IOC Reputation Checker

User guide

https://mba360.net

Check IP addresses, IP ranges, domains and file hashes against VirusTotal, AbuseIPDB and other threat-intelligence sources in one place. Results are grouped by IOC type, summarised by AI, and ready to copy into tickets and reports.

By: Mahdi Bathallath

October 2026

1. Overview

The IOC Reputation Checker takes a list of indicators of compromise (IOCs), looks each one up on the threat-intelligence sources you select, and shows a verdict per source. An AI model running on the site's own server then writes a short threat intel summary of each IOC. IPv4 ranges up to /24 can be entered too: every address in the range is checked and the range gets one summary. Lookups run on the site's server, so API keys never reach your browser unless you enter your own.

Anyone can open the site and use it with the free sources or their own API keys. The VirusTotal and AbuseIPDB keys saved on the server, and the AI summaries, are unlocked with an access password (section 9).

The page has these areas:

- **Main page:** paste IOCs, run the check, and read the results and AI summaries.
- **Thresholds menu** (top left): decide when VirusTotal and AbuseIPDB results count as malicious, suspicious or clean.
- **User guide button** (top right): opens this guide in a new tab.
- **API keys menu** (top right): add your own keys and choose which extra sources to use.

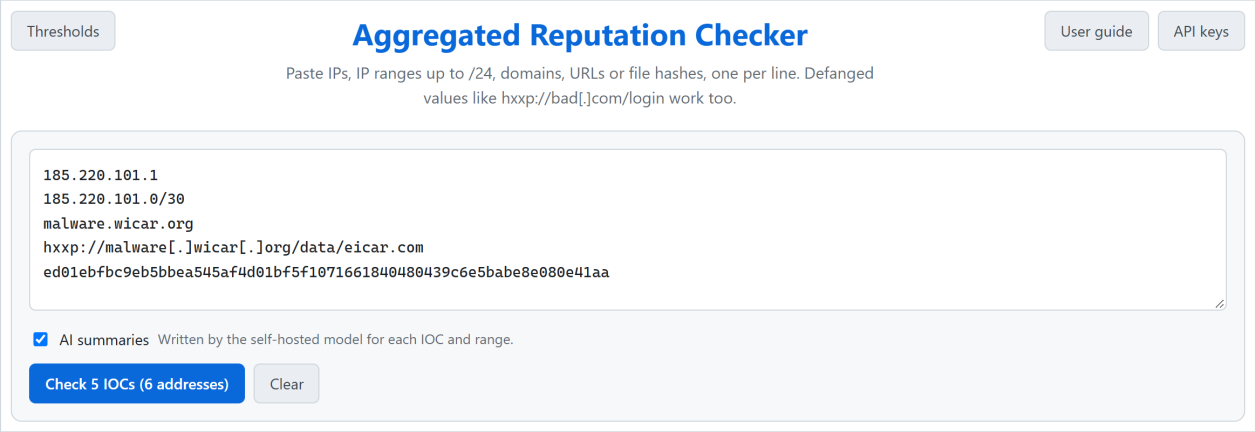


Figure 1. The main page with an IP, an IP range, a domain, a defanged URL and a file hash ready to check. The AI summaries box turns the AI on or off; the button counts the IOCs and, when a range is included, the addresses that will be looked up.

2. Checking IOCs

Paste one IOC per line into the box and press **Check**. The button shows how many IOCs will be checked, with duplicates merged so nothing is looked up or listed twice: an IP inside a range you also entered is listed under that range, a range inside a bigger one is folded into it, the same IOC written two ways (8.8.8.8 and 8.8.8[.]8) is checked once, and a domain with its mail forms (gmail.com, @gmail.com, bob@gmail.com) is one @gmail.com check covering both its reputation and its email security.

Accepted input

Type	Examples	Notes
IPv4 / IPv6	185.220.101.1 2001:4860:4860::8888	Private and reserved IPs (10.x, 192.168.x, 127.x...) only go to VirusTotal and AbuseIPDB.
IPv4 range	198.51.100.0/24 185.220.101.0/29	Ranges from /24 (254 addresses) to /32. Every address is checked; see section 4. Larger ranges such as /16, and IPv6 ranges, are refused.
Domain	example.com	Checked for reputation.
Mail domain	@gmail.com user@gmail.com	A leading "@" (or an email address) checks the domain's reputation plus its email security: SPF, DKIM and DMARC. The "@" is left out of the results. See "Email security checks" below.
URL	https://bad.com/login bad.com/login	Checked as the full URL by the sources that support URLs (VirusTotal, OTX, ThreatFox, URLhaus, urlscan.io, Pulsedive, Safe Browsing, Recorded Future, Site Review). Shown with domains, using the domain template.
File hash	MD5, SHA-1 or SHA-256	Each source only runs on the hash types it supports.
Defanged	hxxps://example[.]com/a 198.51.100[.]0/24	hxxp, [.], (.), {.} and [:] are converted back automatically.

Email security checks

For a mail domain (entered as `@gmail.com`), an **Email security** row appears under the domain, for example *3 of 3 found* with *SPF found (~all), DKIM found, DMARC found (p=none)*. Click it to open the three checks, each marked **found** or **not found** with its details; the dot colour shows whether the setting found is strong (green) or weak (amber). The records are looked up in public DNS:

Check	What it looks at	Strong	Weak, or not found
SPF	Which servers may send the domain's mail, and what happens to others (its "all" rule). A redirect= is followed. Also lists the domain's mail servers (MX).	-all (reject) or ~all (mark as suspect)	No record, ?all, +all, several records, or more than 10 DNS lookups
DKIM	Signing keys under 38 common selector names (Google, Microsoft 365, Mailchimp, Zoho, Proton...).	A key is published	Not found means none at those names; selectors cannot be listed, so the domain may use another (shown grey, not as a failure)
DMARC	The policy for mail that fails SPF and DKIM, from _dmarc (or the parent domain's).	p=reject or p=quarantine	No record, p=none (monitoring only) or pct below 100

These checks show how well the domain's mail is protected against spoofing; they do not change its reputation status, and their Output line is a fixed sentence rather than the domain template. The AI summary mentions them. Like other results they are cached for 24 hours.

Defanged results

In the **IOC column** (and the IOC column of the CSV export), domains and URLs are defanged so they cannot be clicked or resolved by accident: `hxxp://google[.]com/login`, `google[.]com`. The **Output** column, AI and range summaries, Copy and Copy all keep them as they are. IP addresses, ranges and file hashes are never changed.

While a check runs, the progress text on the right shows which IOC is being checked (and which range it belongs to). IOCs are checked one at a time with no pause between them. Press **Stop** to end early; results gathered so far stay on screen.

With the **AI summaries** box ticked (under the input box), each IOC's AI summary is written in the background once its lookups finish, while the next IOCs are checked (section 5). The progress text says *Writing the AI summaries* while the last ones finish. Untick it for a faster check without AI; the choice is saved in your browser.

