COST · VELOCITY · OWNERSHIP

Agency vs Dedicated Engineering Pods

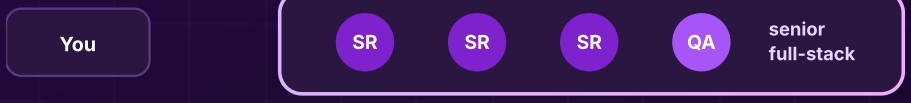
A **quantitative teardown** of traditional agency billing models versus dedicated full-stack pods with **zero junior delegation** — where the money goes, where the time goes, and when each model wins.

AGENCY · HAND-OFF CHAIN

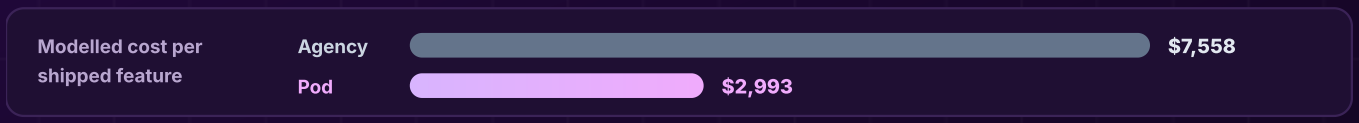


4 hand-offs
every requirement

DEDICATED POD · DIRECT LINE



0 hand-offs
engineers in your channel



–60%
modelled cost per shipped feature with a senior pod

2.1×
more features shipped per month in the base case

26 → 7
calendar days from request to release, modelled

~40%
pod utilisation where the cost advantage disappears

THE MODELS COMPARED

What you're actually buying

Agencies sell hours or outcomes and staff them from a shared bench. Dedicated pods sell a fixed, named team that works only on your product. The billing line looks similar; the delivery mechanics are not.

AGENCY

Three billing models

Time & materials	Hourly rates per role; you pay for every hour logged, including coordination and rework.
Fixed price	Quote = estimate + risk buffer; every scope change becomes a change request.
Retainer	Block of hours per month; unused hours often expire, overflow billed extra.

- Team drawn from a shared bench, rotating across clients
- Seniors sell and review; juniors do much of the build
- Account and project managers sit between you and code

DEDICATED POD

One model: a team, by the month

Composition	2–4 senior full-stack engineers + part-time QA/DevOps.
Billing	Fixed monthly fee per pod; no change requests, scope flexes inside capacity.
Control	You own the backlog and priorities; engineers join your tools and stand-ups.

- **Zero junior delegation** — the people you interview write the code
- Same named engineers month after month
- Direct line from product owner to engineer

INSIDE THIS TEARDOWN

03 **Anatomy of an invoice** — Where each dollar goes in both models

04 **The delegation problem** — Staffing pyramids and hand-off chains

05 **Velocity teardown** — Request-to-release, step by step

06 **12-month cost model** — Every assumption on the table

07 **Sensitivity & break-even** — When the answer flips

08 **Hidden costs** — Six line items that never appear on invoices

09 **Measuring delivery** — Metrics to hold any vendor to

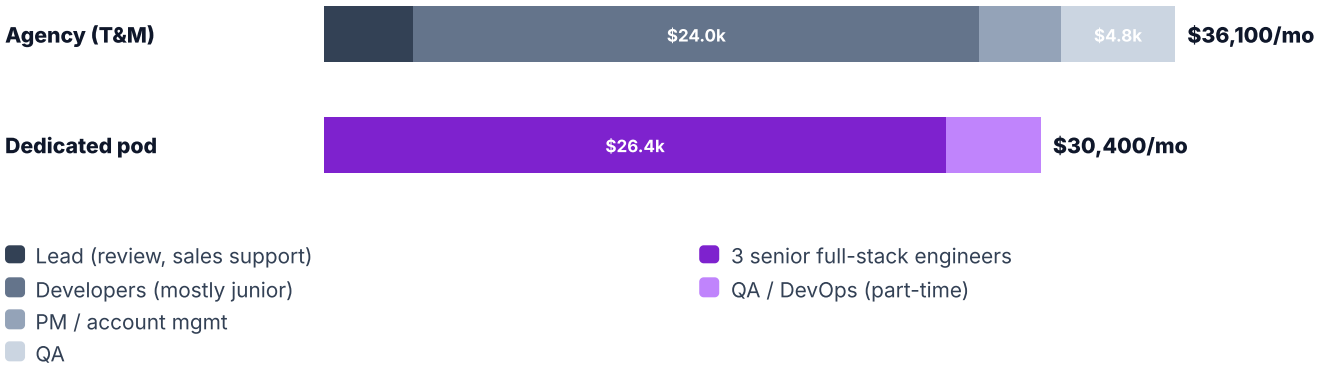
10 **Decide & contract** — Decision matrix and vendor checklist

