ARES
Server Guardian

ARES SUITE

# ARES Server Guardian

## Product Description

Detects attacks on your servers and blocks them at the firewall.

| | | |
|---|---|---|
| **VERSION** | **PRODUCTION** | **DATE** |
| 1.6.0 | ares-guardian.nacino.net | 23 November 2026 |
| **AUTHOR** | **DEMO** | **CLASSIFICATION** |
| Alfredo Nacino | ares-guardian-demo.nacino.net | Internal |

# Contents

01 Who it is for

02 Two ways to run it

03 Main features

04 A typical workflow

05 Accounts and security

06 Platforms

07 Limits

ARES Server Guardian watches a Linux server for attacks and blocks the attackers in the server's own firewall. This document describes every feature from the point of view of the people who use it. The screenshots show demo data: documentation addresses (192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24) and example.com sites.

## 1. Who it is for

| Person | What they do with ARES |
|---|---|
| **Server administrator** | Installs ARES, sets the whitelist, switches the firewall from dry-run to enforcing, adds rules. |
| **Operator** | Watches alerts, bans or unbans addresses, blocks an attacked web application. |
| **Viewer** | Follows the dashboard, the attackers and the reports. |
| **Fleet manager** | Manages many servers from ARES Central; the servers can run as agents with no dashboard of their own. |

## 2. Two ways to run it

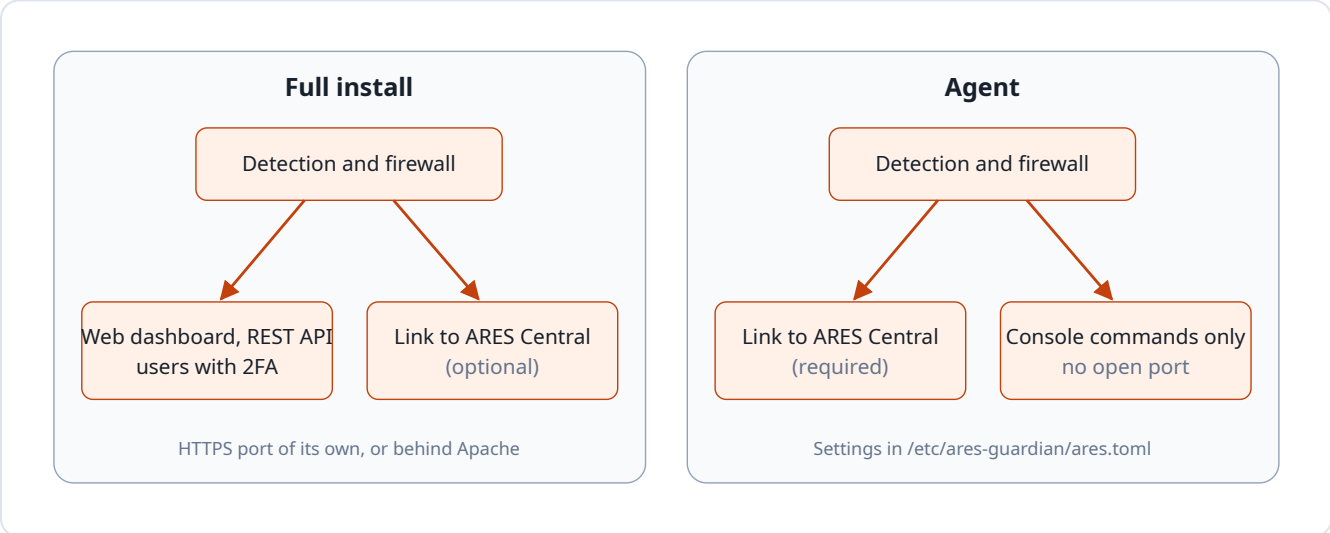



**Figure 1.** Full install and agent, side by side

- A **full install** has the dashboard described below, served on its own HTTPS port or behind Apache.
- An **agent** has no dashboard, no users and no open port. It detects and blocks attacks the same way and is seen and managed only in ARES Central. Its settings are in `/etc/ares-guardian/ares.toml`.

