GraphOrbis **CAD Engine**

GraphOrbis is a parametric CAD engine you **embed in your own product** — sketching with a constraint solver, parametric solids, assemblies, 2D drawings, and DXF/DWG interchange, behind a REST API and TypeScript SDK. You license it per **concurrent session** — actual modeling workload — not per named seat. Typical enterprise utilization is 40–60% of seats, so session pricing runs **60–75% below per-seat incumbents** (AutoCAD 2026 ≈ $2,310/seat/yr; Onshape Pro ≈ $2,500/seat/yr).

What you pay for: a session, not a seat

A **session** is one active, concurrent modeling session — a live connection holding an open document while someone designs. It opens on **createSession** and auto-releases after **15 minutes idle**. You're billed on the **month's peak** number of sessions open at the same instant — **not** registered users, saved designs, API calls, or the developers who build your product. Rule of thumb for a live app: peak sessions ≈ your monthly active designers × 1–5%.

Tiers

TIER	BEST FOR	PLATFORM BASE	PER SESSION / MO	CONCURRENCY
Developer Preview 30-day Docker trial	Evaluation	**Free**	—	2
Startup OEM Early-stage ISVs, <$5M ARR	Getting to market	**$2,400/yr** first year waived	**$18**	≤ 25
Growth OEM Scaling ISVs, $5–50M ARR	In market	**$24,000/yr**	**$14 / $11** graduated	≤ 200
Growth Plus OEM to 1,000 High-engagement mid-market	Scaling hard	**$24,000/yr** same base — no step	**$14→$5** graduated	≤ 1,000
Enterprise OEM >$50M ARR / on-prem	Regulated / at scale	**$120,000/yr**	**$8–10** negotiated	200+
Reseller / Marketplace PLM & platform partners	Sub-licensing	**Rev-share**	10–18%	—

Growth Plus graduated rate: 1–100 **$14** · 101–200 **$11** · 201–350 **$9** · 351–600 **$7** · 601–1,000 **$5** per session/mo. Above 1,000 → Enterprise. Growth includes assemblies & constraint API, STEP/IGES pack, SSO (SAML/OIDC), 99.9% SLA and quarterly reviews; Enterprise adds on-prem/private-cloud, source escrow, SOC 2, 99.95% SLA.

Scaling to 1,000 sessions

PEAK CONCURRENT	ALL-IN / YR	PER SESSION/MO	SHARE OF A $4.5M-ARR APP
200 (Growth cap)	$54,000	$22.50	1.2%
350	$70,200	$16.71	1.6%
600	$91,200	$12.67	2.0%
1,000	$115,200	$9.60	2.6%*

The rate is engineered to stay at or below **~2% of your product's ARR** through 1,000 concurrent sessions — a small, predictable share as you scale. *At a fixed seat count, 1,000 concurrent means >30% of seats live at once; usage-based revenue grows with it, keeping the share low.

Worked example — SMB ISV

A web "Designer" on Azure, built by 4–5 developers, Startup tier. Cost tracks peak simultaneous designers — the developers add $0.

YEAR	PEAK SESS.	BASE	YEAR TOTAL
1 · build + beta	~8	$0*	≈ $1,700
2 · in market	~15	$2,400	≈ $5,600
3 · scaling	~22	$2,400	≈ $7,150

Three-year total ≈ **$14,500**; the 4–5 developers add **$0** as headcount. *First-year base waived. Sustained >25 → move to Growth.

How it's metered

The engine enforces the unit it bills: a per-tenant concurrent-session cap at session-create (a hard limit at your tier), 15-minute idle auto-release, and daily peak tracking. Each month the peak concurrent count is reported to billing and invoiced — base plus metered sessions. Licensing is validated continuously; you host the engine as source + SQL under your own brand.

Request an evaluation build → celestryll.com/graphorbis.html · inquiries@celestryll.com