How to read the numbers: figures come from a transparent, illustrative model — not a survey. Every input is listed on page 6 so you can swap in quotes from real vendors and rerun it.

01 · ANATOMY OF AN INVOICE

Where each dollar goes

Two teams quoted to build the same product for a month. The invoices are close. What they buy is not: the useful question is how much **senior-quality engineering** each dollar produces.



FROM HOURS BILLED TO USEFUL OUTPUT

191 h

senior-equivalent engineering hours per month, from 490 billed hours

320 developer hours × 0.70 junior productivity	224 h
− 12% rework from review findings	197 h
+ lead's hands-on time (half of 40 h)	217 h
− 12% hand-off & coordination loss	**191 h**

406 h

senior-equivalent engineering hours per month, from 560 billed hours

480 senior engineering hours × 1.0	480 h
− 8% rework	442 h
− 8% coordination (stand-ups, reviews)	**406 h**
QA/DevOps hours support quality, not counted	—

$189 vs $75
cost per senior-equivalent hour

4.8 vs 10.2
features shipped per month (40 h each)

$7.6k vs $3.0k
cost per shipped feature

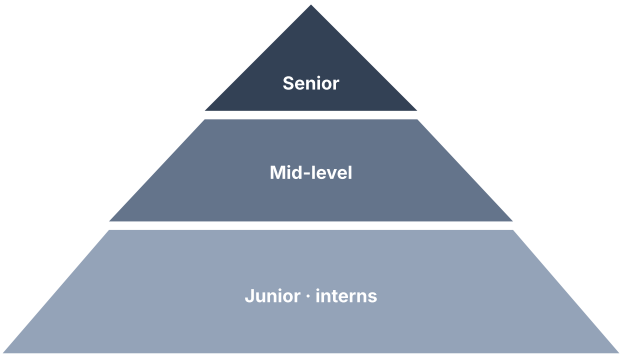
All rates and factors are illustrative assumptions (page 6). Replace them with your real quotes.

02 · THE DELEGATION PROBLEM

Who really writes your code?

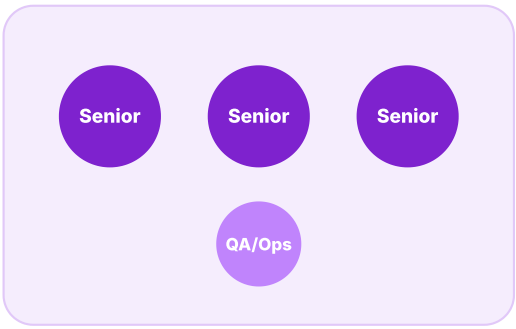
Agency economics reward leverage: one senior's credibility sells the project, several juniors deliver it. That is rational for the agency — and the main reason output per dollar drops for the client.

