



SENTRA
DOCUMENTATION

USER GUIDES

SENTRA AI Threat Shield — User Guide

Reading the shield's console day to day, for website owners and analysts

A plain-language guide for website owners and customer staff with read-only (Analyst-type) access — signing in, reading the Overview and Live pages, what each verdict means, looking into an event, reports, asking for changes, everyday workflows and a short quiz.

Version 2.26.0 27 October 2026

Contents

01 Welcome

02 Reading the Overview

03 Watching Live

04 What the verdicts mean

05 Looking into an event

06 Reports

07 Asking for changes

08 Everyday workflows

09 FAQ


10 Quiz

Welcome

What SENTRA AI Threat Shield does for your website, what you can see with your account, and how to sign in.

Your website is protected by **SENTRA AI Threat Shield** (shown as **Web App · ATS**). It sits in front of the website and checks every visit before it reaches your web server. Ordinary visitors pass straight through; attacks, abusive bots and scrapers are slowed, checked or stopped. Every decision is recorded with its reasons.

This guide is for website owners and their staff who **read** the shield's console: what is happening, what was stopped, and what the reports say. Changing the protection is done by your administrator (see [Asking for changes](#)).



The Overview: what the shield did today, at a glance


What you can do with a read-only account

You can	You cannot
See the Overview, Live page, events, incidents, IDS views and reports	Change the protection mode, rules or settings
Open any decision and see why it was made	Ban or unban addresses
Build reports, print them and download CSV files	Manage users, licences or updates

You can	You cannot
Customise your own Overview	
Change your own password and set up two-factor sign-in	

A **READ-ONLY** tag in the header reminds you that your account views and reports only. Some organisations give staff extra permissions; if a button is locked for you, your role does not include it.

Signing in



The sign-in page

1. Open the console address your administrator gave you.
2. Enter your **Username** and **Password**, then **Sign in**.
3. If two-factor sign-in is on, enter the **6-digit code** from your authenticator app.


If the page says **Sign in with your GUARD account**, your organisation's accounts are managed centrally: use the same username, password and code as on your SENTRA GUARD.

TIP

After 8 wrong passwords from one place in 15 minutes, the console asks you to wait 15 minutes. If you forgot your password or lost your phone, ask your administrator to reset it.

Your profile

Click your name at the top right to change your password, set up **two-factor sign-in** (scan the QR code with any authenticator app and enter the code), or **Sign out**. Sessions end after 12 hours.



Your profile

Where to go next


- [Reading the Overview](#)
- [Watching Live](#)
- [What the verdicts mean](#)
- [Looking into an event](#)

- [Reports](#)

Reading the Overview

What each number on the Overview means, and how to arrange the board for yourself.

The **Overview** is the first page you see. It refreshes every few seconds.



The Overview

The line at the top says whether the shield is **ENFORCE** (stopping attacks) or **MONITOR** (only recording them while the website is being tuned).

The figures

Figure	What it tells you
Inspected today	How many requests the shield checked since midnight, and how many per second right now
Stopped (or Would stop in Monitor)	Requests stopped today, and how many looked suspicious
AI agents verified	Search engines and AI assistants that proved who they are, and impostors caught pretending
Swarms tracked	Groups of many addresses acting as one attacker


Figure	What it tells you
Decoy pages	Fake pages served to scrapers that kept coming back
Active bans	Addresses currently refused completely
Inspection time	How long the shield takes to decide — usually far less than a millisecond
Swarm signatures	Recognised attacker toolkits being stopped on sight
Protected sites	Your websites, and whether each one's server is answering
Top threats today	The kinds of attack seen today
Value saved (30d)	An estimate of the money the shield saved in the last 30 days
Latest decisions	The newest decisions: time, address, page, and pass , check or drop

NOTE

A high **Stopped** number is not bad news: it is the shield doing its job. Worry when **Protected sites** shows a site that is not answering, or when your administrator tells you about an incident.

Arranging your board


Press  **Customize** to make the board yours: drag widgets or use the arrows to reorder them, make them **Wide** or **Narrow**, hide them with **X**, or add hidden ones from **Add widgets**. **Save** keeps your layout for your account, on any computer. **Reset to default** returns to the standard board.