# 3. Main features

## 3.1 Dashboard

The first page shows the server's health (CPU, memory, disk), the active bans, the attackers and incidents of the last 24 hours, the web attacks, the open alerts and a chart of security events per hour. A banner says when the firewall is still in dry-run.

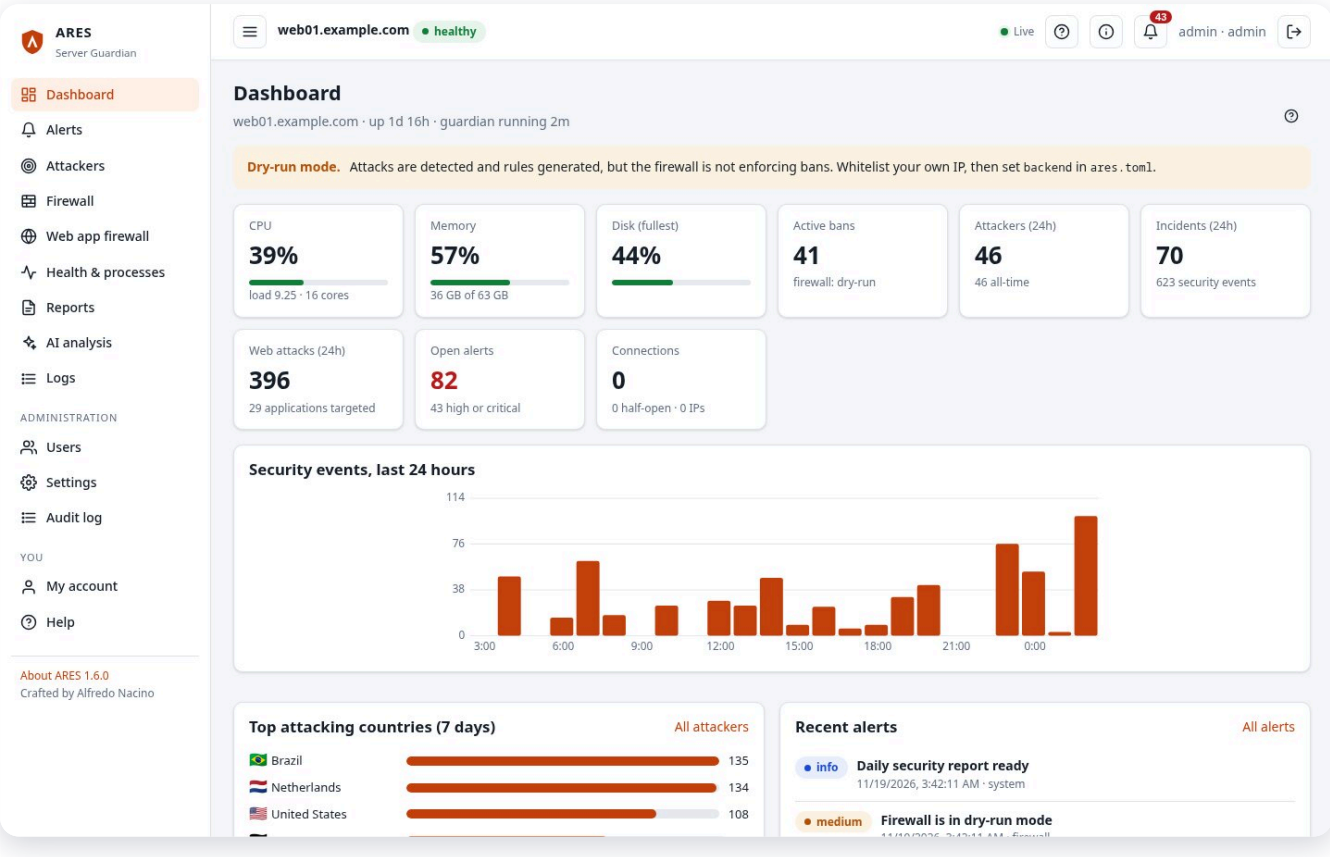



**Figure 2.** Dashboard (demo data)

## 3.2 Attack detection and automatic bans

ARES recognises eight kinds of attack and acts on each with its own threshold:

