CTI as a Centralized Collaborative SOC

How OneFirewall Alliance turns distributed member detections into a single, continuously enforced intelligence layer — eliminating the siloed defence problem at scale.

Cyber Threat Intelligence | Collective Defence | Automated Enforcement | STIX 2.1 / TAXII 2.1 | WCF Agent | Crime | DeceptionGrid | IPS / SIEM / XDR

DOCUMENT VERSION 3.0 — February 2026

CLASSIFICATION Restricted / NDA Required

AUDIENCE SOC Teams · CISOs · Security Architects

5 Greenwich View Place · London E14 9NN · United Kingdom
Co. No. 11150273 · VAT GB305742714 · DUNS 223612138
onefirewall.com · support@onefirewall.com · +44(0)2038078020

Table of Contents

1. The Problem: 98% of Organisations Defend Alone

1.1 Fragmented Intelligence

Every SOC generates threat intelligence — blocked IPs, detected domains, malware hashes, phishing URLs. But in the conventional model, that intelligence stays inside the organisation that generated it. When a threat actor attacks Company A, Company B has no way of knowing until it is attacked too. Each organisation independently rediscovers the same attackers, the same infrastructure, the same campaigns.

The result is a massive duplication of investigative effort, compounded by the fact that adversaries reuse infrastructure. The same C2 server that attacks a bank in Italy on Monday will probe a government agency in the UK on Wednesday — and only the bank's SOC knows it.

1.2 The Adversary Advantage

Organised threat actors operate as networks. Criminal syndicates, nation-state APTs, and ransomware groups share tooling, leaked credentials, and target lists. They coordinate attacks, recycle infrastructure, and iterate on campaigns across many victims simultaneously.

> **Asymmetry:** Cyber criminals operate as an alliance. 98% of organisations defend themselves in isolation. OneFirewall exists to close this asymmetry by making collective defence the default for every member.

1.3 Traditional SOC Limitations

A conventional SOC — even a well-resourced one — is bounded by what it can see from its own perimeter:

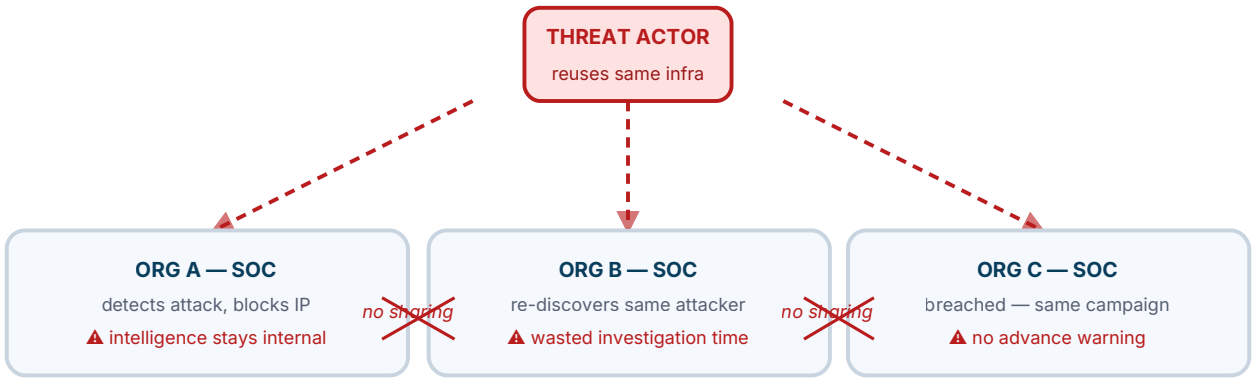

Fig 1. Traditional siloed defence: each SOC independently faces the same adversary with no intelligence sharing.

- Intelligence generated from detections expires inside the generating team — never shared externally.
- Each SOC must independently triage, investigate, and remediate the same attacker infrastructure.
- Time-to-detect is measured in days or weeks per organisation; time-to-share is never.

- Alert fatigue is amplified because analysts must rebuild context from scratch every time.
- Commercial threat feeds provide some coverage but lag production detections by hours or days.

2. OneFirewall as a Centralized Collaborative SOC

2.1 The Collective Defence Model

OneFirewall Alliance inverts the traditional model. Every member organisation contributes anonymised threat indicators it observes in production. Those indicators are validated, enriched, scored, and redistributed in real time to all other members. An attack detected by one organisation becomes a pre-emptive block for all others within the same distribution cycle — typically under 200ms.

This is the Centralized Collaborative SOC model: a single shared intelligence layer that acts as a continuous, crowd-sourced detection and enforcement function operating at the speed of the platform's distribution engine, not the speed of individual analyst review.

