DEVELOPERS

The Carmen API.

One endpoint. Send the thing, choose the depth, get a verdict with its evidence. Sync for links, people and deepfakes, async for everything else.

Request

`POST /api/analyze-content  Authorization: Bearer <api key>`

FIELD	MEANING
`content.text`	The thing to check.
`content.context`	Optional. Where it came from, who received it.
`options.priority`	low, normal or high.
`options.createCase`	Keep the check as a case you can reopen.
`options.reportFormats`	markdown, pdf, docx or json, when you asked for a report.
`metadata`	Anything you want back with the result.

Depth and credits

Depth is chosen per request.

YOU ASK FOR	WHAT YOU GET	CREDITS
Quick check	A bare link, number or address. Takes the fast safety route.	5
Deep analysis	A message, email or file read properly, with reasons.	25
Advanced analysis	A person or company across public records.	250
Full report	Everything written up as a document.	500

DEPTH CONTROL

Choose depth per request.

A link at Quick check is 5 credits. The same link at Full report is 500. You decide what each check is worth.

Example

```
// Deep analysis on a message
POST /api/analyze-content
{
  "content": { "text": "USPS: your package is on hold. Pay the $2.99 fee at parcel-redelivery.info" }
}

// 202, queued. Poll links.result.
{
  "job_id": "3f9c…", "status": "pending", "mode": "async",
  "links": { "status": "/api/jobs/3f9c…", "result": "/api/jobs/3f9c…/result" }
}

// A bare link (https://acme-banking-login.com/login) answers in the same call: 200, sync
{
  "job_id": "8a21…", "status": "completed", "mode": "sync",
  "result": {
    "verdict": "malicious", "confidence": "high",
    "headline": "This page copies the Acme Bank sign-in screen on an address the bank does not own.",
    "recommended_action": "Close this page and do not enter your password, card number or one-time code.",
    "evidence_trail": [
      { "source": "registry", "signal": "domain_age", "weight": 0.9, "observation": "Registered 4 days ago." }
    ]
  }
}
```

Plans

PLAN	PRICE	CREDITS
Dev	$5.99 / month	500
Starter	$49 / month	10,000
Growth	$299 / month	100,000
Enterprise	On contract	Volume by agreement. Overage one cent a credit on every plan.

Response

FIELD	MEANING
`job_id`	Always present. A UUID.
`status`	pending, completed or failed.
`mode`	sync or async. Decided by the search type, never by which keys happen to be present.
`search_type`	What Carmen decided it was: safety, scam, person, deepfake, diligence.
`case_id`	When createCase was set.
`result`	Sync only. The verdict object.
`links.status, links.result`	Async only. Poll these.

The verdict object

FIELD	VALUES
`verdict`	clean, unknown_fresh, suspicious, malicious. Show them as Low, Medium, Medium, High. unknown_fresh means under 30 days old with no history.
`confidence`	low, medium, high.
`headline`	Customer ready copy, under 280 characters.
`recommended_action`	Customer ready copy, under 500 characters.
`evidence_trail[]`	Up to six. Each has source, signal, weight from 0 to 1, and observation.
`context_insights`	Optional short narrative.

Headline and recommended action are already written for a person to read. Put them on screen as they are.

```
// JavaScript
const r = await fetch(`${BASE}/api/analyze-content`, {
  method: "POST",
  headers: { Authorization: `Bearer ${KEY}`,
    "Content-Type": "application/json" },
  body: JSON.stringify({ content: { text: "https://acme-
bank-login.com" } })
});
const { result } = await r.json();
```

Sync or async

Links, people and deepfakes answer in the same call. Messages, emails and anything with words around a link are queued: you get a 202 with a job id and links, then poll. informational is a result type you can get back when the content is a question rather than a threat.

Good to know

- Every deliberate call draws credits by the depth you chose, whichever internal path it took.
- Carmen decides the search type from the content. You do not need to tell it what you are sending.
- Re-checking identical content within the hour returns the same verdict.
- Keys are issued by request for now. One key per integration.

In four languages

```
BASE = https://api.backgrounder.com
```

```
# curl
curl -X POST $BASE/api/analyze-content \
  -H "Authorization: Bearer $KEY" \
  -H "Content-Type: application/json" \
  -d '{"content":{"text":"https://acme-bank-login.com"}}'
```

```
# Python
import requests
r = requests.post(BASE + "/api/analyze-content",
  headers={"Authorization": f"Bearer {KEY}"},
  json={"content": {"text": "https://acme-bank-login.com"}})
print(r.json()["result"]["verdict"])
```

```
// TypeScript: the same fetch, with a typed result
type Verdict = { verdict: "clean" | "unknown_fresh" |
"suspicious" | "malicious";
  confidence: "low" | "medium" | "high"; headline: string;
recommended_action: string };
const { result } = (await r.json()) as { result: Verdict };
```