Customising the Overview

Full screen shows the board alone — handy on a meeting-room screen. Press Esc to leave.

On a phone the board stacks into one column:




The Overview on a phone

Watching Live

The Live page in plain words — the pipeline, the verdict feed, swarms, sites and where traffic comes from.

Live shows the shield deciding, as it happens.




The Live page

The pipeline

Each moving dot is a real request passing through the shield's five checks on its way to your website:

Colour	Meaning
Blue-grey	Arriving
Green	Cleared and sent to the website
Amber	Challenged or slowed down
Red	Blocked
Violet	The AI engine is looking at it



The pipeline and the verdict feed

Beside it, **Verdicts** lists each decision. **ALL** shows everything, **FLAGGED** only what was checked or stopped. Lines tagged **DRILL** are simulated attacks your administrator ran to test the shield — they never reach your website. Click any line to see why the decision was made (see [Looking into an event](#)).

The rest of the page

Panel	What it shows
Top figures	Requests per second, inspected and blocked today, AI agents caught, inspection time, staff signed in
Swarm radar	Groups of many addresses acting as one attacker. Your administrator can contain them with one click
Site cards	For each website: requests per second, today's requests, flagged requests, response times and recent errors. NOT ANSWERING means the website's server did not respond
Threats today	Today's attacks by kind
Intent engine	How much pressure suspicious traffic puts on the shield, minute by minute
Where traffic comes from	Today's traffic by country, or by network name for your office and VPN
Who is on now	Staff signed in to your applications in the last 15 minutes
Busiest endpoints	The most requested pages today

The **wall** button in the header turns Live into a full-screen board for a TV; **exit wall** or Esc leaves it.

What the verdicts mean

Allow, challenge, investigate and block; what was done to each request; Monitor and Enforce; the risk score; the kinds of threat.

The risk score

Every request gets a **risk score** from 0 to 100 %, made from what five checks found: where it came from, what kind of client sent it, how it behaves, what it is trying to do, and your organisation's own rules. Two lines turn the score into a verdict (by default 60 % and 85 %).

The four verdicts

Verdict	Means
ALLOW	Nothing suspicious enough to act on
CHALLENGE	Probably hostile: check it before letting it through
INVESTIGATE	A browser visit the AI is double-checking; the visitor sees a short "being checked" page
BLOCK	Clearly hostile: stopped

What was done

The event list also says what happened:

Shown as	Meaning
allowed	Sent to the website
blocked	Stopped; the visitor saw a "Request blocked" page with a reference number
challenged	A browser had to pass an invisible check (it solves a small puzzle by itself; people notice nothing)
passed the browser check	The visitor passed the check earlier and was let through
slowed down	A script or API call was delayed a little, then let through
let through (signed-in staff)	A signed-in staff member was let through despite a borderline score
monitor mode — not stopped	The shield is in Monitor mode: this would have been stopped, but was only recorded
would block / would challenge	The same, for drills and Monitor mode

Monitor and Enforce

- **MONITOR** — the shield checks and records everything but stops nothing. New websites run like this for a few days.
- **ENFORCE** — the shield stops attacks.

The mode is shown at the top of every page.

Kinds of threat


Kind	In plain words
Exploit attempt	Trying to break into the website with malicious input (SQL injection, scripts, path tricks)
Vulnerability scanner	An automated tool looking for weaknesses
Path probing	Looking for hidden files such as backups and configuration
Credential stuffing	Trying many stolen usernames and passwords
AI crawler	A bot collecting content for AI training
Autonomous AI agent	An automated agent working through your site
Headless / automation	A program pretending to be a browser
Spoofed browser	A client lying about which browser it is
Script client	A program (not a browser) asking for pages
Rate flood	Far too many requests from one address
Blocked network	An address that is banned, on a block list, or from a blocked country or network
Prompt injection / Hidden instructions	Text trying to manipulate an AI feature on your site
Data exfiltration	Reading large amounts of data, especially sensitive areas
Policy violation	Breaking one of your organisation's own rules

Looking into an event

Finding decisions in the event log, filtering them, and reading one event's full record.