24-hour cache: each source's answer for an IOC is kept on the server for 24 hours, so checking the same IOC again is instant and does not use up API quota. VirusTotal and AbuseIPDB results are kept per IOC **and per your thresholds** for that source: change a threshold and check again, and the source is asked afresh, while the results for the old and the new thresholds both stay cached. With **Rescan if older than** set, the cached result's last scan date is compared with it, and VirusTotal is asked again (and rescans) when the scan is too old. Errors such as rate limits are never cached, so Retry always asks the source again. Opening a cached row says when the source was actually asked.

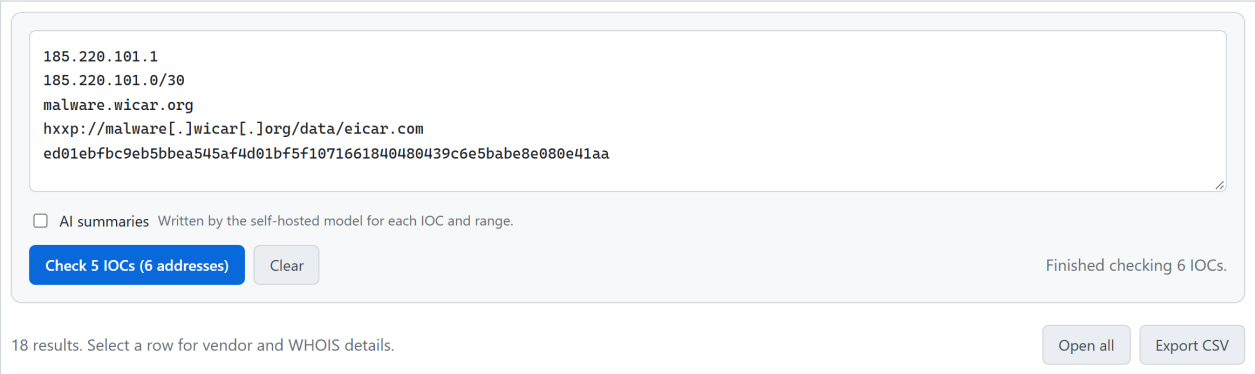


Figure 2. A finished check. The bar below shows the result count, Open all and Export CSV. If any lookup fails, a Retry button also appears here (see section 10).

3. Reading the results

Results are grouped into sections: IP addresses, Domains and URLs, File hashes and Not checked (input that is not a valid IOC, or that no selected source supports). Each source gives one row per IOC. When AI summaries are on, an **AI summary** row with a blue edge sits above each IOC's source rows (section 5), and each IP range is a single row that opens to list its addresses (section 4).

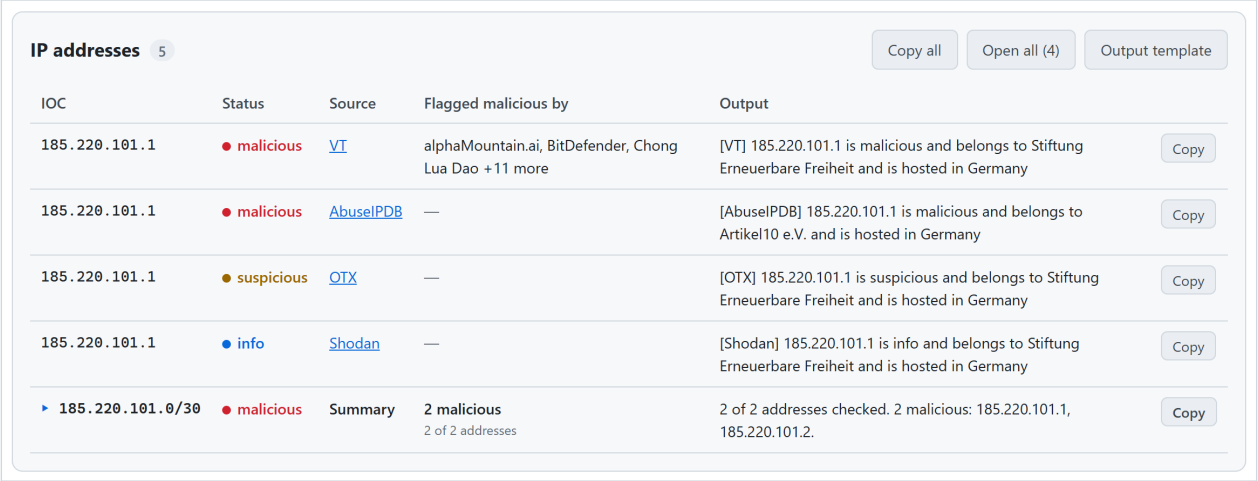


Figure 3. An IP checked against four sources (AI summaries off), and a range below it. OTX and Shodan report no owner, so their rows show the owner and country the other sources found.

Column	Meaning
IOC	The value checked (long hashes are shortened; hover to see the full value). Ranges show a small blue arrow; click the row to open them.
Status	The verdict for this source (see the table below). AI summary and range rows show the most serious status among their sources or addresses.
Source	The source name. Click it to open that source's own report in a new tab. AI-written rows say "AI summary".
Flagged malicious by	For VirusTotal: the security vendors that detected it as malicious. For a range: how many addresses are malicious or suspicious, and how many were checked.
Output	A one-line summary built from the section's output template (section 7), or the AI summary.
Copy	Copies that row's output line.

Status values

Status	Meaning
malicious	The source considers it harmful.
suspicious	Some signs of abuse, but not conclusive.
clean	Known good, allowlisted, or not listed by a blocklist.
info	Context only, not a verdict (owner, open ports, Site Review link).
unrated	The source has too little data to judge.
private	A private or internal IP address.
error	The lookup failed; the Output column says why.

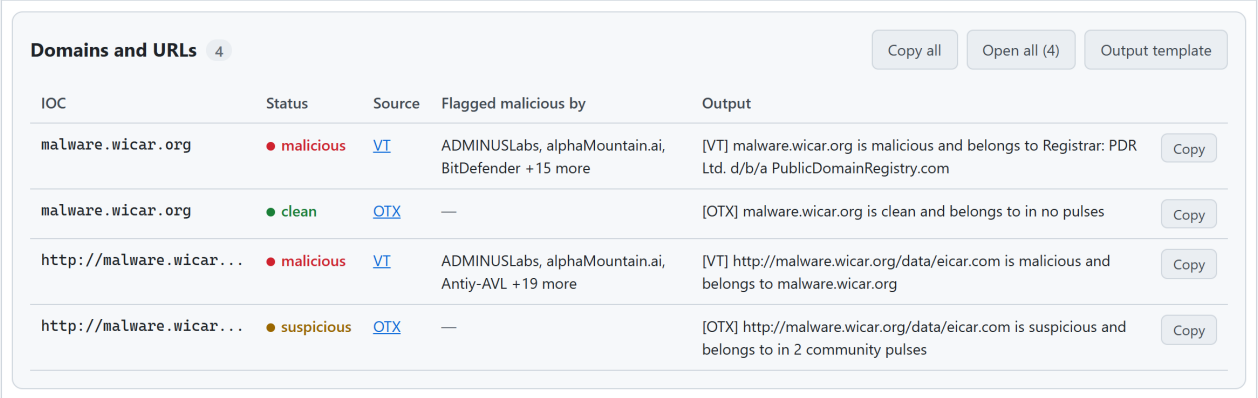


Figure 4. Domains and URLs: the domain and the defanged URL on it are checked separately, the URL as a full URL.