| Attack | Default trigger | Action |
|---|---|---|
| SSH password guessing | 5 failures in 5 minutes | ban |
| Web login guessing | 10 failures in 5 minutes | ban |
| HTTP flood (layer-7 DDoS) | 600 requests in 1 minute | ban, plus a rate limit for every source |
| Web vulnerability scan | 20 suspicious requests in 5 minutes | ban |
| Port scan | 15 ports in 1 minute | ban |
| Connection flood | 150 open connections from one address | ban, plus a rate limit |
| SYN flood | 300 half-open connections in total | SYN rate limit |

| Attack | Default trigger | Action |
|---|---|---|
| Guessing ARES's own password | 8 failures in 10 minutes | ban |

Bans last 60 minutes by default and double with each repeat; the third automatic ban of an address is permanent. After **one** failed login the brute-force guard already limits that address to 6 new connections a minute for 30 minutes, so a password-guessing tool is slowed down before the ban.

## 3.3 Alerts

Every attack, ban, overload and account event becomes an alert with a severity. On sign-in, each person sees the alerts raised since their last visit. Alerts stream live, can be acknowledged or resolved, and can be sent to a webhook (Slack, Discord, ntfy).

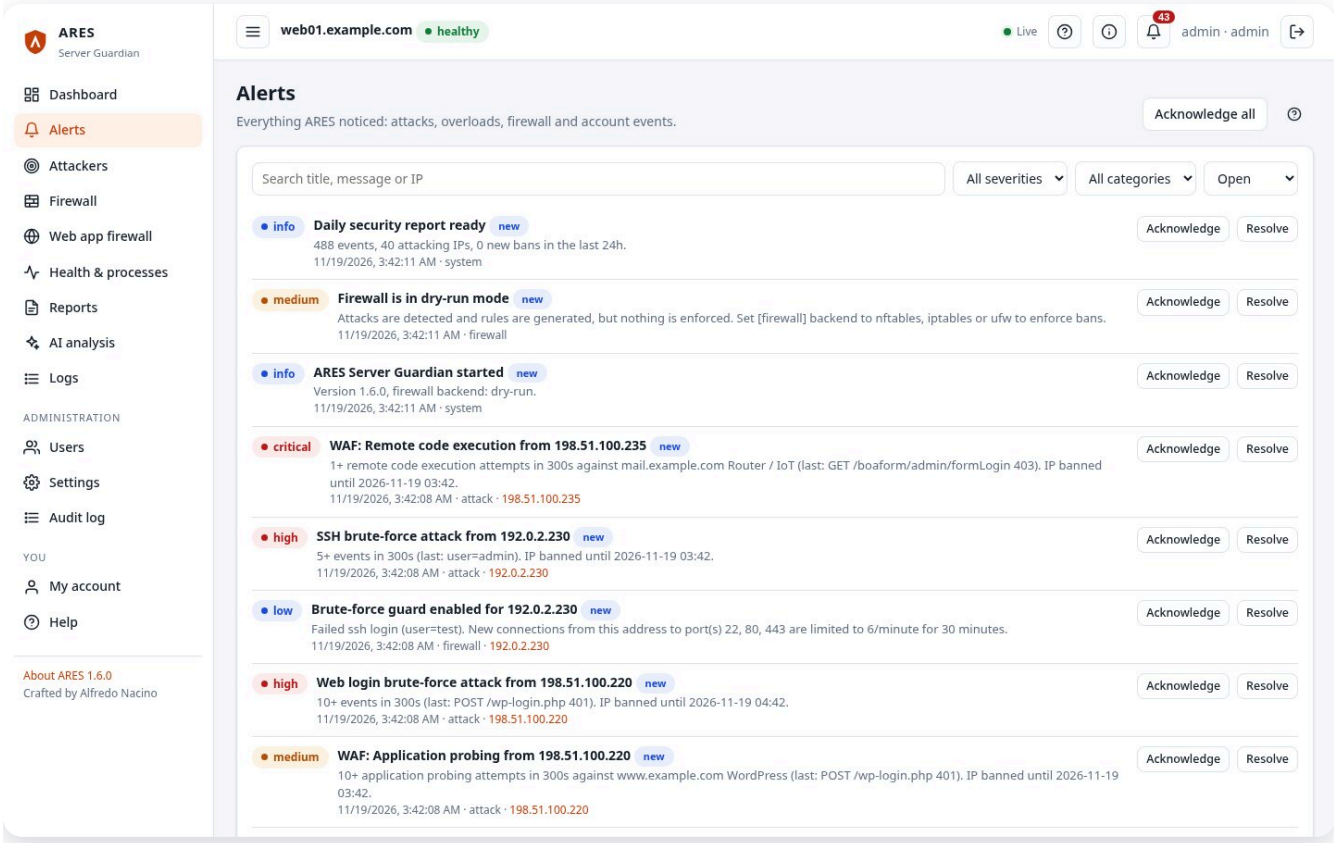



**Figure 3.** Alerts (demo data)

## 3.4 Attackers

Every address that attacked the server is listed with its country, city, network, the types of attack, the number of events and bans, and its status. The list can be searched, filtered by country, and exported to CSV; a second tab groups attackers by country.