AGENCY PYRAMID



seen in the pitch ↑ · writing the code ↓

DEDICATED POD



the people you interview write the code

THE HAND-OFF CHAIN, AND WHAT EACH LINK COSTS

HAND-OFF	WHAT GETS LOST	TYPICAL SYMPTOM
You → account manager	Business context compressed into a ticket	"That's not what I asked for"
Account → PM	Priorities re-ranked against other clients	Your urgent fix waits for a sprint slot
PM → tech lead	Requirement becomes an estimate, not a design	Change requests at every surprise
Lead → junior	Architecture intent behind the task	Rework in review; inconsistent code

What "zero junior delegation" means

- Every engineer on the invoice is senior and named in the contract
- No silent substitution with trainees or subcontractors
- Replacement only with equal seniority, with overlap at the vendor's cost

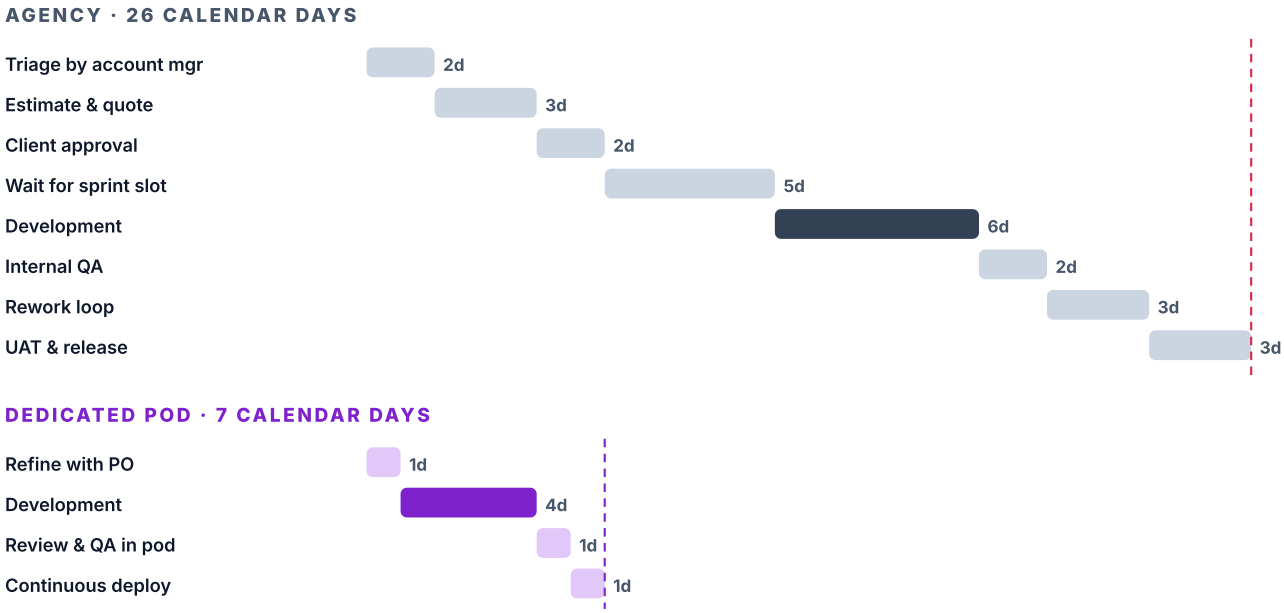
Juniors aren't the problem — hidden juniors are

Mixed teams can work well when you know who is on them and pay junior rates. The cost gap comes from paying a **blended senior rate** for work mostly done by juniors, without visibility into who did it.

03 · VELOCITY TEARDOWN

Request to release: 26 days vs 7

Coding is rarely the slow part. Calendar time disappears in queues, approvals and loops between people. Modelled path of one medium feature:



23%
of agency lead time is actual development (6 of 26 days)

57%
of pod lead time is actual development (4 of 7 days)

3.7×
faster feedback loop — more learning cycles per quarter

Why the agency path is longer

- Every change must be estimated and approved before work starts
- Shared engineers juggle several clients' sprints
- QA and rework happen after hand-off, not during
- Releases batched to reduce coordination cost

Why the pod path is shorter

- No re-quoting: capacity is already paid for
- Engineers talk to the product owner directly
- Seniors review each other the same day
- Continuous deployment ships small changes daily

Step durations are illustrative for a medium feature; measure your own with the metrics on page 9.