Row details

Click any source row to expand it. VirusTotal rows list every vendor that flagged the IOC and the name each gave it, plus the scan date and WHOIS or file details. For file hashes the details also show whether the file is signed and by whom, its product and copyright, and VirusTotal's threat label. Other sources show their tags and raw details. Click the row again to collapse it.

For a **file hash**, the expanded row starts with **all hashes of the file** (MD5, SHA-1 and SHA-256) with VirusTotal's result, whichever one you checked; the one you entered is marked "checked".

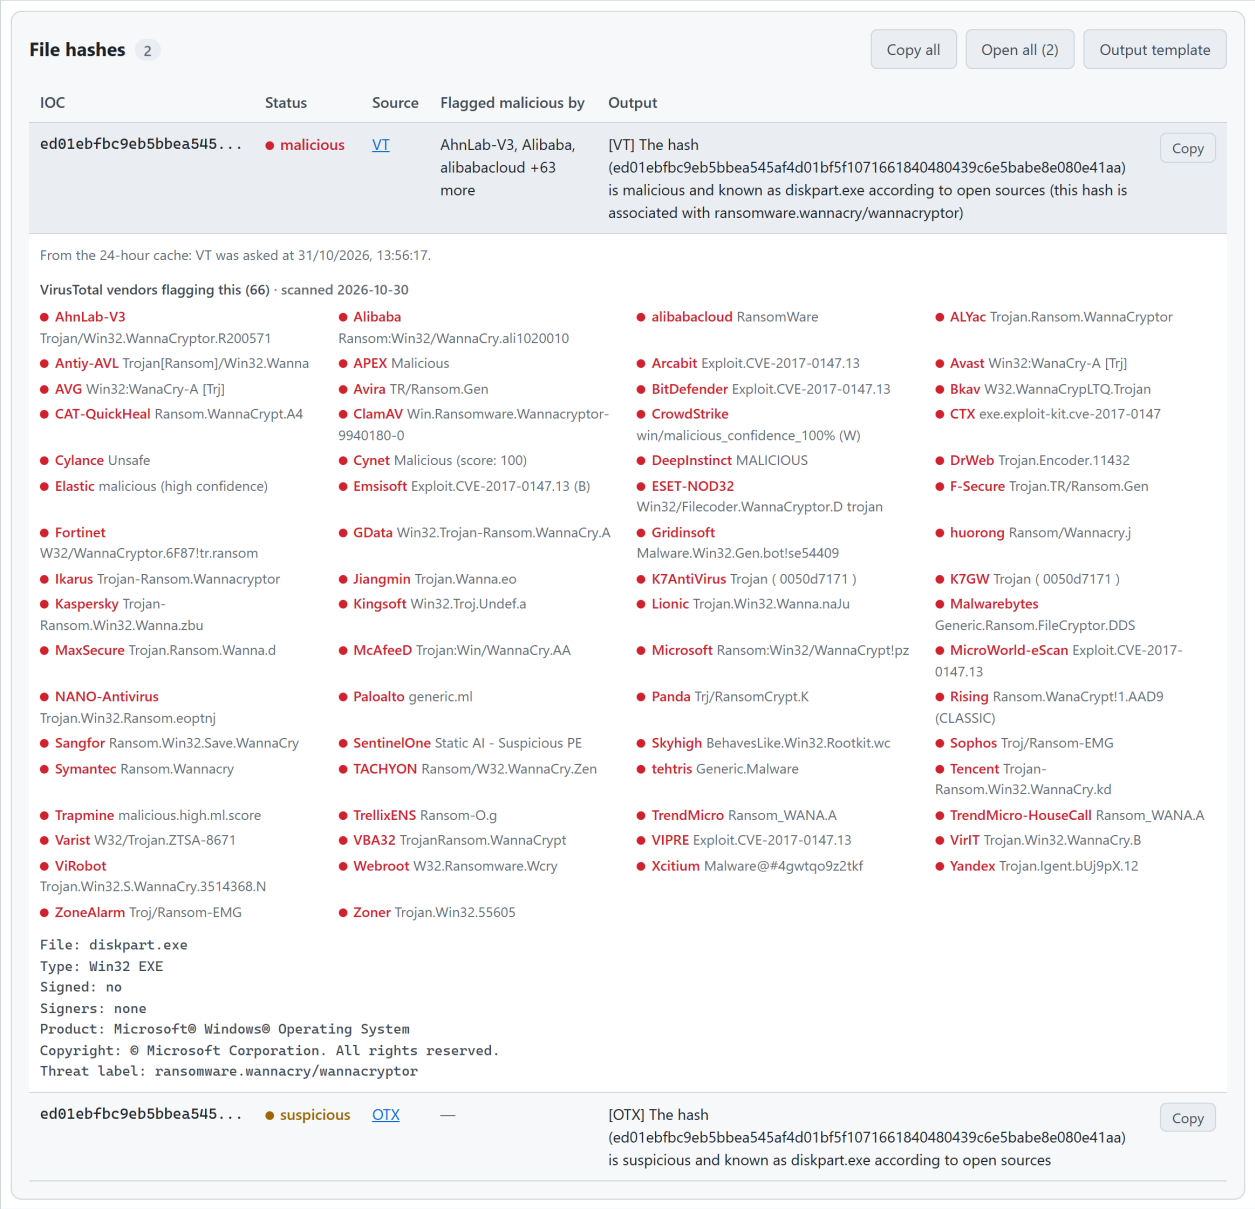


Figure 5. An expanded VirusTotal row for a file hash, listing every detecting vendor and the file details. The first line shows it came from the 24-hour cache.

4. IP ranges (subnets)

Enter an IPv4 range in CIDR form, such as `198.51.100.0/24`, on its own line. The page expands it into its addresses and checks each one like a single IP, with the sources you selected.

- **Size:** /24 (254 addresses) is the largest range accepted; /25 to /32 work too. A /16 or an IPv6 range is listed under Not checked with the reason.
- **Addresses checked:** the network and broadcast addresses are skipped for /30 and larger ranges, so a /24 is 254 lookups and a /29 is 6. A range written from the middle, like 198.51.100.77/24, is treated as 198.51.100.0/24.
- **Check button:** shows the number of addresses that will be looked up, e.g. *Check 3 IOCs (4 addresses)*.

