reflect fully-loaded in-house engineer cost. A release-ready deliverable is modeled to intent, tolerance-closed, DFM-cleared, and released clean.

The case study (MD-069) is anonymized under NDA; volumes and dates are generalized, figures rounded. Readiness, first-article, and ECO effects are audit-verified; the scrap-avoided line reflects removed prototype and tooling rework, and the launch-value line is a conservative estimate. ROI = net first-year value ÷ all engagement-related client costs. PITON-Global's advisory fee is provider-paid and appears at $0 in the client cost base; neutrality safeguards are documented at piton-global.com.

Market context draws on IBPAP reporting and PITON-Global's proprietary provider mapping (1,000+ operations; release-readiness and first-article performance re-verified May 2026). Design support runs under the client's engineer of record and design authority; nothing here transfers that authority offshore. Estimates and projections are PITON-Global analysis and should be cited as such. The full series is available at piton-global.com.