

■ COMPANY • PRESS — 2026

Press kit.

The company, the platform, the numbers, and the people behind them. Everything media need to cover Blackbox, ready to use.

■ PRESS: ENTERPRISE@BLACKBOX.AI • V1 • AUGUST 2026

Blackbox at a glance

Core company context for reporters, analysts, and partners.


Blackbox is the high-trust inference layer for teams that deploy AI at scale.

Blackbox (Blackbox AI Technologies Inc.) is the high-trust inference layer for teams that deploy AI at scale. The platform routes requests across 300+ open and closed models through one OpenAI-compatible endpoint, with third-party-verified output speed, end-to-end encryption, and zero data retention. Blackbox sells to enterprises: it runs the open-weight model that a customer chooses as a dedicated deployment, isolated to that customer. Blackbox has offices at 535 Mission Street, San Francisco.

[Company mission](#) ↗

[Who we serve](#) ↗

[Company LinkedIn](#) ↗



Company snapshot

The company at a glance. Every figure links its public source.

THE COMPANY

LEGAL NAME	Blackbox AI Technologies Inc., d/b/a Blackbox AI
INCORPORATION	Delaware, United States
HEADQUARTERS	535 Mission Street, San Francisco, CA
FOUNDERS	Robert Rizk, CEO · Richard Rizk, CTO · Roger Rizk, COO
MODELS SERVED	300+ open and closed models, one OpenAI-compatible endpoint
PRODUCTS	Secure Enterprise Inference platform
SEGMENTS	Large-cap tech · regulated industries · developer-led · high-spend teams
DEVELOPERS	5M+ developers building with Blackbox developer tools
SPEED	454 t/s, Nemotron 3 Ultra, #1 provider (Artificial Analysis, Jul 2026)

OFFICIAL CHANNELS

WEBSITE	blackbox.ai
BLOG	blackbox.ai/blog
DOCS	docs.blackbox.ai
LINKEDIN	linkedin.com/company/blackbox-ai-infra
X	x.com/blackboxai

PRESS CONTACT

enterprise@blackbox.ai

Interviews, embargoed material, and fact checks. We respond within one business day.

SOURCES: BLACKBOX.AI · BLACKBOX.AI/PRESS · ARTIFICIALANALYSIS.AI

Product portfolio

Secure enterprise inference, anchored by the Router and the agent surfaces around it.

Blackbox Router

300+ open and closed models through one endpoint. The gateway enforces zero data retention and no training, wherever the provider API supports it.

[BLACKBOX.AI/API](#)

Agents API


Send coding agents to any repository over HTTP: create runs, stream logs, and open PRs.

[BLACKBOX.AI/AGENTS](#)

Command Line

An agentic terminal that reads the conventions of your repository.

[BLACKBOX.AI/CLI](#)



SECURE INFERENCE FOR ENTERPRISE

Enterprise Inference

The open-weight model that you choose, on capacity reserved for you: no shared pools, and no other customers on the instance.

[TALK TO SALES → BLACKBOX.AI/CONTACT](#)

Performance and scale

Independently measured and publicly comparable. Every figure links its source below.

01

454 t/s

NEMOTRON 3 ULTRA · #1 PROVIDER
· ARTIFICIAL ANALYSIS, JULY 2026 [1]

02

300+

OPEN + CLOSED MODELS · ONE
ENDPOINT [2]

03

5M+

DEVELOPERS · MARKETPLACE
INSTALLS, JULY 2026 [3]

04

SF

OFFICE · 535 MISSION STREET [4]

[1] [ARTIFICIALANALYSIS.AI](#) [2] [BLACKBOX.AI/MODELS](#) [3] [PUBLIC MARKETPLACE INSTALL COUNT](#) [4] [TERMS OF SERVICE](#)

What Blackbox does

IN ONE SENTENCE

Blackbox is the high-trust inference layer for teams that deploy AI at scale.

The platform routes requests across 300+ open and closed models through one OpenAI-compatible endpoint, with third-party-verified output speed, end-to-end encryption, and zero data retention. For enterprises, Blackbox runs the open-weight model a customer chooses as a dedicated deployment, isolated to that customer.

FOR INFRASTRUCTURE AND AI REPORTERS

Applications call OpenAI-compatible endpoints (/v1/chat/completions, /v1/embeddings, /v1/images/generations) at enterprise.blackbox.ai with one API key. The gateway routes each request to the selected model, enforces retention suppression through provider terms and per-request flags (wherever the provider API supports it), and falls back automatically when a provider has an outage. Data is encrypted at rest with AES-256. On Enterprise, PII is removed before a prompt reaches a closed model; workloads that cannot accept any third party run single-tenant.