The range row

Each range appears as one row in the IP addresses section. It shows the most serious status of any address in it, how many addresses are malicious or suspicious, how many of the range's addresses were checked, and a summary of the whole range in the Output column. When you are signed in, the summary is written by AI (section 5) and covers every address in the range, not a sample; otherwise the page writes a plain count, such as *6 of 6 addresses checked. 2 malicious: 185.220.101.1, 185.220.101.2.*

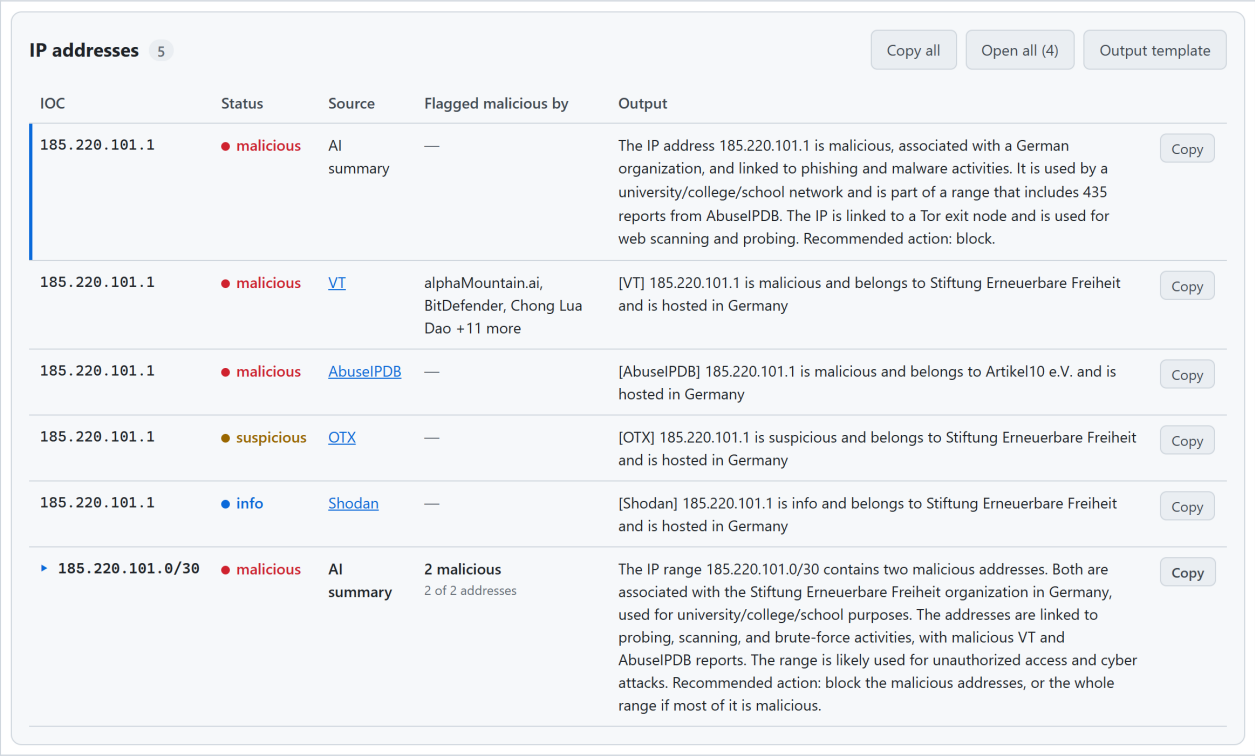


Figure 6. An IP with its AI summary above its source rows, and a /30 range as one row with an AI summary of both addresses.

Click the range row to open it. Every address is listed under it, indented, with one row per source, so you can see the reputation of each address. Click an address row for its vendor and WHOIS details, and click the range row again to close it.