The event log

Monitor > Events lists flagged requests, sign-ins and a sample of ordinary traffic, for the last 40 days.



Monitor > Events

- The chart at the top shows requests and flagged requests over 24 H or 30 D.
- Filter by verdict (**Block**, **Challenge**, **Allow**), by website, by kind of attack, and by **Real only** / **Real + drills** / **Drills only**.
- **Load more** shows older events.

Monitor > Who lists every address seen in the last 30 minutes, riskiest first; click one to see its events.

Monitor > Incidents groups related events into cases, with a timeline and the shield's assessment.

One event

Click any event to open its record.

BLOCK

EVENT #548

CLOSE

Address banned until 12:41 - 3 blocked requests in 5 min (Script client)

blocked · Blocked network

risk 100% · challenge at 60% · block at 85%

REQUEST

WHEN

2026-10-27 12:12:05 MYT

PLATFORM

example.com

FROM

198.51.100.214 · Vietnam

BANNED

USER

not signed in

REQUEST

GET /products?page=6

ANSWER

200 · 4.1 KB out

REPEATS

1 similar requests from this address in the 10 s before were counted but not stored separately

INSPECTION

400 µs

CLIENT

python-requests/2.32.3

WHY - THE CHECKS

NETWORK

Address banned until 12:41 - 3 blocked requests in 5 min (Script client) · Blocked network

100%

FINGERPRINT

Script client asking for a web page - not a browser: python-requests/2.32.3 · Script client

90%

THIS CLIENT

REQUESTS

6 since 12:11:49

FLAGGED

6 · Script client 3, Blocked network 3

HEADERS

Accept-Encoding: identity

Host: www.example.com

X-Real-IP: 198.51.100.214

User-Agent: python-requests/2.32.3

Accept: */*

ASK COPILOT ABOUT THIS EVENT

UNBAN 198.51.100.214

ALLOW-LIST THIS ADDRESS

BLOCK-LIST THIS ADDRESS

ALL FROM THIS ADDRESS

An event record

Part	What it tells you
Top	The verdict, the reason in one sentence, and the risk compared with the challenge and block lines
Request	When, which website, from which address and country, whether the visitor was signed in, which page
Answer	What the shield or website sent back
Client	The program or browser that made the request
Why — the checks	Each finding, which check made it and how much it weighed
AI analyst	The AI's opinion, if it reviewed the request

Part	What it tells you
This client	What the same address did recently
Headers / Body	The technical details (passwords are removed)

All from this address shows every event from the same address.

A customer says they were blocked

1. Ask for the **reference number** on the "Request blocked" page, the time, and what they were doing.
2. Find the event around that time (filter on **Block**, or search the address in **Who**).
3. Read **Why — the checks**.
4. Send your administrator the event number and the reason (see [Asking for changes](#)).

Reports

Building a report, printing or saving it as PDF, downloading CSV, and reading the shadow review and value report.

Tools › Reports turns the shield's records into reports you can share.

Building a report

The screenshot displays the SENTRA AI Threat Shield interface. The top navigation bar includes 'OVERVIEW', 'LIVE', 'MONITOR', 'SHIELD', 'TOOLS' (selected), and 'ADMIN'. The 'TOOLS' dropdown menu is open, showing 'REPORTS' (selected), 'Build a report', 'Saved and scheduled', 'Archive', and 'Overview'. The 'BUILD A REPORT' section is active, showing a 'pick what goes in, then build' prompt. Below this, there are 'SECTIONS' and 'Quick picks' (EXECUTIVE, SECURITY OPERATIONS, COMPLIANCE AND AUDIT, OPERATIONS, EVERYTHING, CLEAR). A list of sections includes EXECUTIVE SUMMARY, THREATS, INCIDENTS AND CAMPAIGNS, BANS, BOTS AND AI AGENTS, AI ANALYST, TRAFFIC PER SITE, SYSTEM PERFORMANCE, and ADMIN ACTIVITY. Below the sections, there are filters for PERIOD (Last 7 days), SITE (All sites), CLIENT ADDRESS (OPTIONAL) (203.0.113.7), SEVERITY (All severities), ROWS IN TOP LISTS (10), and REPORT TITLE (OPTIONAL) (Security report). Buttons for BUILD REPORT, PRINT / SAVE AS PDF, DOWNLOAD CSV, SAVE THIS REPORT, and KEEP A COPY IN THE ARCHIVE are visible. A note states: 'Times are in the shield's time zone (UTC+8). Request totals are exact; attack details come from recorded security events. Reports never contain passwords, keys, tokens or request contents.'