REQUEST PATH



THE INFERENCE PLATFORM

The Router at enterprise.blackbox.ai plus dedicated enterprise deployments. Sold to enterprises, metered per token.

THE DEVELOPER TOOLS

The Blackbox CLI and Agents API: the agent surfaces five million developers build with.

Enterprise capabilities

Security, privacy, and control: the platform that security teams approve. Primary sources at right.

CAPABILITIES	
ENCRYPTION	End-to-end encryption, with data encrypted at rest with AES-256.
DATA RETENTION	We retain zero data and provide Enterprise with a single tenant endpoint.
TRAINING	Training opt-out is on by default.
PII MINIMIZATION	On Enterprise, we remove PII before a prompt reaches a closed model: the provider sees the task, not the person it is about.
CUSTOMER ISOLATION	Dedicated deployments run the open-weight model you choose as a single tenant, isolated from every other customer.
IDENTITY & ACCESS	SAML SSO and advanced security controls on Enterprise; Azure AD SSO for Microsoft-stack teams.
ABUSE MONITORING	Abuse monitoring and human review can be disabled, negotiated into the contract, then enforced technically.
GOVERNANCE	Governance, security, a single-tenant deployment, and dedicated support are included on Enterprise, not sold as an upgrade.

PRIMARY SOURCES
SECURITY BRIEF blackbox.ai/#security
PRIVACY ARCHITECTURE blackbox.ai/privacy
ZERO DATA RETENTION, EXPLAINED blackbox.ai/blog/zero-data-retention-how-blackbox-enforces-it
ENCRYPTED INFERENCE blackbox.ai/blog/blackbox-encrypted-ai-secure-llm-inference

SECURITY QUESTIONS

enterprise@blackbox.ai

Security review material and deployment specifics are provided to customers and media through the enterprise team.

Customers and use cases

What enterprises run on the platform, from multi-model inference to dedicated deployments and high-throughput workloads.

APPROVED USE CASES

01 MULTI-MODEL INFERENCE

One endpoint, 300+ open and closed models. The correct model for each task, at verified top speed.

02 DEDICATED OPEN-WEIGHT DEPLOYMENTS

A deployment that is only yours: dedicated, isolated from every other customer, with zero data retention.

03 PROVIDER FAILOVER

Automatic fallback routing keeps requests in service when a provider has an outage; load balancing distributes for the lowest latency.

04 USAGE & SPEND GOVERNANCE

One endpoint and one bill. Routing and caching lower the cost of the same models without a change to your stack.

05 SECURE DEVELOPMENT WORKFLOWS

Coding agents in the IDE, terminal, and CI behind an endpoint that a security team can approve.

06 HIGH-THROUGHPUT WORKLOADS

454 tokens/sec verified output speed, with load balancing across providers for latency-critical and batch workloads alike.

WHO BUYS

LARGE-CAP TECHNOLOGY

The correct model for each task; expose only the data each task requires.

REGULATED INDUSTRIES

Dedicated, isolated deployments with zero data retention.

DEVELOPER-LED ADOPTION

The open models engineers already use, behind an approvable endpoint.

HIGH-SPEND TEAMS

Lower model spend with no change to the stack.

ON THE RECORD

Customer names, logos, and quotes appear only with written permission – the same standard we hold our claims to. To go on the record: enterprise@blackbox.ai.

Founders

Who runs the company, and what each founder speaks to. LinkedIn links go to each founder’s profile. Interviews: enterprise@blackbox.ai – response within one business day.

01



Leads the company and its enterprise business: taking high-trust inference – verified speed, encryption, zero retention – to teams deploying AI at scale.

INTERVIEW TOPICS

Enterprise AI adoption · company strategy · pricing and the token economy

Robert Rizk

CO-FOUNDER & CEO

[LINKEDIN](#)

02



Leads the platform: the routing gateway, dedicated deployments, and the encryption and privacy architecture behind 300+ models on one endpoint.

INTERVIEW TOPICS

Inference infrastructure · model routing · encryption and privacy architecture

Richard Rizk

CO-FOUNDER & CTO

[LINKEDIN](#)

03



Leads operations and go-to-market: the engagements, dedicated deployments, and support behind the platform’s enterprise customers.

INTERVIEW TOPICS

Go-to-market · operations · dedicated enterprise deployments

Roger Rizk

CO-FOUNDER & COO

[LINKEDIN](#)

Company momentum

A year of launches, benchmarks, partnerships, and adoption. Every entry links its source.