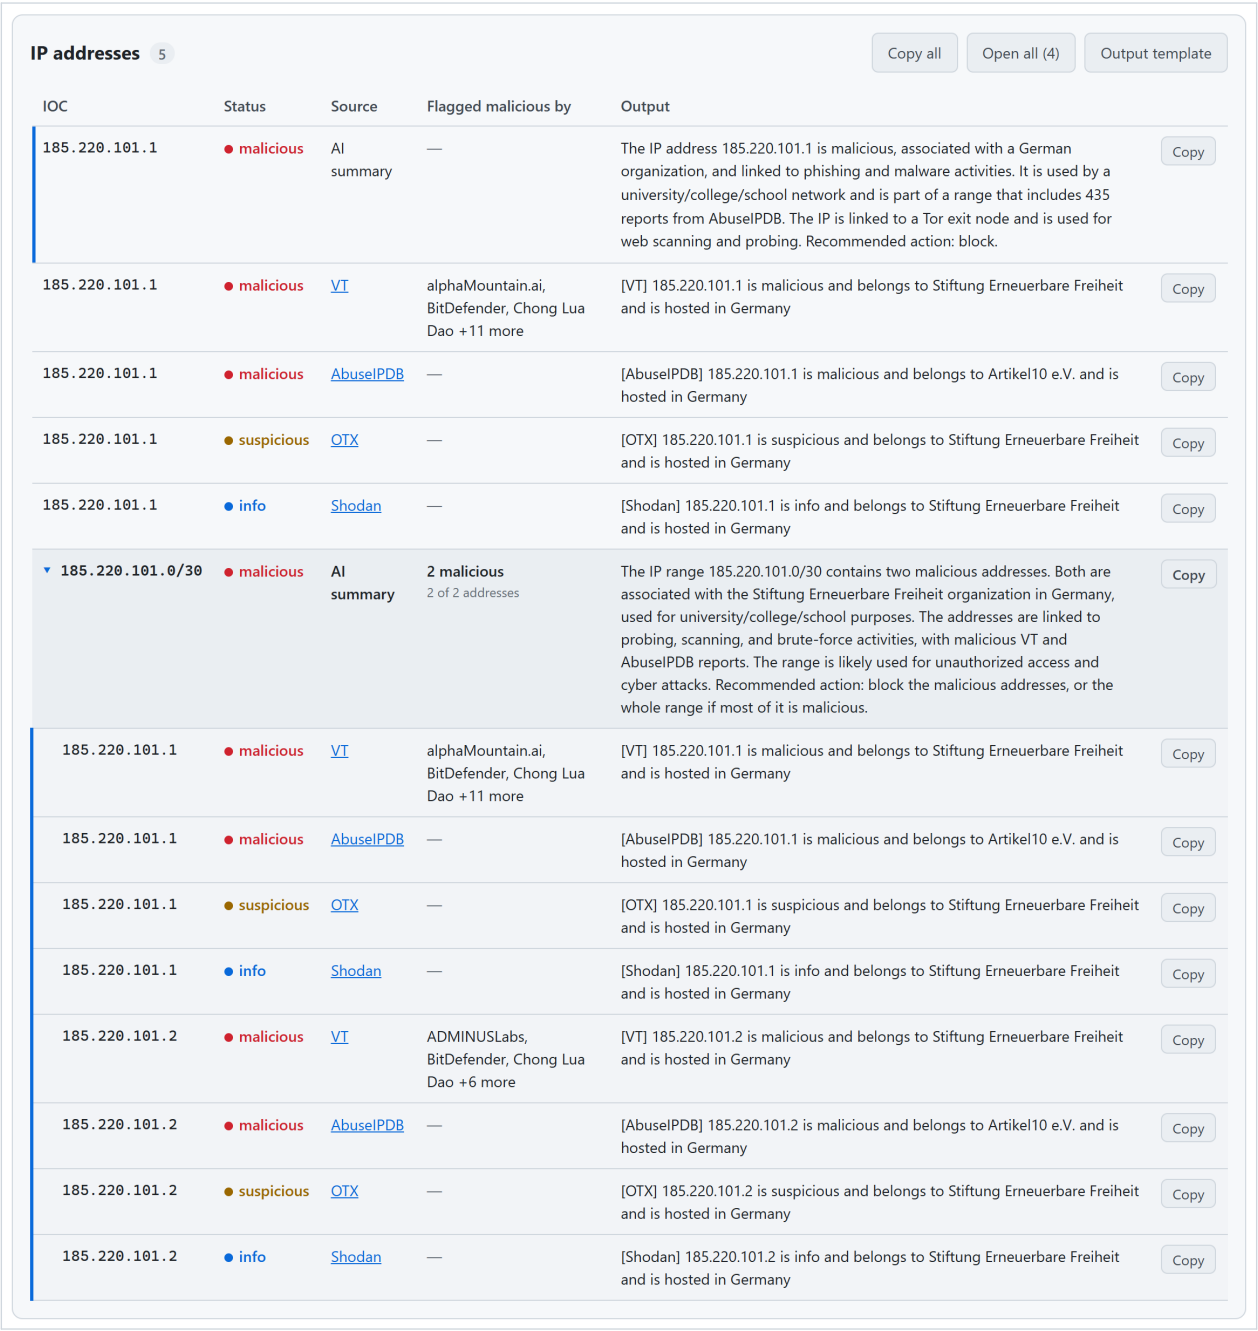


Figure 7. The same range opened: each address with its VirusTotal, AbuseIPDB and Shodan results.

Things to know about ranges

- **Lookups add up quickly.** A /24 is 254 lookups per source. On the server, VirusTotal lookups are shared across several keys (section 10), but very large ranges can still hit the limits; use **Retry failed** for the addresses that failed. A /29 or /28 is a good size for a quick check.
- **Open all** skips the addresses inside ranges, so a /24 doesn't open hundreds of tabs.
- **Copy** and **Copy all** copy the range's summary line, not every address.
- **Export CSV** includes every address, with a Range column naming its range, plus one summary row per range.
- Addresses inside a range don't get AI summaries of their own; the range has one summary covering all of them.

5. AI summaries

After an IOC is checked, the site gives everything the sources returned to an AI model and shows its answer in an **AI summary** row above that IOC's source rows. Ranges get one AI summary of all their addresses.

The model is self-hosted: it is an open-source model (Qwen3 1.7B) running with Ollama on the same server as the site. Your IOCs and their results are never sent to an outside AI service, and there is no per-use cost.

What the model receives

From the sources: every detail each source returned: its verdict and score, the vendors' detection names (often the malware family), tags from OTX, ThreatFox and Shodan, CVEs (also those VirusTotal tags a file with), the owner, country and reverse-DNS hostnames, the file name, type and signer (including a signature that does not verify, such as a revoked or copied certificate), and the start of the WHOIS record.

For a range: every address, grouped. Addresses with the same results (verdict, each source's verdict, owner, country and tags) form one group, written as spans such as 185.220.101.1-14 with the spread of their scores, ports and errors, and what the network's registry record says it is for (e.g. "operating ... Tor relays"). A /24 of similar hosts becomes a few lines, so the whole range is covered and the model still reads it quickly.

From the open web: the official description and CVSS score of each CVE the results name, from NIST's National Vulnerability Database (NVD), and the opening of the Wikipedia article on a malware family the results name (from VirusTotal's threat label or abuse.ch tags), or on the IP or domain itself when it is famous enough to have one. Wikipedia also describes the **organisation behind the IOC**: an IP's network owner or ISP, a domain's company (github.com - GitHub) or a file's signer (Microsoft). An article is used only when it really matches. Published **threat reports** that mention the IOC come from AlienVault OTX, with the sites they cite; the malware families, threat actors (APT groups) and CVEs they are tagged with are taken only from AlienVault's own research, because community reports are often bulk lists tagged with every family in them. The **owner and ISP** come from the official registration record (RDAP, run by the internet registries), e.g. "185.220.101.0/27, registered to Artikel10 e.V.", so they are named accurately. Known malware and threat groups are looked up in **MITRE ATT&CK** (MITRE's own published data, section below). These lookups run in parallel and send only the CVE ID, name or IOC.

Only tools that answer through an API feed the AI: Symantec Site Review (a link only) is left out. So are AbuseIPDB's rough usage label, which often contradicts the registration record. A domain's registrar is left out of the results the model reads, which it would otherwise mistake for the owner, and is given only as a fixed sentence (below). For a clean IOC, threat reports are left out too: they also list harmless addresses, such as public DNS resolvers seen in honeypot logs.

Well-known indicators that the sources misread get a fixed, documented fact instead: the EICAR anti-malware test file, the hash of an empty file (which VirusTotal lists under unrelated file names), the WICAR and Google Safe Browsing test sites, and WannaCry's kill-switch domain (where the advice is *not* to block it, because WannaCry stops when the domain answers).

It also receives the verdict the site worked out from **your thresholds** (the status shown on the AI row) and the matching action: malicious - block, suspicious - investigate, clean - no action, unrated - monitor. The model is told to keep that verdict and action, so the AI never overrules your thresholds; its job is to explain the indicator.

What the summary covers

Each summary is one **PEEL** paragraph: **Point** (the verdict), **Evidence** (each source's result, attributed: "according to VirusTotal"), **Explain** (who is behind it and what it is, and any published reports) and **Link** (the recommended action and why, in MITRE ATT&CK terms). The request is built from a template whose placeholders, such as ***IOC***, ***IOC_TYPE*** and ***SOURCE_VERDICTS***, are filled in for each IOC.

- It opens with the verdict, "This IOC is malicious" (or clean, suspicious...; "This range is ..." for a range), followed by the evidence, each source attributed.
- Facts taken from the results' fields (not from the output lines), in fixed wording: for a **file hash**, its signature ("It is signed by ...", "It carries an invalid signature in the name of ...", or "It is not signed."); for a **domain**, its registrar ("It is registered through the registrar ..."); for an **email domain** (@domain), its SPF, DKIM and DMARC records and what they mean.
- Who is behind it. **IP:** the network owner and AS number, what kind of organisation it is (ISP, mobile carrier, hosting, cloud, Tor, VPN, scanner), what that company does, and the country. **Domain or URL:** the organisation that runs it, what it does and is known for, and how old and popular the domain is (the registrar is given only as the fixed sentence above, so it is never taken for the owner). **File hash:** the file name, type, product and description, and who signed it and what that company is, that its signature does not verify, or that it is unsigned, and its malware family.
- Published reports, with attribution: e.g. "according to 50 AlienVault OTX threat reports, including one citing malcat.fr". The model is told never to mention a report, malware family, threat actor or CVE the sources do not give, and never to guess an attribution.
- The closing action (below), written by the site rather than the AI.