The 'SECURITY REPORT' section shows a date range from 2026-10-20 12:14 to 2026-10-27 12:14, filters for all sites, all clients, all severities, and Mode: Enforce. It includes buttons for PRINT / SAVE AS PDF and CSV, ALL SECTIONS.

The 'EXECUTIVE SUMMARY' section shows headline figures, the change against the previous period and the key findings. It includes a table with the following data:


REQUESTS INSPECTED	FLAGGED AS BAD	BLOCKED	ATTACKING ADDRESSES	INCIDENTS	BANS	DATA SERVED	AI REVIEWS
568 new in this period	117 20.6% of requests - new in this period	107	20 from recorded security events	8 4 high or critical	10	1022.2 KB new in this period	0 0 judged malicious

Below the table, there are bullet points:

- 568 requests were inspected and 117 (20.6%) were flagged as bad, of which 107 were blocked.
- The most common attack was blocked network (10 recorded events).
- 28 events were critical and 4 high severity.

Build a report

1. On **Build a report**, choose the sections — or a quick pick: **Executive**, **Security operations**, **Compliance and audit**, **Operations**, **Everything**.
2. Choose the **Period** (last 24 hours, 7 days, 30 days, this month, last month or a custom range), the website, and optionally one client address, a minimum severity and a title.
3. Press **Build report**.



A built report

- **Print / Save as PDF** opens a print-ready page; use your browser's *Save as PDF*.
- **Download CSV** gives the data for a spreadsheet; each section also has its own **CSV** button.

Reports never contain passwords, keys, tokens or the content of requests. Times are in the shield's time zone.


TIP

For a monthly management meeting, use **Executive** with **Last month**. For the security team, use **Security operations** with **Last 7 days**.

Saving reports, schedules and e-mailing them are set up by your administrator; scheduled reports land in the **Archive** tab, where you can view and download them.

Shadow review

On the **Overview** tab of Reports, **Shadow review** lists everyone the shield **would have stopped** (Monitor mode) or **did stop**, grouped by address and reason.




Shadow review

If you recognise a legitimate address — your payment provider, a partner, your own office — tell your administrator, who can trust it.

Value report

An estimate of what the shield saved: bandwidth and server work not wasted on bad traffic, SMS/OTP abuse prevented, and staff time saved handling incidents. Every line shows its calculation. **Open printable report** gives a one-page summary for management.




Value report

Asking for changes

What your administrator can change, what to include in a request, and what to do in an emergency.

With a read-only account you see everything but change nothing. Changes to the protection go through your administrator (your organisation's own admin, or your SENTRA service provider), so that every change is approved and recorded.



A read-only account sees the policy but cannot change it

Common requests

You need	Ask for
A customer or partner is being blocked	Trust their address or network, or a narrower rule for the affected page
An address keeps attacking	A ban or a block-list entry
No traffic from a country or hosting provider	A geo-fence or network (ASN) block
A search engine or AI assistant should (or should not) read the site	A change to its agent policy
An admin area should only open from the office	An admin zone for those pages, with your office network

You need	Ask for
A weekly report by e-mail	A scheduled report
More access for a colleague	A user account or a different role

What to include

- **What happened** — in a sentence.
- **Event number or block reference** — from the event record or the "Request blocked" page.
- **Address, time and website** involved.
- **What you want** — and for how long (a temporary rule can expire by itself).
- **Who approves** — if your organisation requires it.

In an emergency


If the website is down or under heavy attack, call your administrator or service provider straight away rather than writing. Tell them what you see on the **Overview** (a site *not answering*, a jump in **Stopped**) and on **Live**.

