Portfolio Brief

Autonomous kernel · Operator architecture · 3D engineering agent benchmark

A self-built system by an independent developer. This 6-page brief is the primary submission; the full edition runs 18 pages.
It argues with only two kinds of material: **structures you can open** and **numbers you can recompute**.
No demo posture, no screenshots standing in for results, no estimates filling gaps.

Narrow, in the middle

No upstream generation engine, no downstream scene showroom — only the two mandatory checkpoints of a 3D production line: agent capability benchmark and asset intake quality gate.

Evidence can be rechecked

Every key claim carries a URL or a ledger pointer. What the GPU-less environment cannot measure is stated as not done, never written as "supported".

Boundaries stated first

Positioning, structure and benchmark method are public. Core decision algorithms, weights and client data are not.

27 / 10	172	118	4	110,557	7
operators / topology layers	meta-rules	skill packages (46 self-built)	Apache-2.0 engines	cloud truth records	federated nodes, RPO 1h

Structure of this brief

Page	Section	What you are looking at
P.01	Cover and stance	Position, three principles, six metrics, fixed public wording
P.02	What this system does	Positioning logic, six asset domains, three things not done, platform fit
P.03	Hard evidence · 3D agent benchmark	Six-task design, scores, discriminative validity, paired quality gate
P.04	Hard evidence · operators and engineering	Ten-layer topology, three real incidents with root cause and verification
P.05	Inventory and verifiability	Four open-source venues, scale counts, mirror exposure tiers
P.06	Boundaries and verification	What is not available, five self-serve channels, how to verify further

Fixed public wording: 172 meta-rules (38 of them registered in a queryable registry) + 8 categories of system documents + 6 technical articles + 3 commercial proposals + 15 skill packages. Kernel snapshots, rule text details and credentials are excluded.

Tan Bopiao · independent system architect, AI agent product manager

I What this system does

> **In one line** Turn "AI produced something and nobody checked it" into an executable gate: no upstream generation, no downstream showroom — only **capability benchmarking** (can the agent actually do the job) and **intake quality gating** (can the output enter the pipeline), plus the autonomous kernel and operator architecture that carry them.

Asset domain	What it contains	How it can be checked
Kernel & operators	ZONGYUAN-ROOT autonomous kernel, 27 operators across a ten-layer dependency topology, three-dimensional steady-state decision weights	Topology table, per-layer outputs, ledger pointers
Product lines	Huodou Cloud AIOS enterprise agent platform; Kunlun Dongtian multimodal content engine	Live services, 60+ running service list
Open source	4 Apache-2.0 zero-dependency engines, publicly clonable from four venues	Repository URLs, passing case counts
Knowledge assets	110,557 cloud truth records; 26 bound whitepapers; site index of 155 pages	Online search, index page counts
Evaluation assets	ExamForge item pipeline; six-task 3D agent benchmark with a public leaderboard	Item packages, scoring sheets, leaderboard page
3D & content	GLB intake quality gate, asset library, industrial digital-twin viewer	7 online pages, all HTTP 200

Three things deliberately not done

No demo posture	A deliverable either runs, opens, or recomputes. Concept decks and screenshots are not results
No outsourced judgement	Scoring is not "looks fine": binary criteria per checkpoint, so a third party recomputes the same score from the same sheet
No boundary swallowing	Missing capabilities are written in the open (chapter 17 of the full edition); reviewers can question them one by one

Fit against six submission platforms

Xpert (ByteDance), Alibaba Siriser and Tencent ExpertHub weigh **item design and recomputable criteria**; overseas platforms such as Mercor weigh English-readable structured material; TalentsAI, BElite and Yimian Qianshi weigh **real engineering history and checkable outputs**. This brief is organised as method — evidence — boundaries, so each of those three asks has a landing point.

How the system is layered (four tiers, the lower the harder to change)

Tier	Content	Mutability and check
Zero-state base	System-wide anchor and entitlement identifier	Immutable, traceability only
Meta-constitution	Three axioms: internal/external isolation, append-only ledger, single-source governance	Change requires version locking
Meta-rules / rules	172 meta-rules (38 registered in a queryable registry) and layered rule sets	Append-only, numbered, searchable
Manifestation	Product lines, pages, content and 3D assets as replaceable instances	Replaceable; destroying instances leaves the本体 intact

Numbers on this page are measured: 27 operators / 10 layers come from the topology definition; 172 meta-rules from three whitepaper editions, of which 38 are registered; 110,557 from the cloud truth store (collected 2026-10); 60+ counts live running services.

II Hard evidence: 3D engineering agent benchmark

BENCHMARK

It measures judgement and decision-making, not code arithmetic. All six tasks come from failures actually hit during generation and intake; the noise is what the production environment already contained, not traps built by hand.