The closing action: MITRE ATT&CK, APT groups and CVEs

The last sentence is put together by the site from the verdict and the evidence, so the small model never has to work it out, and it is always there exactly as built:

- **The action:** malicious network IOCs, block at the network boundary; malicious files, block and quarantine; ranges, block the malicious addresses or the whole range; suspicious, investigate related activity in your logs; clean, "No action is needed."
- **Known malware** named in the results (detection names, threat label) or in AlienVault's tags, with its MITRE ATT&CK ID: WannaCry (S0366).
- **The threat group (APT)** that MITRE ATT&CK records as using that malware, or that the evidence names: Lazarus Group (G0032), Sandworm Team (G0034), APT29 (G0016).
- **CVEs** it exploits: from the detection names (Exploit.CVE-2017-0147), VirusTotal's tags, AlienVault's tags, or the exploits MITRE ATT&CK says the malware uses (EternalBlue is CVE-2017-0144). Each CVE's NVD description is also given to the model.
- **Techniques and tactics:** the two most telling MITRE ATT&CK techniques, with their tactic: those the malware or group uses, or, without one, those the evidence shows (a Tor exit node: Proxy: Multi-hop Proxy, T1090.003, Command and Control; a scanner: Active Scanning, T1595; a document exploit: Exploitation for Client Execution, T1203).

For example: *Recommended action: block and quarantine the file, as it is associated with WannaCry (MITRE ATT&CK S0366), used by the threat group Lazarus Group (MITRE ATT&CK G0032), which exploits CVE-2017-0147, CVE-2017-0144 and uses Service Stop (T1489, Impact) and Data Encrypted for Impact (T1486, Impact).*

Under the summary, **MITRE ATT&CK** lists up to six techniques with their tactics. Each technique ID is a link that opens that technique's own page on attack.mitre.org (for example T1090.003 opens attack.mitre.org/techniques/T1090/003/) in a new tab, and nowhere else.

Summaries are one concise paragraph of at most four sentences and 80 words, or five sentences and 100 words for a range. Each fact is said once: a sentence that mostly repeats an earlier one is removed. If the model pastes the results instead of writing sentences, the server asks it again. When a changed threshold changes an IOC's verdict, the AI **revises its previous summary** rather than writing a new one.

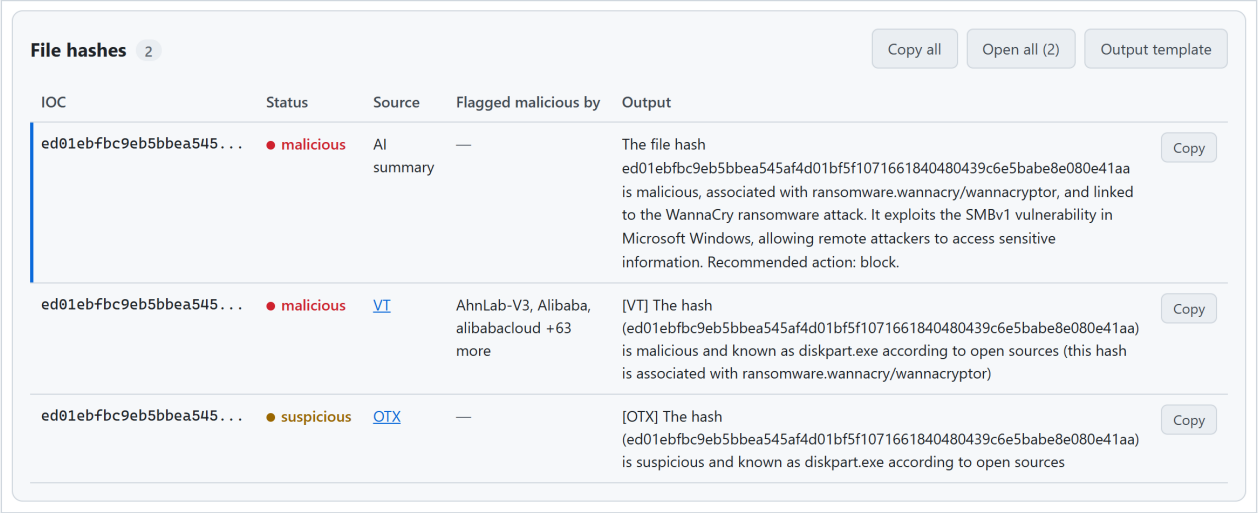


Figure 8. The AI summary of a WannaCry sample: the file it poses as, the ransomware family from the vendor detections, and the recommended action for a malicious verdict.

Good to know

- **Timing:** the model runs on the server's processor, so a summary takes from several seconds to about a minute. **Three summaries are written at once** (in parallel) while the lookups carry on, and a range is summarised as soon as its last address has been checked; the row says *Writing the AI summary...* until it is ready. To keep this fast, the model is loaded when the server starts and stays loaded, every request shares the same instructions so their processing is reused, the results are sent without repeats, and summaries are kept short.
- **Repeats are instant:** checking the same IOC again with the same results returns the stored summary straight away (for 24 hours, like the source results).
- **Turning it off:** untick **AI summaries** under the input box for checks without AI. Ranges then get the plain count summary.
- **Which IOCs get one:** only IOCs that returned data. Private IPs, invalid input and IOCs where every source failed are skipped.
- **Access:** AI summaries are for signed-in users: enter the access password in the API keys menu. Without it, or if the model is not running on the server, IOCs get no AI row and ranges get the plain count summary.
- **A small model:** it is fast enough for a server without a graphics card but can still misread or add details, for example describing a company it knows little about. That is why the verdict, the action, the MITRE ATT&CK entries, threat groups and CVEs are decided by the site, not the model. A larger model (such as Qwen3 4B) is more accurate but slower; it can be set on the server at no extra cost.

- **Thresholds:** the AI text is written once. Changing a threshold re-rates the rows and the status at once, but the summary keeps the wording it was written with; check again for a new one.
- **Check the facts:** AI can be wrong. Treat the summary as a starting point and confirm important findings in the source rows and reports before acting on them.
- **Copy and export:** the AI row's Copy button, Copy all and Export CSV all include the summary.

6. Copying, opening and exporting

Button	Where	What it does
Copy	Each row	Copies that row's output line, AI summary or range summary to the clipboard.
Copy all	Each section	Copies every output line in that section, one per line, including AI summaries and range summaries.
Open all (n)	Each section	Opens each source's report page for that section in new tabs. Addresses inside ranges are left out.
Open all	Results bar	Opens report pages for every result on the page, except addresses inside ranges.
Export CSV	Results bar	Downloads all results as a spreadsheet: IOC, type, status, source, output, flagging vendors, country, owner, last scan date, report URL and range, plus one row per AI summary and per range summary.

Pop-ups: browsers block a page from opening many tabs at once. The first time you use Open all, a message may say tabs were blocked; allow pop-ups for mba360.net in your browser's address bar, then press Open all again.

7. Output templates

Each section has its own output template that controls the Output column, Copy, Copy all and the CSV export for source rows (AI and range summaries are not affected). Press **Output template** on a section to edit it. Click a field button to insert it at the cursor; the preview updates as you type. **Reset** restores the default. Templates are saved in your browser.

