


EON Switchboard

Best value money can buy. Frontier quality where it matters.
Security you never compromise on the customer's behalf.

The intelligent model-routing and governance layer beneath every EON product

July 2026



The ground has shifted under enterprise AI



The frontier is slowing — and going dark

Top-tier models are being pulled under export controls. Fable 5 and Mythos were suspended days after launch; GPT-5.6 was held back.



Open weight is within a hair

Open models now sit within ~0.2 points of the frontier on core benchmarks — at up to 30x lower cost per token.



Routing is now standard infrastructure

Teams running a routing layer report 40–85% lower bills with no quality drop. Gartner calls AI gateways critical infrastructure.



The moat moved

Progress is shifting from the model itself to proprietary data and the work it represents — exactly where EON lives.

Renting one frontier model for everything is the wrong architecture



Expensive at scale

You pay frontier prices for the majority of tasks that never needed a frontier model.



Brittle & single-vendor


One outage or a suspended model takes you down. Fable 5 vanished days after launch.



A sovereignty risk

Large enterprises won't accept dependency on models they can't govern — and can't prove what ran where.

At Work-Intelligence scale — millions of inferences across many products — this model breaks.



THE CATEGORY INSIGHT

Work Intelligence and physical AI need more than an LLM

Traditional LLMs answer questions. They don't understand the work — the equipment, the procedure, the competence record. That takes a composable ontology plus orchestration across many models, gated for safety. The model is a commodity input. The understanding is the product.

Orchestrate the commodity. Own the compound.

INTRODUCING

EON Switchboard

The intelligent routing and governance layer beneath every EON product. It takes each task, makes a fast judgment, and sends it to the right model — under firm rules you set — so customers get the best value money can buy without ever trading away quality or safety.



Value

The cheapest model that clears the bar — every time.



Quality

Frontier where it matters; nothing ships below the line.



Safety

Provenance and data posture are gates, not preferences.

Three things we never trade against each other



Value

Every task runs on the cheapest model that clears the quality bar for that job. Up to 85% lower cost on customer inference.



Quality

The build plane is always frontier. Nothing customer-facing ships below the quality line, regardless of price.



Safety

A model's provenance and a customer's data posture are hard gates evaluated before cost — never traded for a cheaper route.

Other routers optimize cost and quality. We add safety as a first-class, customer-controlled dimension.

Two planes, one principle



BUILD PLANE

Creating the products

- **Always best-in-class frontier.**
- Never cost-optimized — cheap models here are false economy.
- Opus 4.8 / GPT-5.5 / Fable 5 when available.