Task	Scenario	The real failure
T1	Robot arm sunk into floor	Origin never placed on the ground plane; the whole rig sinks through the floor
T2	Intake pipeline	Validation order mis-sequenced; compliant parts rejected in batch
T3	End-to-end loop	Deciding which steps to keep across detect → repair → export
T4	Lost texture path	Relative path broken; surfaces render fully white
T5	Inverted normals	Face orientation flipped; the whole surface goes black under light
T6	URDF out of range	Joint parameters exceed limits; the assembly falls apart

Scores (max 100)

Model	T1	T2	Note
minimax-m3	40.4	55.0	Most stable on long items; streaming call
Spark-X2.6	16.0	55.0	Enough length, incomplete criteria
GLM-4-Flash	0	37.5	Missed the key criterion on T1
telechat3-438b	0	42.5	T1 empty response, counted as missing, never estimated from an average

Discriminative validity (T5 probe, strong-model ceiling 28)

27	26	25	24	23
gemini-2.5-flash	gpt-4o-mini	minimax-m3	gpt-3.5-turbo	qwen-flash

Five models spread across a 23–27 band instead of flattening, which shows **the items separate strong from weak**. This is the claim reviewers probe hardest and the hardest one to fabricate. A flat band fails the items and forces a rewrite; that rule is written into the scoring process.

The paired intake gate The same judgement applied on the asset side: six-dimensional structural inspection (18 industrial GLBs all PASS) + three real repairs (origin lift, normal recomputation, bounding-box normalisation) + a geometric topology proxy in a GPU-less environment (zero degenerate faces, zero non-manifold edges).

Three preconditions that let a third party reproduce the scores

Frozen item set	Task stems, clues and criteria lists fixed at a version; any change records a new version
Answers on file	Raw response and latency kept per task (empty responses included), so a reviewer recomputes independently
Published scorer	Criterion-by-criterion comparison; the same answer scores the same twice; no subjective weighting

Boundary: the sandbox has no GPU and WebGL is disabled, so pixel rasterisation was not run; visible holes still need a rendered visual check, and this brief does not claim “rendering-level QA supported”. The leaderboard and all 7 online pages were measured at HTTP 200.

III Hard evidence: operator architecture and engineering

Layer	Responsibility	Representative operators	Checkable output
L1	Truth filtering and distillation	P4 truth reconciliation · P7 external anchoring · distillation	Purity score, compression ratio
L2	Manifold metrics and steady-state mapping	Riemann metric · SM-BS bidirectional mapping · drift patrol	Drift rate, three-level alerts
L3	Knowledge-graph ordering	Entity-relation extraction · graph completion · cross-document linking	Triples, density and connectivity
L4	Causal reasoning	Causal chain tracing · singularity prediction · do-calculus intervention	Causal paths, four-level warnings
L5	CTE trinity adaptation	CausalToTruth · TruthToEvolution · EvolutionToCausal	Loop closure, fit score
L6	Structured archiving	Four-layer split · nine meta-classes · Merkle-DAG append check	Standard object, root hash
L7	Asset reconciliation	Resource side ↔ M9 ledger bidirectional tracing	Difference report
L8	Security and entitlement	eFuse branch isolation · ZKP integrity check	Isolation result per branch
L9	Decision and compliance shield	Option appraisal · risk grading · structured opinion	Three-dimensional steady-state score
L10	Business entity constraints	Simulation constraints · character meta-rules · keyframe consistency	Simulation snapshot

Three hard scheduling rules: run strictly along the topology; if an upstream output is not consumed downstream, the chain is declared broken and alerts fire; decision weights are fixed at **benefit 40% / risk 35% / cost 25%**, with risk holding a veto.

Three real production incidents (root cause and verification)

Incident	Root cause	Fix and verification
Intermittent connection failures	nginx worker_connections 2048 exceeded the fd soft limit of 1024	rlimit raised to 8192; 89 of 90 samples succeeded (98.9%)
Unbounded memory on full traversal	Recursive walk with no ceiling, 311MB per 4s round	Incremental traversal + cgroup 512M; memory 97% → 58%
Node listed online but unresponsive	Heartbeat wrote events only, never wrote back the registry timestamp	Timestamp written inside the heartbeat; 12 nodes restored, false-liveness to zero

Measured cloud baselines (collected 2026-10)

0 ms local knowledge-card answer	**~0.5 s** semantic retrieval round trip	**181 ms** gateway reporting	**4.8 / 5.0** steady-state score (12 dims)

All three incidents are real online events with before/after sample records; what could not be measured is marked N/A rather than filled with estimates. The 4.8/5.0 steady-state score is an internal 12-dimension benchmark, not third-party certification.

