Web, API and AI protection · Runs itself · Explains itself

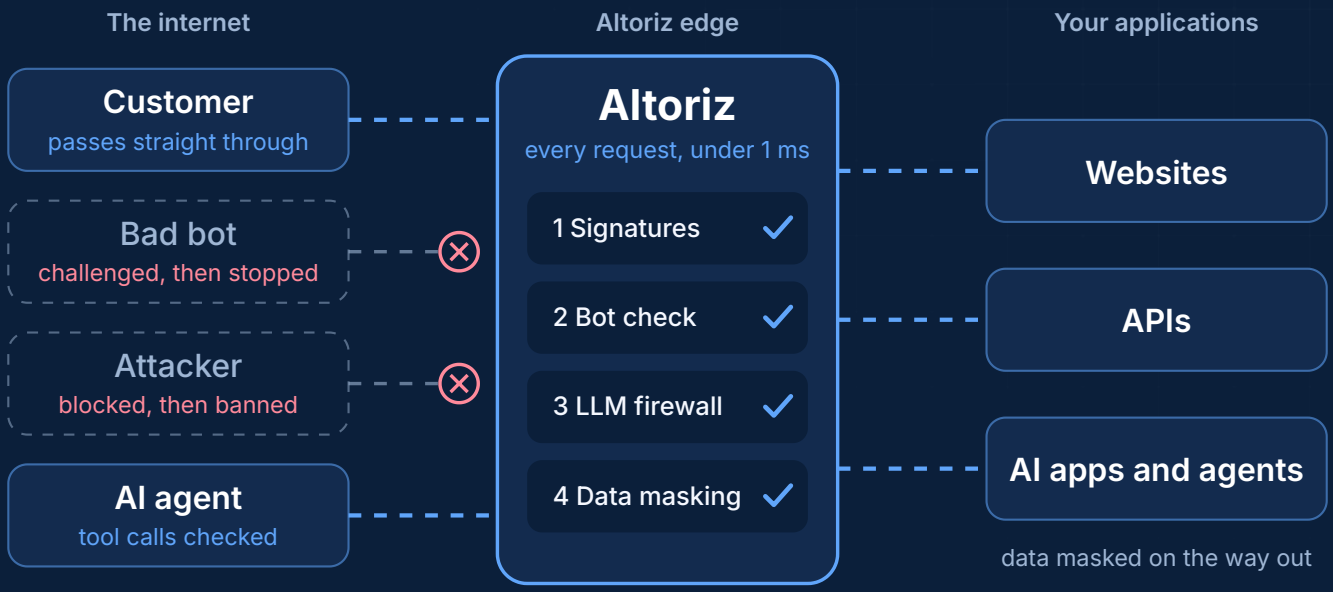
# The firewall that defends itself.

Altoriz protects your websites, APIs and AI applications with an AI agent that stops attacks on its own, and explains every decision in plain language.

**< 1 ms**
added to each request, checked inline at the edge

**7 days**
to learn your traffic, then block attacks on its own

**6 hours**
between checks for new CVEs, patched at the edge




Customers pass straight through. Bad bots are challenged, attackers are blocked and then banned, and AI traffic is checked before it reaches your model.

**Runs itself.** **Explains itself.** **Keeps your AI private.**

Websites · APIs · AI apps

Managed edge · Self-hosted

**Chennai, India**

Why change now

# Attackers have automated. Most WAFs haven't.

The web application firewall was built for hand-written attacks and hand-tuned rules. Today attacks arrive by machine, around the clock, and five gaps stay open.

**Verizon DBIR, 2026:** for the first time, an exploited vulnerability is the #1 way breaches begin (31%), ahead of stolen passwords.

| The WAF gap | How Altoriz closes it |
|---|---|
| 01 **It only watches** Fear of false positives keeps many WAFs in detect-only mode. | **Protected by default** Learns your traffic for 7 days, then switches itself to blocking. Suspected false positives are held for review. |
| 02 **New flaws outrun the rules** 131 new CVEs a day in 2025, each able to slip past static rules. | **Patched at the edge** Every 6 hours the agent checks for new CVEs, tests a virtual patch, then enforces it. |
| 03 **Alerts, not answers** Endless alerts, and rule tuning that needs scarce specialists. | **Attacks arrive investigated** Every few seconds the agent links attacks to the attacker, bans repeat ones and explains why. |
| 04 **Blind to APIs and bots** Shadow APIs, scraping and bots posing as customers slip through. | **Sees APIs and bots** API discovery, object-enumeration (BOLA) detection, schema checks and bot defense on one platform. |
| 05 **Blind to AI apps** Prompt injection, data leaks and agent tool abuse go unseen. | **Guards AI apps and agents** An LLM firewall, prompt redaction, and policy for agent-to-agent and MCP tool calls. |

# Attackers don't knock. They automate.

**53% of web traffic is now bots.** A WAF that waits for a person to read an alert is already behind. Altoriz acts on its own, under a policy you set, and tells you what it did.