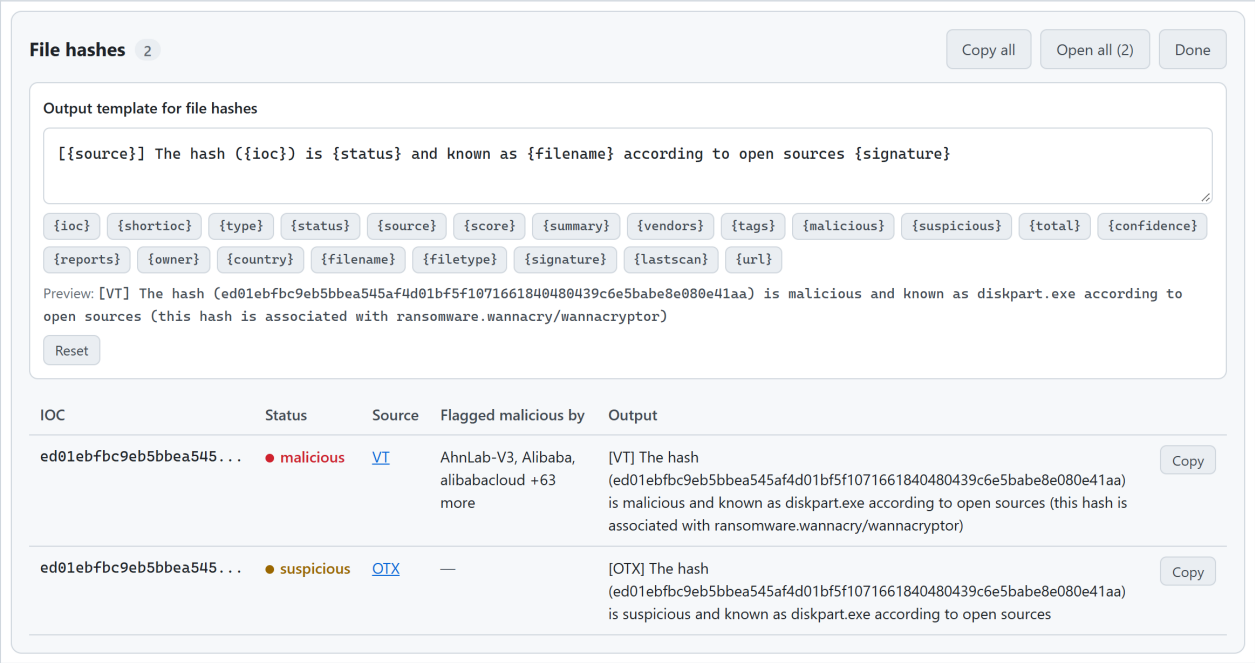


Figure 9. The file-hash template editor with field buttons and a live preview. {signature} adds whether the file is signed and by whom.

Field	Inserts
{ioc} / {shortioc}	The IOC / the IOC shortened to 16 characters

Field	Inserts
{type}	IP, DOMAIN, URL or FILE
{status}	malicious, suspicious, clean, ...
{source}	Source name, e.g. VT or AbuseIPDB
{score}	VT: 12/91 · AbuseIPDB: 100%, 350 reports · others: their own score
{summary}	One line of detail from the source
{vendors}	VirusTotal vendors that flagged it malicious
{malicious} {suspicious} {total}	VirusTotal counts
{confidence} {reports}	AbuseIPDB confidence score and report count
{owner} {country}	Network owner, registrar, URL host or file name; country. A source that reports neither shows what another source found for the same IOC.
{filename} {filetype}	File name and type (hashes); the file name also comes from other sources
{signature}	File hashes only: the signature status. "(this hash is signed by Microsoft Corporation)", "(this hash has an invalid signature in the name of ...)" when the signature does not verify (a revoked or copied certificate), or "(this hash is not signed, associated with ...)" with VirusTotal's threat label. Blank for rows without VirusTotal file data.
{registrar}	Domains only: "(registered through the registrar MarkMonitor Inc.)", the company the domain was bought through (not its owner). Blank for URLs and other sources.
{tags}	Tags reported by the source
{lastscan}	Date of the VirusTotal scan
{url}	Link to the source's report

Default templates

Section	Default template
IP addresses	`[{source}] {ioc} is {status} and belongs to {owner} and is hosted in {country}`
Domains and URLs	`[{source}] {ioc} is {status} {registrar}`
File hashes	`[{source}] The hash ({ioc}) is {status} and known as {filename} according to open sources {signature}`

Every line says what the IOC is: *[VT] google.com is clean (registered through the registrar MarkMonitor Inc.)*, or *[VT] The hash (84c82835...) is malicious and known as diskpart.exe according to open sources (this hash is not signed, associated with ransomware.wannacry/wannacryptor)*. If you never edited a template, the new defaults apply automatically; edited templates are kept, and **Reset** switches one to the default.

Other lines

Lines that are not built from a template follow the same pattern:

Row	Copied line
AI summary	*[AI summary] 8.8.8.8: This IOC is clean, as ...*
Range	*[Range] 185.220.101.0/29 is malicious: 6 of 6 addresses checked...*, or with AI *[Range] 185.220.101.0/29: This range is malicious, as ...*
Email security	*[Email security] microsoft.com has these mail authentication records: SPF found (-all), DKIM found (selector selector2), DMARC found (p=reject)*
A failed lookup	*[AbuseIPDB] 9.9.9.9 could not be checked: AbuseIPDB returned HTTP 500*
No data	*[VT] aaaa...: not found in VirusTotal*

Windows system files are often signed through a Windows catalog rather than inside the file. VirusTotal cannot see those signatures, so such files show "not signed, associated with" instead of "signed by". If you edited a template before {signature} existed, press Reset or add the field yourself.

8. Thresholds

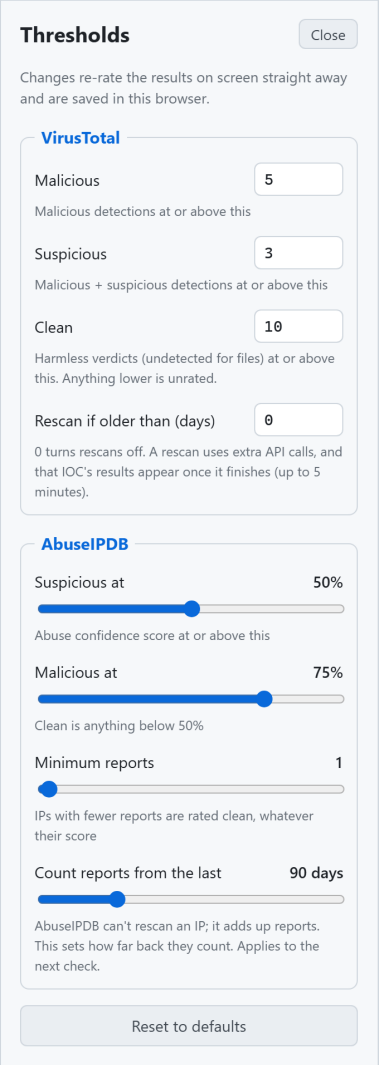


Figure 10. The Thresholds menu.

Press **Thresholds** (top left) to open this menu. Changes re-rate the results already on screen immediately, including range statuses and plain range summaries, and are saved in your browser.

VirusTotal

- **Malicious:** malicious detections at or above this number.
- **Suspicious:** malicious plus suspicious detections at or above this number.
- **Clean:** harmless verdicts (undetected for files) at or above this number. Anything lower is unrated.
- **Rescan if older than (days):** if VirusTotal's last scan is older, it is asked to rescan and that IOC's results appear when the rescan finishes (up to 5 minutes). 0 turns this off. Rescans use extra API calls.