IV Inventory: can you verify it yourself

Open source: four zero-dependency engines, four clonable venues

Component	Verification	Repository
Daotong operator set v1.1	19 operators, 100+ cases all green	github.com/tanbopiao/daotong-ops | atomgit.com/zongyuangen/daotong-ops
Yuheng v0.1.0	5 modules, 10 cases	gitee.com/huodou-cloud-intelligence-aios/daotong-ops
Xuangui M1	22 cases	modelscope.cn/models/zongyuanroot/daotong-ops
Huaxia numerics kernel v0.4	25 cases; seven anchors off by about 9 minutes	Same four venues, released in sync (Apache-2.0)

Scale (countable definitions only)

Skill packages	118 skill directories = 46 self-built + 72 platform, of which 15 exist in a shippable form
Knowledge and docs	26 bound whitepapers; site index of 155 .html pages; 110,557 cloud truth records
Runtime	60+ cloud services; 529 online HTML pages; 7-node federated cluster, RPO 1h
3D assets	18 industrial test pieces (all PASS); 17 of 25 Kunlun ritual pieces (8 in progress); early工程 assets 48 repaired + 17 raw
Operations base	Shared brain: 65 ledgers (meta-rules 38 / capabilities 36 / protocols 27 / breakpoint SOPs 20 / components 9)

Mirrors and exposure tiers (primary / cold / deliberately shrunk)

Public primary	AtomGit, 384 files / 21 MB, publicly clonable
Cold storage	Tencent CNB, 1,875 files; GitHub 356 files (private)
Public docs repo	ModelScope, 16 files (813 MB media commit removed, down to 528 KB)
Deliberate shrink	216 IP keyframes pulled from the public repo into a local vault; removal is not loss, it is exposure reduction

Three reading paths, matching what a platform asks for

What they weigh	Where to read it
Item design and criteria	P.03 six-task design and score ledger; criteria are recomputable (binary judgement per checkpoint)
Real engineering history	P.04 three incidents with root cause and before/after data; P.05 runtime scale and live pages
Checkable output	P.05 four repository URLs with passing case counts; P.06 five channels, each openable on its own

Skill counts are directory counts (46 under .user_skills, 72 under platform .skills). “26 bound whitepapers” and “155 index pages” are two different definitions, listed side by side and never merged. Triple and ledger counts are as of this scan and move with each write.

V Boundaries and verification entry points

Stated plainly: what is not available

Hands-on dual-backend kernels	CUDA / NPU custom kernels were never written first-hand; involvement is at the item-design level and does not substitute for actual pass development
Compiler passes	MLIR / TVM passes never completed independently; only defect localisation and item design
Rendering-level QA	No GPU in the sandbox; pixel-level rendering and hole inspection not run online
English edition	This document is the English brief; the full English edition is still being assembled
Third-party assessment	4.8/5.0 is an internal benchmark, not an external certification

Five channels you can open yourself

#	Channel	What you will see
1	modelscope.cn/models/zongyuanroot	Whitepaper repo and open-source components, publicly readable
2	atomgit.com/zongyuangen/huodouai	Public primary repository, 384 files clonable
3	www.huodouai.com/cyberforge/leaderboard.html	Public 3D agent leaderboard and scoring criteria
4	www.huodouai.com/whitepaper/	Whitepaper index, 155 pages openable one by one
5	www.huodouai.com/omega-brain.html	System overview visualisation, structure on one page

If you need to verify further

Any claim can be matched with raw records on request: sampling logs from incident handling, raw benchmark answers and score ledgers, per-asset inspection reports, Merkle receipts from the ledger. What is handed over is **reproducible steps and receipts**, not core algorithm weights, client project data or credentials. Item-based collaboration additionally requires confirming whether task ideas may be shown in a personal portfolio; that is negotiated per platform rule.

Three kinds of work I can take on

Direction	Existing basis	To be agreed
Item design and evaluation	Six-task benchmark, scoring sheets, scoring runner and ledger	Copyright and portfolio-display rights
Intake quality gate	Six-dimension inspection, three repairs, GPU-less proxy checks	Integration form (plugin / API / offline batch)
3D asset pipeline	Asset library, batch ledgers, 7 live pages	Target engine and acceptance criteria

Entitlement: DID-BR-000002 | system anchor $\Omega_0 \subset \odot \subset \infty \subset \Omega$. All series, concepts, protocols and assets in this portfolio were originated and led by Tan Bopiao; key results carry append-only hash-ledger receipts and can be traced item by item. This is the 6-page English brief; the full edition has 18 pages and 26 chapters. Numbers follow the fixed public wording and are not re-added or merged. Version V1.0-EN.