### Signs it is time to move

- [x] Your WAF is still in detect-only mode
- [x] A critical CVE is waiting on a code fix
- [x] You are launching an LLM feature or AI agent
- [x] An API has exposed personal data
- [x] Nobody watches the WAF at night

What standing still costs

**₹25.5 crore**
average cost of a data breach in India
IBM, 2026

**$4.99M**
average cost of a data breach worldwide, a record
IBM, 2026

**+113%**
rise in daily API attacks in a single year
Akamai, 2026

**₹250 crore**
maximum DPDP penalty for weak security safeguards
DPDP Rules, 2025

What Altoriz is

# Protection that runs itself, and explains itself.

A web application firewall with a built-in AI security agent. It sits in front of your websites, APIs and AI apps, checks every request, and runs an investigation loop in the background.

Virtual patching

### New CVE? Patched at the edge.

The agent drafts and tests a virtual patch for the technology you run, then enforces it before developers ship a fix.

AI and LLM protection

### AI apps and agents, guarded.

Prompt injection, jailbreaks and personal data in prompts are stopped, and agents may call only the tools you allow.

## How it works: a security team that never sleeps

Five steps, running all the time.

1

### Inspect

**A guard at every door**

Every request is checked inline, in under 1 ms.

2

### Detect

**An alarm that knows why**

Attacks become security events; leaks, bots and AI misuse are flagged.

3

### Investigate

**A detective on every case**

Every few seconds, events are linked to the attacker behind them.

4

### Act

**A bouncer who remembers faces**

Under your policy, repeat attackers are banned and patches deployed.

5

### Explain

**A report anyone can read**

Each action gets a plain-language reason, evidence and one-click undo.

**Always on, always explained.**
Fast where it must be, careful where it counts.

**< 1 ms**
added to each request

**Seconds**
between agent investigation cycles

**1 click**
to undo any automatic action

**365 days**
of audit log, for every customer

## Who benefits: boardroom to front line

No rule-tuning for any of them.

Board and CEO

### Risk you can show the board

- A security score from 0 to 100
- Scheduled PDF reports
- A full audit trail

Security team

### Attacks arrive investigated

- Kill chain and MITRE ATT&CK mapping
- Automatic bans under your policy
- Fewer alerts to triage

Platform and IT

### Fast, and fails safe

- Under 1 ms per request
- Keeps the last good policy if anything fails
- Terraform provider and full API

Product and AI teams

### Ship AI features safely

- Prompt-injection defense
- Personal data redacted from prompts
- No code changes

**Privacy and audit:** all AI runs inside the platform, so your traffic never goes to an outside AI service. Every console change is logged for 365 days and can be exported when auditors ask.

Where it runs

# At our edge or in your network. Your call.

Point your DNS at Altoriz and we run it, or run it yourself so payloads never leave your network. Either way, AI stays private and security events flow to the tools you already use.

Run by us

### Managed edge

Point your DNS at Altoriz. We run the edge and the control plane, with per-customer isolation.

Run by you

### Self-hosted

Run the data plane, and optionally the control plane, in your own environment with Docker.

Works with your tools

### Integrations

Splunk, Elastic, Microsoft Sentinel, Google SecOps, Datadog, syslog, Kafka, S3 and Terraform.

The return on acting

## Security AI and automation cut the cost of a breach.

IBM's 2026 study of real breaches: organisations that use security AI and automation extensively pay far less and recover faster.

Lower cost per breach

**₹10.3 crore**
less per breach in India

**$1.93M**
less per breach worldwide

Average breach cost in India

With and without security AI and automation

With extensive security AI and automation — ₹21.3 cr

Without it — ₹31.6 cr

**65 days** shorter breach lifecycle with it, worldwide

Where Altoriz adds to it

- **Acts in seconds, not shifts** — The agent investigates and bans around the clock.
- **Closes the patch gap** — Virtual patches every 6 hours, before code fixes ship.
- **Keeps blocking on** — Protected by default, so it never sits in watch-only mode.

Industry averages from IBM Cost of a Data Breach Report 2026 (India and global figures), not Altoriz-specific results.

## A gradual, low-risk path

1. **Demo** — See the agent investigate attacks and explain each decision.
2. **Pilot on one application** — It learns for 7 days, then switches itself to blocking.
3. **Extend to every application** — One product, all features included. No tiers to upgrade.

Next step

## See it defend your own traffic.

- Stop tuning, start blocking
- Patch new CVEs at the edge
- Protect your AI features

**Our recommendation:** book a demo, then pilot Altoriz on one customer-facing application.

Book a demo →

**Altoriz**
Chennai, India

Altoriz

The web application firewall that investigates and defends itself.

**Runs itself. Explains itself.**
Web, API and AI-application protection

Sources: Verizon DBIR 2026; Imperva Bad Bot Report 2026; CVE program data 2025; IBM Cost of a Data Breach Report 2026 (global and India); Akamai State of the Internet 2026; DPDP Rules 2025. Third-party figures are industry data, not Altoriz results. Altoriz is designed to support SOC 2 controls; it is not yet certified.