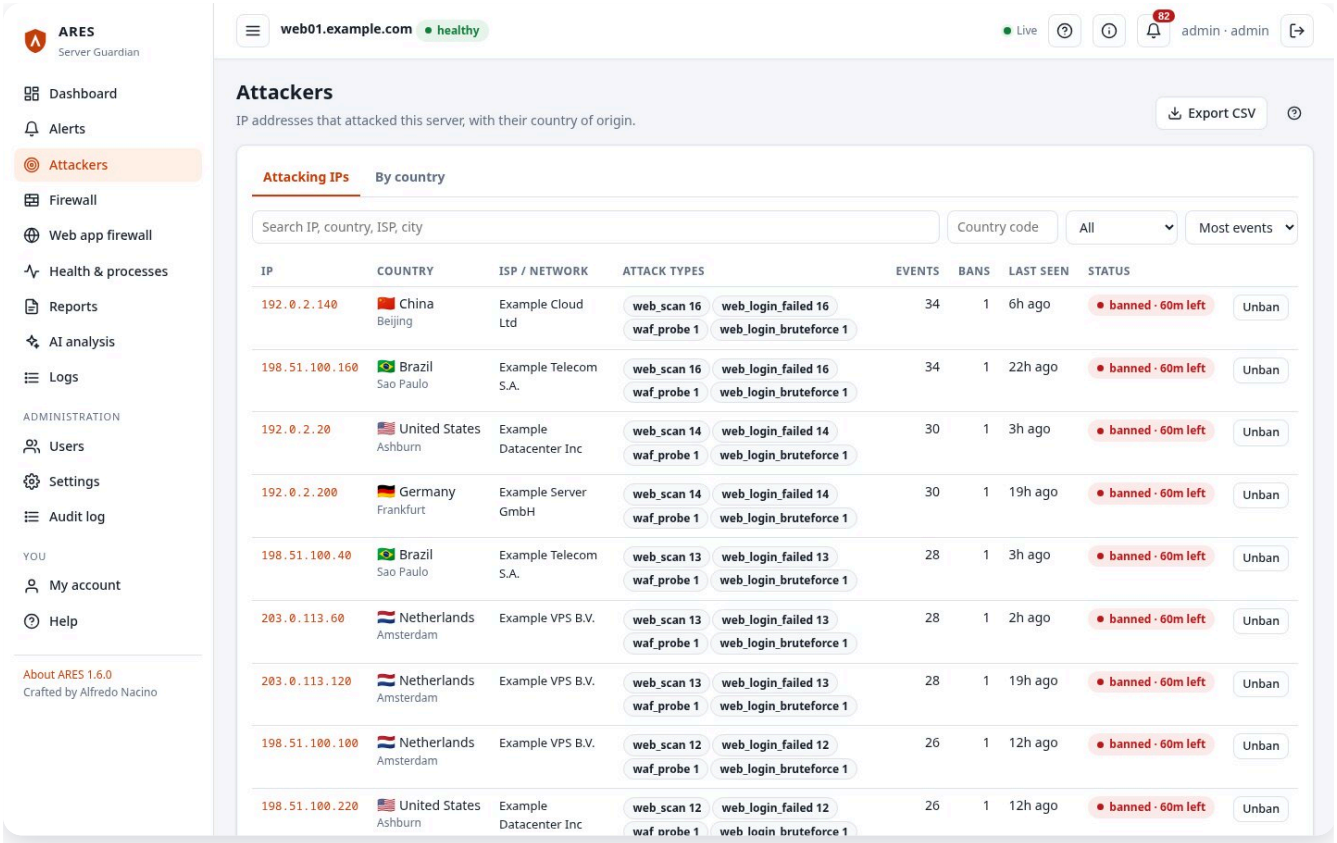



**Figure 4.** Attackers (demo data)

## 3.5 Firewall

The firewall page shows the backend in use, the active bans and rules, and the visitor's own address with a **Whitelist my IP** button. Tabs list bans, rate limits, port blocks, allow rules, the whitelist, allowed and denied DNS names, the detection rules and the ruleset ARES generates. Operators ban and unban addresses, add rules and apply one-click **anti-DDoS presets** (web, SSH, all TCP, all UDP).