04 · THE 12-MONTH COST MODEL

Every assumption, on the table

A 12-month product build. Change any input to match your quotes — the method matters more than our numbers.

Agency (time & materials)

ROLE	HOURS/MO	RATE	$/MO
Tech lead	40	$95	3,800
Developers ×2	320	$75	24,000
PM / account	50	$70	3,500
QA	80	$60	4,800
Total	490		**36,100**

Factors: junior productivity 0.70 · rework 12% · hand-off loss 12% · lead 50% hands-on

Dedicated senior pod

ROLE	HOURS/MO	RATE	$/MO
Senior full-stack ×3	480	$55	26,400
QA / DevOps (0.5)	80	$50	4,000
Total	560		**30,400**

Factors: senior productivity 1.0 · rework 8% · coordination 8% · 100% backlog utilisation

12-MONTH RESULT

METRIC	AGENCY	POD	DIFFERENCE
Total spend	$433,200	$364,800	**−16%**
Senior-equivalent hours delivered	2,293	4,876	**+113%**
Features shipped (40 h each)	57	122	**2.1×**
Cost per shipped feature	$7,558	$2,993	**−60%**

The formula

```
effective_hours = Σ(hours × productivity) × (1 − rework) × (1 − coordination)
features = effective_hours ÷ hours_per_feature
cost_per_feature = monthly_cost ÷ features
```

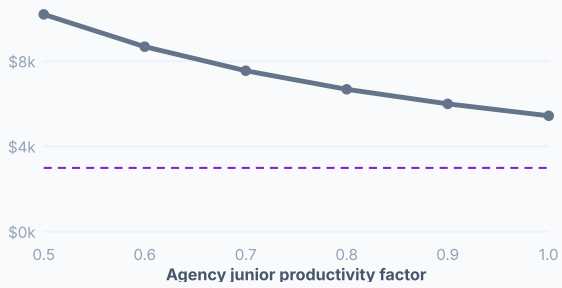
Sources for your inputs: vendor rate cards, staffing plans in proposals, timesheets from past projects, and your own defect and rework logs.

05 · SENSITIVITY & BREAK-EVEN

When does the answer flip?

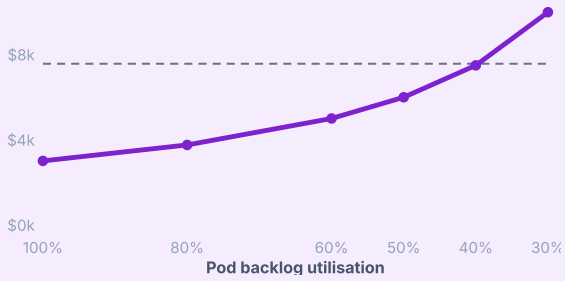
A model is only useful if you know where it breaks. Two inputs move the result most: how productive the agency's delivery team really is, and how busy you can keep a pod.

Agency cost per feature vs junior productivity



- - - pod at $2,993. Even with an all-senior agency team (1.0) cost per feature stays at $5,441 because of PM layers and higher blended rates.

Pod cost per feature vs utilisation



- - - agency at $7,558. Below roughly **40%** utilisation, an idle pod costs more per feature than the agency.

SCENARIOS WHERE EACH MODEL WINS

SITUATION	BETTER FIT	WHY
Small, fully specified one-off (landing site, campaign)	AGENCY FIXED PRICE	No backlog to keep a pod busy; scope won't change
Short burst of specialist work (audit, design sprint)	AGENCY T&M	Pay only for the hours; no team to sustain
Product with an evolving roadmap for 6+ months	POD	High utilisation; learning compounds in one team
Scale-up replacing or extending an in-house team	POD	Works inside your process and tools
Unclear requirements, frequent pivots	POD	No change-request tax on every pivot

WHAT-IF SCENARIOS (COST PER SHIPPED FEATURE)

CHANGE ONE INPUT	AGENCY	POD	GAP
Base case	$7,558	$2,993	**−60%**
Pod senior rate rises from $55 to $75/h	$7,558	$3,937	**−48%**
Agency staffs all-senior (productivity 1.0)	$5,441	$2,993	**−45%**
Pod only 50% utilised	$7,558	$5,986	**−21%**

06 · HIDDEN COSTS

Six costs that never appear on an invoice

Rates are easy to compare. These are not — but they often decide the real total. Estimate each one for your situation before signing.

