

Production Web Retrieval for Financial AI Agents



Join our stack →

One API that gives financial AI agents secure, real-time access to the live web, with the accuracy, speed, and controls regulated environments require.

\$100M+

Projected business impact across risk & compliance

Leading payments network

15 min vs 7 days

Time to produce investment research reports

Global investment manager

100s

Of investigators using AI for AML case work

Top-5 Canadian Bank

How financial institutions use Tavily

🔍 Investment Research

Equip research agents to retrieve filings, earnings-call transcripts, market news, and disclosures in real time, so analysts spend less time gathering facts and more time testing theses.

- Surface filings, news, and earnings calls in real time
- Research thousands of companies at scale
- Generate source-cited briefs for analyst review

🛡️ KYC, AML, & Compliance

Build compliance agents to retrieve adverse media, sanctions exposure, and public risk signals continuously, helping investigators triage entities, resolve alerts, and preserve an audit trail.

- Retrieve adverse media and entity intelligence
- Monitor customers, vendors, and counterparties
- Produce source-backed findings for audit review

📊 Risk Management

Enable risk agents to retrieve regulatory updates, market developments, and counterparty signals as they emerge, helping teams brief leaders sooner and act with confidence

- Track news, filings, and regulatory changes as they happen
- Detect emerging risks before they escalate
- Ground decisions in verifiable sources for review

Why Tavily

/accuracy

Retrieve precise, grounded context from trusted web sources. Ranked #1 on SimpleQA for high-stakes financial AI.

/speed

Keep agent workflows fast with retrieval latency as low as 180 ms, without sacrificing retrieval quality.

/context quality

Return more relevant context with fewer tokens, reducing inference costs while improving downstream reasoning.

/freshness

Access continuously crawled and indexed web context, so agents act on current information, not stale snapshots.

/enterprise-scale

Scale production AI with high-throughput retrieval, enterprise reliability, 99.99% uptime, and infrastructure powered by Nebius.

/interoperability

Works with your AI stack, from LangChain and LangGraph to AWS, Databricks, Snowflake, IBM, and your own enterprise infrastructure.



Join our stack →

Production Web Retrieval for Financial AI Agents

/security & compliance

The most secure web search API on the market,
on the platforms you already use:



✓ Data Protection

Zero Data Retention

Queries and results are never stored

PII detection

Automatically strips PII before it's
included in queries

Encryption in transit

All traffic is encrypted end to end

✓ Access & Governance

SSO / SCIM

Enterprise identity and automated
user provisioning

Exclude domains

Block specific domains from results

SafeSearch

Filters out unwanted or unsafe
sources

✓ Agent & Network Security

Prompt injection defense

Blocks malicious reprogramming
attempts hidden in web content

IP rotation

Searches route through Tavily's
rotating IPs, never your
infrastructure

Proprietary Index

Built on Tavily's own search index to
reduce third-party data exposure

/free

Free 1000
API credits/month

/pay-as-you-go

Starts at \$8 per
1000 credits

/project

\$5-7.5 per 1000 credits with
monthly, prepaid credits

/enterprise

Custom API volume, rate
limits, and enterprise support

Join the ranks of the world's leading financial institutions

Try Tavily Free →

Connect with Team →