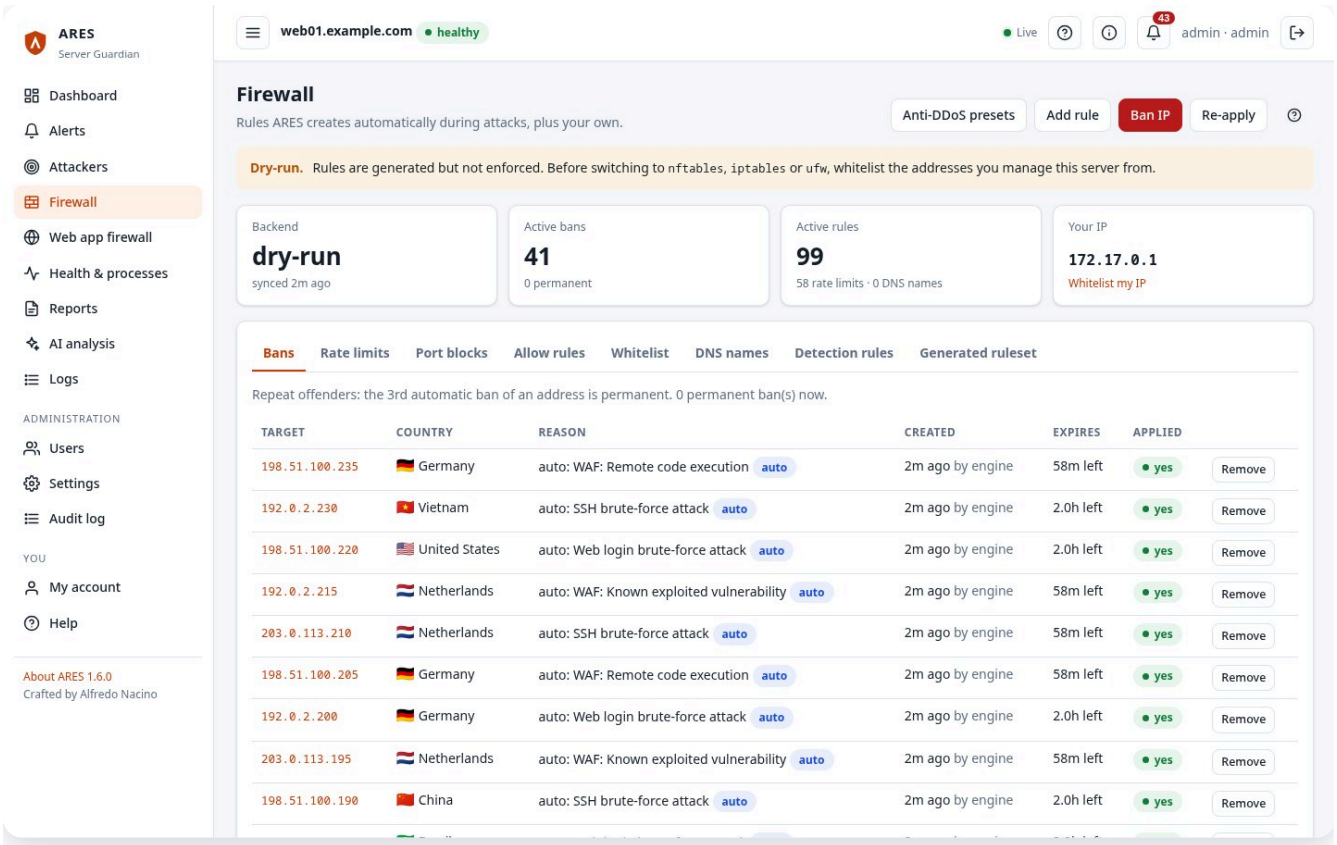



**Figure 5.** Firewall (demo data)

Allowed and denied **DNS names** are made for dynamic addresses: a name is looked up every 5 minutes, and when its address changes the old rule is replaced by one for the new address.

## 3.6 Web application firewall

Every request in the web server's access logs is checked against ten groups of signatures: known exploited vulnerabilities (CVEs), server-side request forgery, template and code injection, remote code execution, SQL injection, path traversal, cross-site scripting, sensitive files, application probing and scanner tools. A table shows which applications on which sites are under attack; **Block** creates a rule for one application on one site with the pattern filled in.