1 Change-request tax

Fixed-price quotes carry a risk buffer, then charge again for each change. Track: CR value ÷ original contract value.

Pod: scope changes absorbed within monthly capacity.

2 Knowledge loss from rotation

Each new person on a rotating bench re-learns your domain. Track: ramp-up days × number of rotations per year.

Pod: named engineers stay; domain knowledge compounds.

3 Context switching

Engineers split across clients lose focus time at each switch. Ask: how many clients does each engineer serve this month?

Pod: one product, full attention.

4 Your own management time

Hours your team spends writing specs, chasing status and checking deliverables. Track: internal hours × internal cost.

Pod: lighter specs; status visible in your tools.

5 Quality debt

Code that passes acceptance but is hard to change raises the cost of every future feature. Track: defects found after release, time to onboard a new engineer.

Pod: seniors own long-term maintainability.

6 Lock-in & exit cost

Missing documentation, vendor-owned infrastructure or unclear IP make switching expensive. Ask: could another team take over in 30 days?

Pod: your repos, your cloud, your docs from day one.

Add them to the model

True cost = invoices + change requests + (rotations × ramp-up cost) + internal management hours + post-release defect fixing + expected exit cost. Even a rough estimate of these six lines usually moves the decision more than a $5 difference in hourly rate.

HIDDEN-COST WORKSHEET — FILL IN FOR YOUR PROJECT

LINE ITEM	HOW TO ESTIMATE	YOUR 12-MONTH $
Change requests	Expected CRs × average CR value	$ _____
Rotation & ramp-up	Rotations × ramp-up days × daily cost	$ _____
Internal management	Hours / week × 52 × internal hourly cost	$ _____
Post-release fixes	Defects × average fix hours × rate	$ _____
Exit / handover	Probability of switching × handover cost	$ _____

07 · MEASURING DELIVERY

Hold any vendor to these numbers

Whatever model you choose, measure outcomes, not hours. The four DORA metrics are a widely used standard for software delivery performance; add two commercial ones.

SPEED

Deployment frequency

How often working changes reach production. Target for a pod: several times a week.

SPEED

Lead time for changes

Commit to production. Plus our request-to-release measure from page 5.

STABILITY

Change failure rate

Share of releases causing an incident or hotfix.

STABILITY

Time to restore

How long a production incident lasts before it is fixed.

COMMERCIAL

Cost per shipped feature

Monthly cost ÷ features accepted in production.

COMMERCIAL

Rework ratio

Hours on fixes and redo ÷ total hours billed.

A MONTHLY VENDOR REPORT — ONE PAGE

METRIC	THIS MONTH	LAST MONTH	TARGET	TREND
Features accepted in production	11	9	≥ 10	▲
Median request-to-release	6 days	8 days	≤ 7 days	▲
Deployments	24	19	≥ 12	▲
Change failure rate	4%	6%	< 10%	▲
Rework ratio	7%	9%	< 10%	▲
Cost per feature	$2,760	$3,380	trend down	▲

Example values for illustration.

Ask for it in the contract

Make the monthly report a deliverable, sourced from your own Git, CI and issue tracker — not from the vendor's spreadsheet.

Red flag

A vendor that reports only hours logged and "tasks completed" — and resists measuring what reached production.

08 · DECISION MATRIX

Which model fits you?

Answer eight questions. Count your “yes” answers.

