Know what the internet can see — before an attacker does

ThreatWorx EASM · a platform SKU, not a point tool

Every organisation's real perimeter is larger than the one it maintains. Forgotten staging boxes, a microsite on someone's personal cloud account, an expired SaaS CNAME, a subdomain pointing at a de-provisioned bucket, an API key baked into a JavaScript bundle.

These are the assets that get breached — precisely because no one is looking at them.

ThreatWorx EASM discovers your internet-facing footprint from the outside in, starting from nothing more than a domain name, and continuously watches it for the changes that create risk. It is a full SKU of the ThreatWorx platform — discovery, attribution, exploit-aware prioritisation, monitoring, alerting, and remediation workflow — not a standalone scanner. Scans are executed by the open **twigs** collector or the scheduled discovery app; everything they find lands in the same inventory, correlation, and alerting pipeline you already run for the rest of your estate.

10 passive subdomain sources, each failure-tolerant	**11** vetted public resolvers, NXDOMAIN-tested	**12** cloud & CDN providers attributed by prefix	**0** agents on assets, paid API keys, or daemons required

WHY THREATWORX EASM

One seed, whole surface

Give it a domain — or a hostname, IP, CIDR, or ASN — and it expands outward: subdomains, sibling domains, co-hosted sites, cloud netblocks, and the services running on all of them.

Passive-first, low-noise

The bulk of discovery is passive — CT logs, DNS, public datasets, archives. Active probing is bounded, rate-aware, and every module can be switched off.

Exploit-aware prioritisation

Findings are re-ranked against CISA KEV and FIRST.org EPSS, so "patch this" means "this one is being exploited in the wild" — not a 900-row CVSS spreadsheet.

Flexible scan execution

Scans run wherever you want them to — laptop, cron box, container, or ephemeral CI runner — via the twigs collector or the scheduled discovery app. No agent on the assets themselves; no paid API keys for core functionality.

One platform, one pipeline

External assets, tags, and config issues land in the ThreatWorx inventory model and console alongside your hosts, containers, cloud accounts, and code — with the same correlation, alerting, ticketing, and reporting the rest of the platform gives you.

HOW IT WORKS

01 **Seed** One or more of: domain, hostname, IP, CIDR, ASN — auto-classified.	02 **Expand** Passive + active enumeration, ASN/CIDR expansion, reverse-WHOIS and reverse-IP pivots.	03 **Resolve** Multi-resolver pool resolves candidates; wildcard-only responses filtered out.	04 **Attribute** Each IP mapped to cloud/CDN provider, region, and service — or flagged target-owned.
05 **Assess** Per host: tech, versions, panels, headers, cookies, methods, JS bundles, ports. Per domain: DNS/email hygiene, takeover, buckets, typosquat.	06 **Prioritise** Findings enriched and re-ranked against KEV and EPSS.	07 **Report** Assets, tags, and config issues emitted into the ThreatWorx platform inventory and console.	08 **Watch** On the next scheduled run, incremental CT polling and re-assessment surface what changed.

CAPABILITIES

Discovery & enumeration

Multi-seed input — domain, hostname, IP, CIDR, or ASN, auto-classified; CIDR and ASN seeds expanded to live hosts under a configurable host cap.

Active brute force with tiered bundled wordlists (~130 / ~5k / ~20k) or a custom list.

Permutation scanning — affix, number-bump, and environment-token mutations (dev / stage / qa / uat…) derived from names already found.

Org-ASN derivation + PTR sweep — infers the organisation's own ASNs from its IPs (rejecting shared hosters), then reverse-resolves every address in the announced prefixes.

Reverse-IP / virtual-host discovery — co-hosted hostnames on an IP, plus Host-header probing for apps reachable with no current DNS.

Passive subdomain enumeration from 10 independent sources — crt.sh, Certspotter, AnubisDB, HackerTarget, AlienVault OTX, urlscan.io, RapidDNS, Subdomain Center, Wayback Machine, Common Crawl.

Wildcard-DNS fingerprinting — brute force still runs on wildcard domains; wildcard-only responses are subtracted.

Vetted multi-resolver pool — 11 public resolvers, each tested for NXDOMAIN honesty, round-robin with cross-resolver retry.

Reverse-WHOIS pivot — organisation and label pivots surface sibling and related domains.

Recursive discovery — optionally re-runs the full assessment against related domains found via reverse-WHOIS, co-hosting, and JS analysis, tapering by depth.

Attribution & context

Cloud / CDN IP attribution — AWS, GCP, Azure, Microsoft 365, Oracle OCI, Cloudflare, Fastly, DigitalOcean, Linode, Vultr, GitHub, Rackspace, via published prefix files plus a ~60-entry ASN map. Tags `CLOUD:`, `CLOUD_REGION:`, `CLOUD_SERVICE:`. IPs matching nothing are tagged `CLOUD:unattributed` — a "likely target-owned, scan harder" signal.

Discovery provenance on every asset — which seed and which source(s) it was found from (`EASM_SEED`, `EASM_SEED_TYPE`, `EASM_ROOT_DOMAIN`, `DISCOVERY_TYPE`…).

Technology & product fingerprinting, with version recovery — curated version-disclosure probes for products detected without a version.

Continuous monitoring

Incremental Certificate-Transparency polling — a persisted per-domain high-water cursor plus a time-window fallback for cold starts and ephemeral runners. crt.sh primary, Certspotter augment.