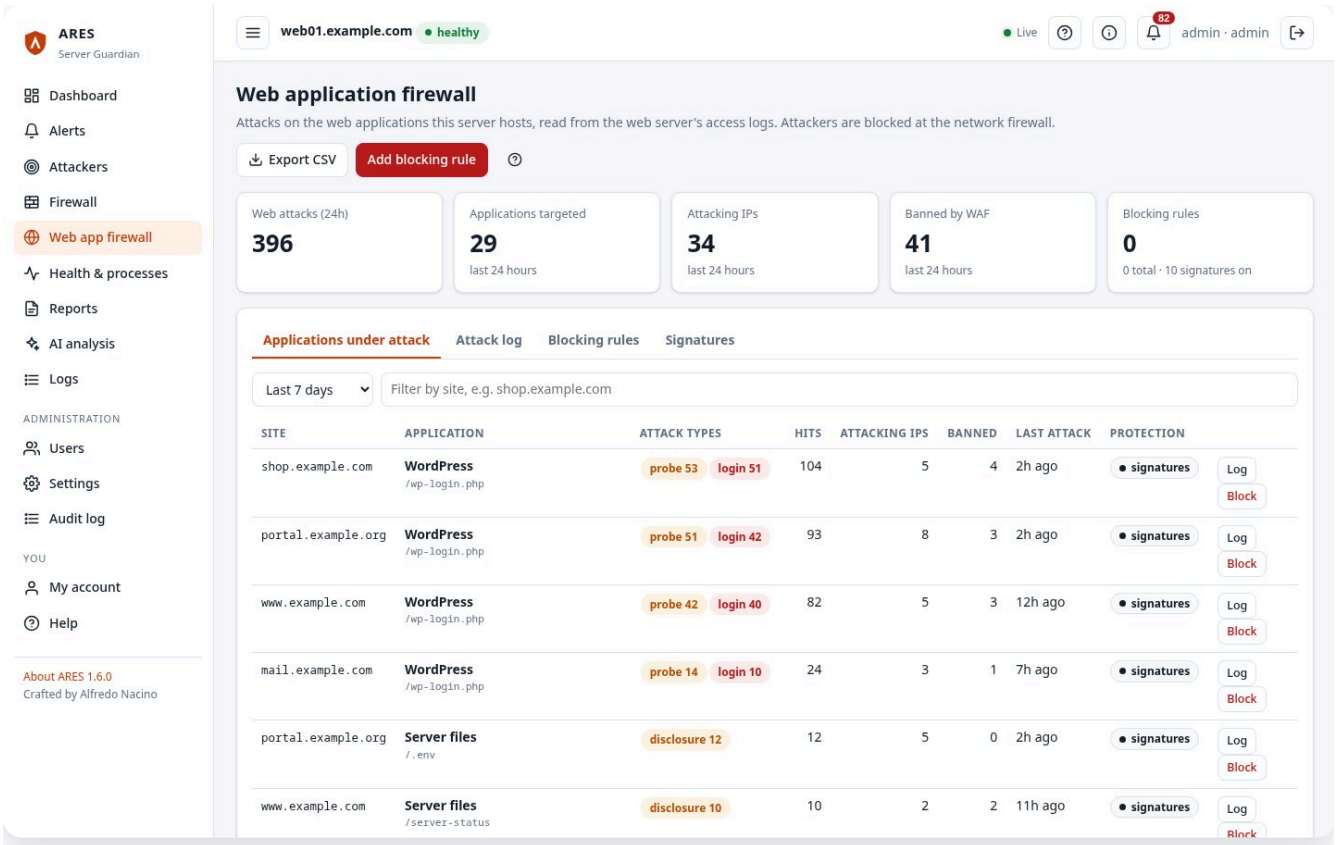



**Figure 6.** Web application firewall (demo data)

## 3.7 Logs and attack tracing

The Logs page shows everything ARES records (activity, security events, web attacks, alerts, firewall changes, AI runs and its own application log) as a table, raw lines or JSON, with search, dates, live follow and downloads. Any address can be **traced** across all logs. The trace page adds an analysis made by ARES itself, without AI: whether a ban is enforced, logins that succeeded after failures, attacks that were answered with success.