Everyday workflows

A daily check, a weekly review, handling a blocked customer, and spotting an attack.

Daily check (5 minutes)

1. **Overview** — is every site in **Protected sites** *answering*? Anything unusual in **Top threats today**?
2. **Monitor** › **Incidents** — any new **High** or **Critical** incident? Read its summary.
3. **Live** — glance at the swarm radar.
4. Note anything odd and tell your administrator.



A read-only Overview

Weekly review (20 minutes)

1. Build a **Security operations** report for the **last 7 days**; save it as PDF.
2. Look at **Shadow review** for addresses you recognise as legitimate.
3. Compare **Stopped** and the top attack kinds with last week.
4. Share the PDF with your team.

A customer was blocked

1. Get the reference number, time and what they were doing.
2. Find the event (**Monitor** › **Events**, filter **Block**) and open it.
3. Send the event number and reason to your administrator, asking for a narrow fix.

Something looks like an attack

Signs: **Stopped** jumps, a swarm appears on the radar, a site stops answering, a **Critical** incident opens.

1. Note the time, the incident or campaign number and what you see.
2. Contact your administrator straight away.
3. Keep watching **Live**; do not try to contact the attacker's addresses.

Monthly

Build an **Executive** report for **last month** and attach the **value report** for management.

FAQ

Short answers for website owners and analysts.

Does the shield slow my website down? No noticeable delay: it decides in well under a millisecond on average (see *Inspection time* on the Overview).

Will real customers be blocked? The shield is tuned to stop attacks, not people. Browsers that look borderline get an invisible check they pass automatically. If someone is blocked by mistake, the block page shows a reference number so it can be fixed quickly.

What are the DRILL entries? Simulated attacks your administrator runs to prove the shield works. They never reach your website.

Why do I see Google, OpenAI or Anthropic bots? Search engines and AI assistants read websites. The shield checks they are genuine; fakes are stopped as impostors. Whether AI *training* crawlers may read your site is a policy your administrator sets.

Is my data sent to a cloud? No. The shield runs on your organisation's own server. Reports never contain passwords or the content of requests.

Why can't I change anything? Your account is read-only so that every change is approved and recorded. See [Asking for changes](#).

Can I use the console on my phone? Yes; the pages adapt to small screens.

How far back can I look? Events are kept for 40 days; reports can cover any period within that.

Quiz

Six questions to check you can read the shield's console.

QUIZ Check your understanding

6 questions

1. The top of the Overview says SHIELD · MONITOR. What does that mean?

- A Attacks are being stopped
- B **Everything is checked and recorded, but nothing is stopped** ✓ answer
- C The shield is switched off
- D Only monitoring staff may sign in

Monitor mode records what would be stopped without stopping it. Enforce mode stops attacks.

2. An event is tagged DRILL. What is it?

- A A real attack from overseas
- B **A simulated attack run to test the shield, which never reached the website** ✓ answer
- C A failed sign-in by staff
- D A website error

Drills are safe self-tests, inspected by a separate copy of the engine and never sent to the website.

3. A customer says they were blocked. What do you need from them first?

- A Their password
- B **The reference number on the block page and the time** ✓ answer
- C A screenshot of their desktop
- D Nothing; ask the administrator to switch the shield off

The reference number and time let you find the exact event and its reason.

4. Which part of an event record tells you why it was stopped?

- A Headers
- B **Why — the checks** ✓ answer
- C Answer
- D This client

Why — the checks lists each finding, which check made it and its weight.

5. Where do you find everyone the shield would have stopped during Monitor mode?

- A Admin › System
- B **Shadow review in Tools › Reports** ✓ answer
- C The sign-in page
- D Shield › Agents

Shadow review groups would-be-stopped traffic by address and reason, so legitimate visitors can be trusted before enforcing.

6. You need a monthly summary for management. Which quick pick fits best?

- A Operations
- B **Executive, period Last month** ✓ answer
- C Compliance and audit, Last 24 hours
- D Everything, custom range of one day

The Executive pick gives the summary, threats, incidents and value — the right level for management.