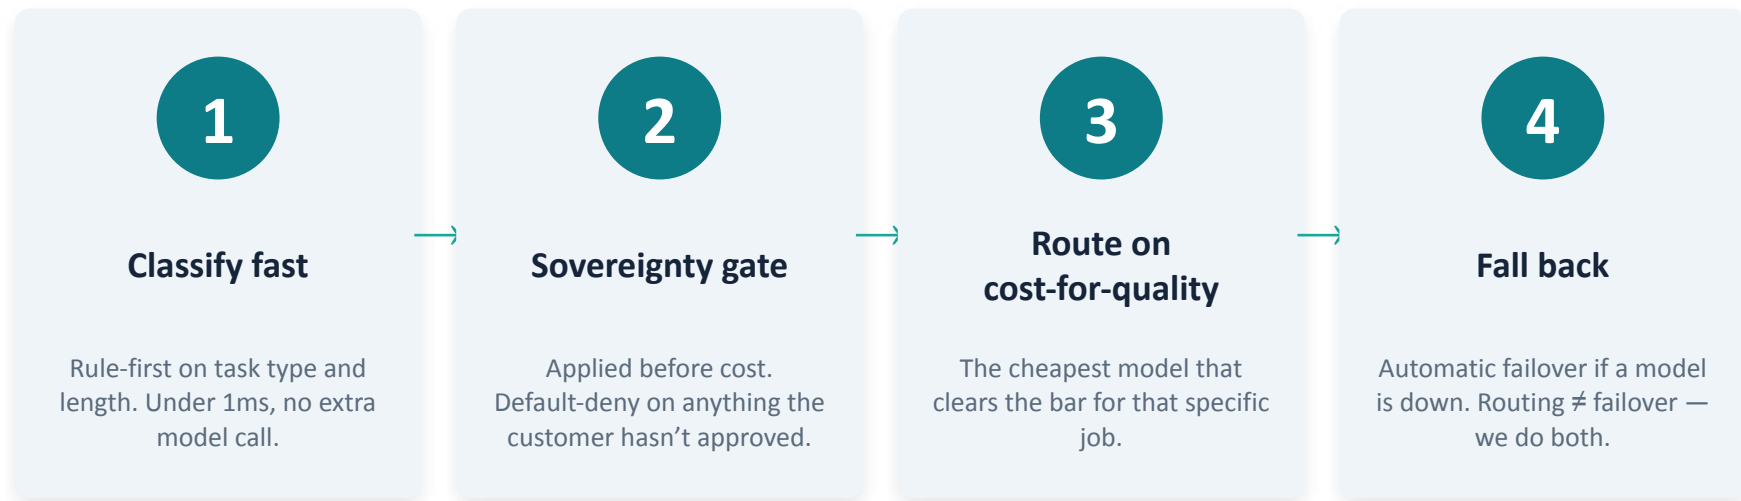
RUN PLANE

Customer inference at scale

- **Cost-optimized and compliance-gated.**
- Sovereignty switch applied per customer.
- Routed by job and posture — where value compounds.

Spend on the build. Commoditize the run.

How a task gets routed — in under a millisecond



Every routing decision is logged and auditable — the audit trail is the security story.



The switch no one else gives the customer



US-only

American models exclusively. No foreign-provenance or export-restricted model touches the tenant.



Allied / open-weight

US plus allied and permissively-licensed open-weight models.



Global-South permissive

Best price-performance, Chinese open-weight allowed where the customer wants it.

Per-tenant • Default-locked to what you agreed • Flippable by an authorized admin • Every flip logged

Most routers optimize cost and quality. Almost none let the customer choose their own security posture. That is our line.

The right model for every job

Job	Best money can buy	Value pick
Code creation (build)	Opus 4.8 / GPT-5.5	Sonnet 4.6
Agentic / terminal	Codex + GPT-5.5 / Claude Code	Kimi K2.7
GUI / frontend	Opus / GPT-5.5	Qwen 3.7 Max
Images	GPT Image 2 (text & labels)	Gemini 3 Pro Image
Inference — quality	GLM-5.2 / DeepSeek V4-Pro	—
Inference — speed / cost	Gemini Flash / Haiku 4.5	DeepSeek V4 Flash

The bench changes monthly. The Switchboard absorbs that — your products never change. (Landscape as of June 2026.)

Commoditize the run, and the savings compound

40–85%


lower cost on customer inference

<1 ms

routing decision overhead

~30x

price spread we exploit



Most production traffic never needed a frontier model. The Switchboard makes sure it doesn't pay for one — while the build plane keeps buying the best.



WHY THIS IS DEFENSIBLE

The models commoditize. The compound doesn't.

The equipment ontology, the competence records, and the orchestration don't commoditize. Independent analysis projects that as public training data runs out between 2026 and 2032, the frontier shifts to companies' proprietary data for specific use cases. That is precisely the Work-Intelligence thesis — and the Switchboard is how we turn cheap, commoditized models into compounding, owned intelligence.

Stop renting intelligence. Own it instead.

One switch. Every product.

Set the posture once, and every EON product inherits it. The Switchboard is a layer you flip — not a rebuild per product.



Genesis



EON Universal



Field IQ



Assess IQ



Soft Skills



Customer Virtual

The same value, quality, and safety guarantees — everywhere EON runs.

Switchboard, Conductor, Verdict



EON Conductor

Orchestrates the workflow — what work runs, in what order.



EON Switchboard

Decides which model runs each step — for value, quality, and safety.



EON Verdict

The safety gate — decides whether the output is allowed to ship.

Conductor decides the work. Switchboard decides the model. Verdict decides if it ships.

The market is moving our way



Model commoditization

Open-weight models at near-frontier quality and near-zero cost — the input layer is becoming a commodity.



Physical AI

The Nvidia / Jensen Huang thesis: AI is moving into the physical world of equipment, facilities, and work.



Proven where it's hardest

Deployed and validated with two of the world's largest energy companies.

Commoditized models, physical work, proprietary data — the Switchboard sits exactly at that intersection.

What this means for the customer



Best value money can buy

They pay frontier prices only where it earns them.



Frontier quality where it matters

Nothing customer-facing ships below the line.



Security they control

They choose their own sovereignty posture — and we prove it.



No vendor lock-in

Models swap underneath; the product never changes.

EON SWITCHBOARD

Understand the work. Route it right. Own the outcome.

Value, quality, and safety — never traded against each other.

Orchestrate the commodity. Own the compound.