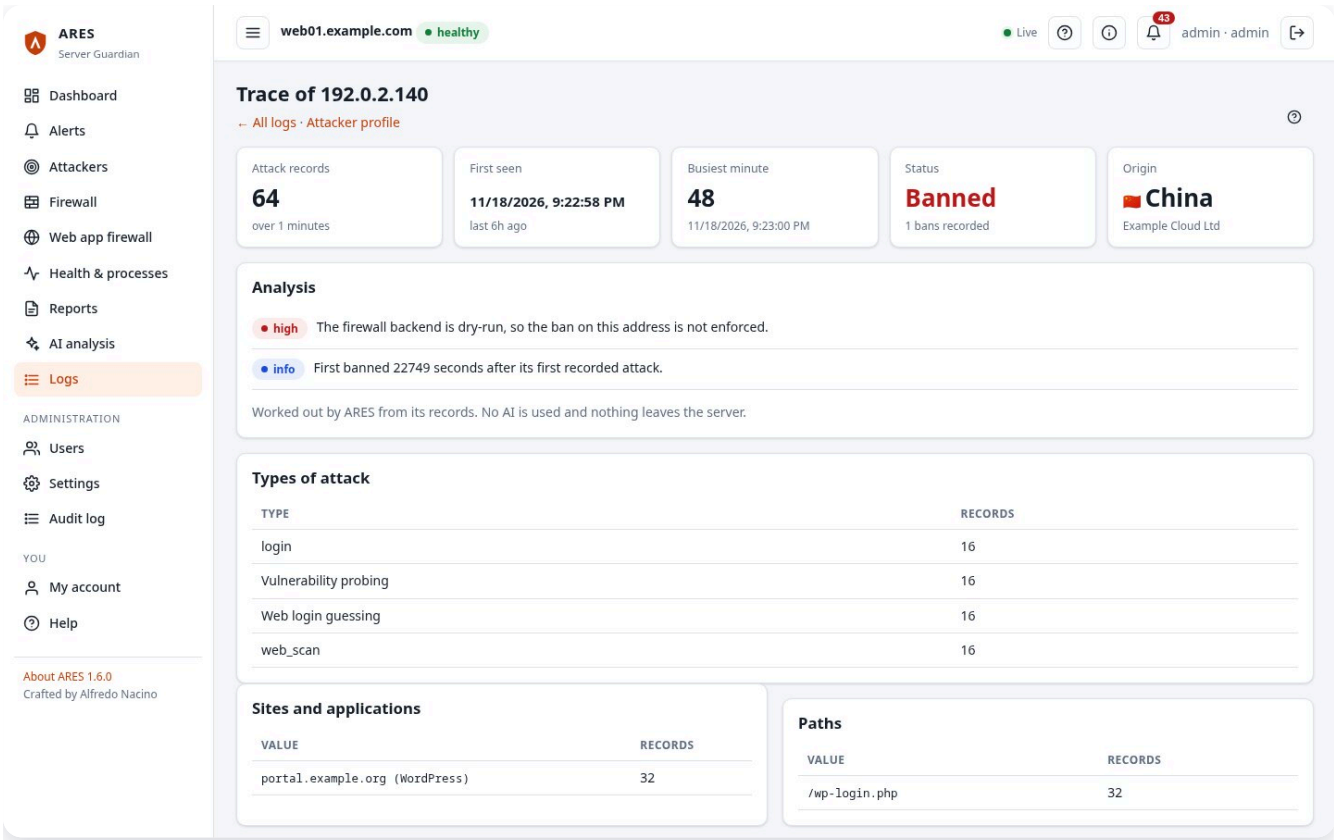



**Figure 7.** Trace of one address (demo data)

## 3.8 Health and processes

ARES samples CPU, memory, swap, disks, load, network and connections, and flags processes that overload the server. An administrator can stop such a process from the dashboard.

## 3.9 Reports, AI analysis and backups

- **Reports**: a daily security report, and reports on demand, as HTML, CSV or JSON, with attack rates per second, minute and hour.
- **AI analysis** (off by default): Claude, through Claude Code on the owner's Claude plan, assesses the recorded attacks and returns a risk score, findings and recommendations, with an emailed PDF on a schedule. It never changes the firewall.
- **Backups**: scheduled or on demand, of the database, keys and configuration; restored from the console.
- **Log forwarding**: events, alerts, the audit log and web attacks to your own log server as JSON over HTTPS, TLS, TCP or UDP. **Log retention** of 1, 2, 3 or 6 months.

## 3.10 ARES Central

A server joined to ARES Central reports its health, alerts, attackers, attacks and firewall rules, and accepts signed firewall commands from Central's operators. Joining is done once on the console with a one-time code; the owner can make the server read-only for Central.

## 4. A typical workflow




**Figure 8.** From a new server to enforced protection

An **agent** shortens this: `scripts/install-remote.sh` installs it on the servers you name, whitelists the server's own addresses and the administrator's SSH address, picks nftables when it is present, and joins Central in one step.