AbuseIPDB

- **Suspicious at / Malicious at:** abuse confidence score sliders. Below the suspicious level is clean.
- **Minimum reports:** IPs with fewer reports are rated clean whatever their score.
- **Count reports from the last:** how many days of reports count (1-365). AbuseIPDB cannot rescan an IP; this applies to the next check.

Reset to defaults restores: VirusTotal 5 / 3 / 10 with rescans off; AbuseIPDB 50% / 75%, 1 report, 90 days.

Other sources rate themselves; their rules are listed in the API keys menu.

9. API keys and sources

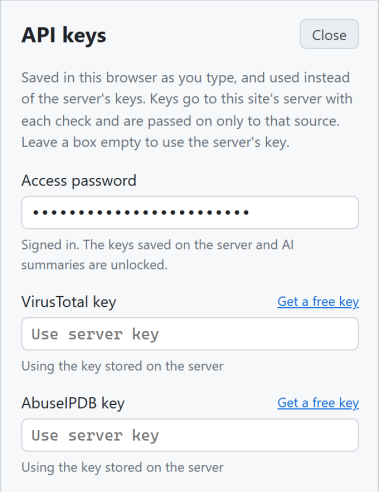


Figure 11. The top of the API keys menu after signing in.

Press **API keys** (top right) to open this menu.

Access password

- Unlocks the VirusTotal and AbuseIPDB keys saved on the server, and the AI summaries. Type it once; it is saved in your browser and the box then says *Signed in*.
- Without it, VirusTotal and AbuseIPDB show "No ... key set" unless you add your own keys, and there are no AI summaries. The free sources work either way.
- Only a salted hash of the password is stored on the server.

API keys

- VirusTotal and AbuseIPDB always run. Once signed in, leave their boxes empty to use the saved keys.
- A key you enter is saved in your browser only, is used instead of the server's key, and is sent only to that service through the site's server.
- You can enter several VirusTotal keys separated by commas; they are used in turn (section 10).
- **Show keys** reveals what you typed; **Clear all saved keys** removes them from this browser.

Threat intel sources

- Tick a source to include it in checks, or use **Select all**.
- Blue text under each source shows the IOC types it supports; it is skipped for others.
- When a ticked source uses a key, a key box appears below it. Required keys must be added or that source shows an error; optional keys raise the free limits.
- ThreatFox, MalwareBazaar and URLhaus share one abuse.ch key.
- More sources give the AI summaries more to work with: OTX pulses and ThreatFox tags often name malware families and campaigns, and Shodan adds open ports and CVEs.
- Your choices are saved in your browser.

Available sources

Source	Checks	Key	How it rates
VirusTotal	IP, domain, hash	Required*	Thresholds menu
AbuseIPDB	IP	Required*	Thresholds menu
AlienVault OTX	IP, domain, hash	Optional	Allowlisted = clean; AlienVault malware = malicious; community pulses only = suspicious
ThreatFox	IPv4, domain, MD5, SHA-256	abuse.ch	Listed at 75%+ confidence = malicious
MalwareBazaar	Hash	abuse.ch	Listed = malicious
URLhaus	IPv4, domain, MD5, SHA-256	abuse.ch	Online URLs = malicious; offline only = suspicious
GreyNoise	IPv4	Optional	GreyNoise's own label
Shodan InternetDB	IPv4	None	Info (ports, CVEs); c2/malware tags = malicious
urlscan.io	IP, domain	Required	Any malicious scan = malicious
Hybrid Analysis	Hash	Required	Sandbox verdict
Pulsedive	IP, domain	Optional	High/critical = malicious; medium = suspicious
CIRCL hashlookup	Hash	None	Known-good file = clean; flagged = malicious
Google Safe Browsing	Domain	Required	Listed = malicious
IPinfo	IP	Optional	Info (owner, location)
Recorded Future	IP, domain, hash	Paid token	Risk 65+ = malicious; 25-64 = suspicious; lower = clean. Details list every triggered risk rule
Symantec Site Review	Domain, email domain, URL (always included)	None	Info: a link to the Site Review page (it needs a CAPTCHA, so it cannot be checked automatically) plus the web categories other vendors assign, via VirusTotal. Categories found for one form of a host (a URL) are reused for the others (the domain, @domain)

* Saved on the server and unlocked with the access password; or use a key of your own.

10. Rate limits, errors and Retry

Free API keys have limits. The most common is VirusTotal's: 4 lookups a minute and 500 a day per key.

- **Server keys:** the server holds several VirusTotal keys and uses them in order, one at a time. Every lookup goes to the current key until VirusTotal stops accepting it; the lookup then moves straight on to the next key, and after the last key the first comes round again. With six keys that is about 24 lookups a minute and 3,000 a day.
- **When a key is refused:** after a rate limit it rests for a minute (the free per-minute window). If it is refused again on its first lookup after that rest, its daily quota is used up and it rests until midnight UTC. A wrong or revoked key rests for an hour.
- **Your own key:** a single key keeps VirusTotal's own limit, so a long list or a range will hit it and the remaining VirusTotal rows show a rate-limit error. Several keys separated by commas are used in turn the same way as the server's.

Retry

- A **Retry n failed** button appears next to Check for errors that can clear up on their own (rate limits, server or connection errors), including addresses inside ranges.
- After a VirusTotal rate limit, the button counts down ("Retry 7 failed in 59s") and unlocks after one minute. Pressing it re-checks only the IOCs that failed, and rewrites their AI and range summaries.
- Once nothing retryable is left, the button disappears.
- Errors a retry cannot fix are not offered: a missing key ("No ThreatFox key set"), a key the service rejected, or input that is not an IOC. Fix these in the API keys menu instead.

Troubleshooting

What you see	What to do
"No ... key set. Add one in the API keys menu."	Add that source's key, or untick the source.
"No ... key set. Enter the access password or add your own key ..."	Enter the access password in the API keys menu, or add your own key for that source.
"Wrong access password"	Check the password in the API keys menu, or clear the box to use your own keys.
"... rejected the API key"	The key is wrong or expired. Re-enter it in the API keys menu.
"... rate limit reached"	Wait for the Retry countdown, then press Retry. For sources with optional keys, adding your own free key gives you your own limit.
No AI summary rows	Enter the access password in the API keys menu. If you are signed in, the AI model is not running on the server; ask the site's administrator.
"AI summary unavailable: The AI model took too long to answer"	The server is busy with other summaries. Check that IOC again later; the source rows are unaffected.
A range under Not checked	It is larger than /24 or an IPv6 range. Split it into /24 ranges.
A range summary that says "Stopped before the AI summary was written"	The check was stopped. Press Check again to finish the range.
Rows under Not checked	The line is not an IP, range, domain or hash, or no ticked source supports that type.
Open all opens one tab only	Allow pop-ups for mba360.net, then press Open all again.
"Could not reach the server"	Check your connection and press Retry.

Your settings (thresholds, templates, sources, the access password and any keys you enter) are stored in your browser. They do not follow you to another browser or computer, and clearing site data resets them.