Emits new-subdomain (LOW) and look-alike-certificate (MEDIUM) findings on each run, with a status / degraded heartbeat.

Designed for scheduled invocation — hourly, daily, or per-CI-run — without a long-lived daemon.

Risk detection

Subdomain / dangling-DNS takeover — CNAME-chain resolution, NXDOMAIN state, and HTTP body/status fingerprinting against the community can-i-take-over-xyz set (auto-refreshed, bundled fallback). Vulnerable service → HIGH; edge case → MEDIUM.

JavaScript bundle analysis — first-party bundles mined for 14 secret patterns (AWS and GCP keys, Slack tokens, Stripe live keys, GitHub/GitLab

Cloud storage bucket discovery — org- and subdomain-derived names probed against S3, Google Cloud Storage, and Azure Blob. Publicly listable → HIGH; exists-but-private → LOW (claimed namespace, future-takeover risk).

KEV / EPSS exploitation enrichment — CISA KEV plus FIRST.org EPSS; monotonic rating escalation, worst-first re-sort, and a per-asset exploitation-

tokens, private-key blocks, JWTs…), API endpoint paths, internal hostnames, and exposed source maps. Live credential → CRITICAL / HIGH.

DNS hygiene — dangling records, weak SPF/DMARC, open resolvers, zone-transfer exposure.

Credential-leak surfacing — organisation mentions in public leak datasets.

Netblock port sweep — masscan → naabu → bounded pure-Python TCP-connect fallback; ~70 top ports per prefix, with risky-port labelling (RDP, SMB, databases, admin panels…).

prioritised summary finding.

Email security posture — SPF, DKIM alignment, DMARC policy strength.

Typosquat / look-alike domains, with ongoing CT watch for newly issued look-alike certificates.

Web exposure checks — exposed admin panels, directory listings, dangerous HTTP methods, open redirects, missing security headers, insecure cookies.

FINDINGS & SEVERITY

CRITICAL	Live cloud credential or private key in a public JS bundle.
HIGH	Subdomain takeover (vulnerable service); publicly listable storage bucket; high-value secret exposure.
MEDIUM	Edge-case takeover; look-alike certificate issued; virtual host serving distinct content with no current DNS; weak email authentication.
LOW	New subdomain observed; private bucket under org naming; exposed source map; historical sensitive path now live.
INFO	Inventory: co-hosted hostnames, discovered subdomains, PTR-named hosts, technology stack, provenance, CT status.

Ratings are dynamic: KEV/EPSS enrichment escalates — never downgrades — CVE-linked findings, and re-sorts each asset's issues worst-first.

DEPLOYMENT & REQUIREMENTS

Licensed as a ThreatWorx platform SKU; findings and inventory live in the ThreatWorx console.

Scans executed by the Python-based twigs collector (`twigs easm`) or the discovery app for scheduled runs.

No agent on the target assets. No inbound access required.

Core discovery needs only outbound HTTPS and DNS. Optional: masscan or naabu on `PATH` for faster port sweeps, an OTX API key for extra passive DNS. Product-specific keys are never required for core features.

Disk cache (default `~/.twigs/cache/easm`, override with `TWIGS_EASM_CACHE_DIR`) shares expensive lookups across runs and persists CT cursors.

Every active module has an opt-out flag; scan rate and wordlist tier are tunable.

Scan execution runs anywhere the collector runs — workstation, scheduler, container, CI.

RUNNING A SCAN

FULL EXTERNAL ASSESSMENT OF A DOMAIN

```
twigs easm --fqdn example.com
```

MULTIPLE SEEDS, LARGER WORDLIST, RECURSIVE

```
twigs easm --seed example.com \
           --seed AS64500 \
           --wordlist_tier large \
           --recursive_discovery
```

LEAN PASSIVE-ONLY PASS

```
twigs easm --fqdn example.com \
           --no_portsweep \
           --no_bucket_discovery \
           --no_js_analysis
```

TWO WAYS TO RUN IT

twigs CLI — run `twigs easm` ad hoc from a workstation, container, or CI job.

Discovery app — deploy github.com/threatworx/discovery_app to schedule regular scans of your external attack surface and keep the inventory continuously up to date.

WHAT SETS IT APART

Depth of a dedicated EASM vendor

With collectors you can read, run locally, and audit line by line.

Exploitation context, not just CVSS

KEV/EPSS re-ranking is built in, not a paid add-on.

CLI-native continuous monitoring

Incremental CT polling with a persisted cursor works correctly from ephemeral runners. No daemon.

One inventory

EASM assets correlate with hosts, cloud accounts, containers, and source code in the same ThreatWorx graph.

Honest coverage signals

Unattributed IPs are labelled as such, so "we don't know" never masquerades as "it's fine".

ABOUT THREATWORX

ThreatWorx is a single platform for continuous, predictive threat detection and prioritised remediation across hosts, cloud accounts, containers, Kubernetes, code repositories, SBOMs — and, with the EASM SKU, everything of yours the internet can reach. Discovery is performed by the open twigs collector and the discovery app; the platform does the correlation, prioritisation, alerting, and remediation workflow.

`threatworx.io`

`threatwatch.io`

`github.com/threatworx/twigs`

`github.com/threatworx/discovery_app`