■ JAN 2026

Zero data retention, explained: how the platform enforces retention suppression by default.
[blackbox.ai/blog/zero-data-retention-how-blackbox-enforces-it](#)

■ MAY 2026

Blackbox Encrypted AI published: confidential LLM inference on the platform.
[blackbox.ai/blog/blackbox-encrypted-ai-secure-llm-inference](#)

■ MAY 2026

Nemotron comes to Blackbox: NVIDIA's open-model family served on the platform.
[blackbox.ai/blog/nemotron-on-blackbox](#)

■ JUN 2026

Orchestrator-executor: the multi-agent architecture behind the coding agents, published.
[blackbox.ai/blog/orchestrator-executor-multi-agent-system-for-swe-tasks](#)

■ JUL 2026

Artificial Analysis verifies Blackbox as the #1 Nemotron 3 Ultra provider at 454 tokens/sec.
[blackbox.ai/blog/fastest-nemotron-3-ultra-provider](#)

■ JUL 2026

Terminal-Bench v2.1 results published: 90% pass@1 with GPT-5.6 Sol + Opus 4.8.
[blackbox.ai/blog/tb-v2-1-blackbox-gpt-5-6-sol-opus-4-8](#)

■ JUL 2026

Blackbox developer tools pass 5M+ installs.
[blackbox.ai/customers](#)

■ AUG 2026

Nemotron 3.5 Lightning announced on the platform.
[blackbox.ai/blog/nemotron-3-5-lightning](#)

PARTNERSHIP

Microsoft partner — AI for the Azure stack, with Azure AD SSO and enterprise features for Microsoft-ecosystem teams. [blackbox.ai/microsoft-partner](#)

Media FAQ & downloads

Approved answers to the questions we hear most, and everything there is to download.

01 WHAT IS BLACKBOX AI?

Blackbox AI is the high-trust inference layer for teams that deploy AI at scale: 300+ models behind one OpenAI-compatible endpoint, plus dedicated deployments for enterprises.

02 INFERENCE PLATFORM OR DEVELOPER-TOOLS COMPANY?

Both, as two businesses that reinforce each other. The inference platform (Router and dedicated deployments) is the enterprise business; the developer tools have put Blackbox in front of five million developers.

03 HOW DOES MODEL ROUTING WORK?

Requests hit one endpoint with one API key. The gateway sends each request to the selected model, falls back automatically when a provider has an outage, and load-balances for the lowest latency.

04 WHAT DOES “OPENAI-COMPATIBLE” MEAN?

The API accepts the same request format as OpenAI’s endpoints (/v1/chat/completions, /v1/embeddings, /v1/images/generations). Teams migrate by changing the base URL and key.

05 WHAT IS A DEDICATED DEPLOYMENT?

Blackbox runs the open-weight model a customer chooses on dedicated infrastructure: one tenant, isolated from every other customer, with end-to-end encryption.

06 WHAT CUSTOMER DATA IS RETAINED?

We provide end-to-end encryption and retain zero data. On routed traffic, retention is suppressed through provider terms and per-request flags, wherever the provider API supports it.

07 HOW ARE PERFORMANCE CLAIMS MEASURED?

Independently, by Artificial Analysis, on its public performance methodology. Figures are directly comparable to any model on the public leaderboard; cite the source and date.

08 HOW IS BLACKBOX DIFFERENT FROM HYPERSCALERS AND AGGREGATORS?

Blackbox sells the trust layer, not just capacity: third-party-verified speed, retention suppression on routed traffic, and single-tenant deployments of open-weight models, behind one endpoint.

BLACKBOX.AI

BLACKBOX.AI



PRESS & BRAND

enterprise@blackbox.ai

HEADSHOTS · SCREENSHOTS · INTERVIEWS · ONE BUSINESS DAY

WORDMARK SVG	blackbox.ai/blackbox/blackbox-wordmark.svg	MARK PNG · LIGHT	blackbox.ai/blackbox/blackbox-mark-black-transparent.png
WORDMARK PNG · LIGHT	blackbox.ai/blackbox/blackbox-wordmark-black-transparent.png	MARK PNG · DARK	blackbox.ai/blackbox/blackbox-mark-white-transparent.png
WORDMARK PNG · DARK	blackbox.ai/blackbox/blackbox-wordmark-white-transparent.png	MARK PNG · SQUARE	blackbox.ai/blackbox/blackbox-mark-white-on-black.png
MARK SVG	blackbox.ai/blackbox/blackbox-mark.svg	ALL LOGOS · ZIP	blackbox.ai/blackbox/blackbox-logo-assets.zip



Everything you need to cover Blackbox.

Press: enterprise@blackbox.ai. We respond within one business day. Brand and asset questions route to the same address.

[BLACKBOX.AI](https://blackbox.ai)

[BLACKBOX.AI/PRESS](https://blackbox.ai/press)