## 5. Accounts and security

- Three roles: viewer, operator and administrator. Two-factor authentication with an authenticator app is mandatory, with ten recovery codes; accounts lock after repeated failures.
- Access to the dashboard can be limited to listed addresses, networks and DNS names. People elsewhere receive a one-time link by email that opens only their own account and still needs password and 2FA.
- Every sign-in and change is in the audit log.

## 6. Platforms

| Where | How |
|---|---|
| Linux server (Debian, Ubuntu and similar, systemd) | `deploy/install.sh`, `deploy/deploy-standalone.sh`, `deploy/deploy-apache.sh` or `deploy/install-agent.sh` from the release bundle |
| Many servers at once | `scripts/install-remote.sh` and `scripts/deploy-remote.sh` over SSH |
| Container | `docker.cybercontrol.tech/ares/server-guardian` with the host's network and processes |
| Browser | Any current browser for the dashboard; a REST API under `/api/v1` for scripts and the mobile app |

There is no desktop application: ARES must run on the server it protects.

## 7. Limits

- ARES protects the server it runs on; protecting a whole network needs one install per server.
- It needs root (or the network administration capability) to change the firewall; the Demo runs unprivileged and therefore only in dry-run.
- An agent has no web settings: users, SMTP, AI and allowed addresses do not apply, and its other settings are edited in `ares.toml`.
- Country lookup uses an online service or a local MaxMind database; without either, countries show as unknown.