01	Will you need new features continuously for 6+ months?	Yes No
02	Is the scope likely to change as you learn from users?	Yes No
03	Do you have a product owner who can steer weekly?	Yes No
04	Does the product need deep domain knowledge to build well?	Yes No
05	Is speed of iteration a competitive advantage for you?	Yes No
06	Do you want to own repos, cloud and know-how outright?	Yes No
07	Could you keep 3 senior engineers at least half busy?	Yes No
08	Is the budget a stable monthly line rather than one project?	Yes No

0–2

Agency project

A well-defined, fixed-price engagement is likely cheaper. Insist on named staff and a clear acceptance spec.

3–5

Start small, then decide

Run a 4–6 week pilot with a 2-engineer pod and compare cost per feature against an agency quote.

6–8

Dedicated pod

Your work is continuous and evolving — the pod's compounding knowledge and speed will pay off.

Hybrid is allowed

Many companies run a dedicated pod for the core product and use agencies for one-off specialist work — brand design, a marketing site, a security audit. Choose per workstream, not per company.

DESIGNING A FAIR PILOT

Same backlog

Give both options comparable, real features — not a toy project.

Same measures

Cost per accepted feature, request-to-release days, defects after release.

Clean exit

4–6 weeks, code in your repos, no lock-in if you stop.

09 · VENDOR EVALUATION & CONTRACT

Questions to ask before you sign

People

- [ ] Who exactly will write the code? Names and CVs.
- [ ] Will the people we interview be the people who build?
- [ ] How many other clients does each engineer serve?
- [ ] Replacement terms: notice, overlap, equal seniority?
- [ ] Any subcontractors? Disclosed in writing?

Delivery

- [ ] How do you report progress — hours or production outcomes?
- [ ] Can we see your Git and CI history on a past project?
- [ ] What's your test coverage and review practice?
- [ ] How fast can a typical change reach production?
- [ ] Who owns architecture decisions?

Commercials

- [ ] How are scope changes priced?
- [ ] What's included: QA, DevOps, PM, hosting?
- [ ] Notice period to scale the team up or down?
- [ ] Pilot period with a clean exit?
- [ ] Rate review terms after 12 months?

Ownership

- [ ] All IP assigned to us on payment?
- [ ] Code in our repositories from day one?
- [ ] Infrastructure in our cloud account?
- [ ] Documentation and runbooks as deliverables?
- [ ] Handover support if we end the contract?

POD ONBOARDING: THE FIRST 90 DAYS

DAYS 1–30
Embed
Access to tools, codebase audit, first features shipped in week 2, baseline metrics.

DAYS 31–60
Accelerate
CI/CD and testing gaps closed, steady weekly releases, first monthly report.

DAYS 61–90
Own
Pod proposes roadmap improvements; cost per feature compared to baseline.

Red flags in proposals

- "Team to be assigned after signing"
- Single blended rate with no staffing plan
- Unlimited revisions promised on fixed price
- Code hosted only in the vendor's accounts
- Progress reported only as hours logged
- Long lock-in with no pilot or exit clause

Steigern Tech

STEIGERNTECH.COM

BUILD WITH A DEDICATED POD

Senior engineers. Your backlog. No hand-offs, no hidden juniors.

Steigern Tech provides dedicated full-stack pods that work inside your tools and process — from web and mobile apps to ERP and AI implementation.

Dedicated full-stack pods

Named senior engineers on your product only, scaling up or down with notice.

Pilot before you commit

A short pilot with clear outcomes, so you can compare cost per feature with your current setup.

Transparent delivery metrics

Monthly report from your own Git and CI — features shipped, lead time, quality.

You own everything

Code in your repos, infrastructure in your cloud, documentation as a deliverable.

WHATSAPP

+91 94275 72571

EMAIL

info@steigerntech.com

WEBSITE

steigerntech.com

© 2026 Steigern Tech LLP. All figures come from an illustrative model with stated assumptions; they are not survey data or a quote.

DORA is a research programme now run by Google Cloud.