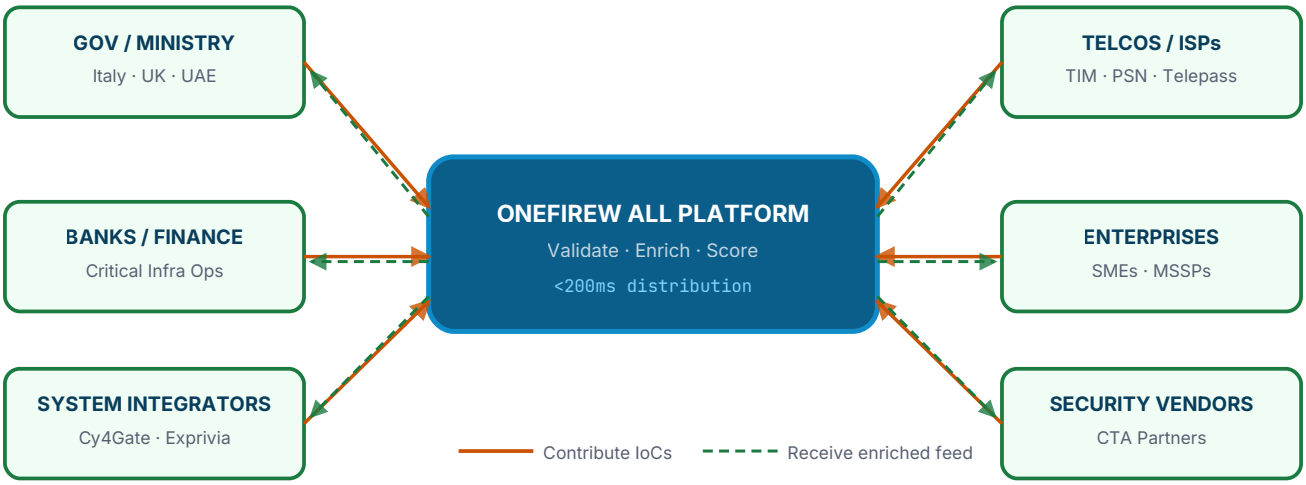

Fig 2. OneFirewall Collective Defence model. Every member contributes and receives intelligence simultaneously.

2.2 Siloed vs. Collective: Side-by-Side

Capability	Traditional Siloed SOC	OneFirewall Collaborative SOC
Intelligence source	Own perimeter only	180+ member organisations
Time to detect known attacker	Hours to days	<200ms if seen by any member
Threat coverage breadth	Limited to own traffic volume	32M+ IPs, 9M+ URLs, 26M+ hashes
Intelligence freshness	Batch imports (hours)	Continuous real-time delta sync
Analyst effort per attacker	Full investigation required	Pre-enriched: Crime Score + ATT&CK tags
Enforcement automation	Manual rule creation	Automatic via WCF Agent
Cost model	Per alert / per traffic volume	Per device — no traffic metering

Capability	Traditional Siloed SOC	OneFirewall Collaborative SOC
Privacy / attribution	–	**Fully anonymised contributions**

2.3 Alliance Membership Overview

180+	12	32M+	8M+	9M+	26M+
Alliance Members	Contributing Countries	IP Indicators	Domain Indicators	URL Indicators	Malware Signatures

Alliance member types include government ministries, national Telcos, critical infrastructure operators (energy, transport, financial market infrastructure), large enterprise security teams, system integrators, and MSSPs. The presence of high-value targets in the alliance means indicators are generated in the most attacked environments — giving all members access to intelligence that rarely reaches commercial feed circuits.

OneFirewall has been a member of the **Cyber Threat Alliance (CTA)** since 2020, contributing to and receiving vetted intelligence globally through the CTA network.

3. Four-Layer Architecture

3.1 Layer Overview Diagram

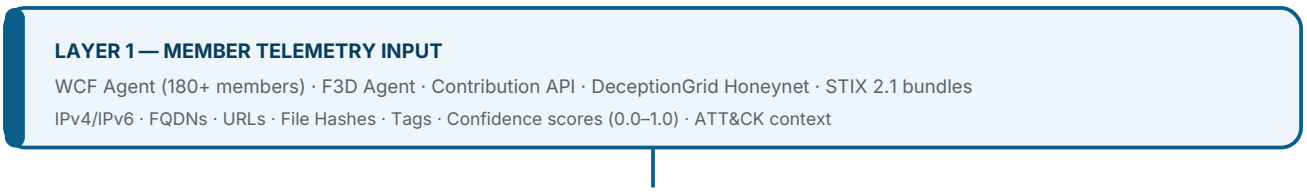

Fig 3. Four-layer OneFirewall architecture: telemetry in, enforcement out, sighting feedback loop.

3.2 Layer Detail

Layer	Components	Key SLA
L1 — Telemetry Input	WCF Agent (member perimeters), F3D Agent, Contribution API, DeceptionGrid nodes, STIX 2.1 ingest	60s batch flush; streaming available
L2 — Processing Engine	Deduplication, format normalisation, trust-weighted validation, GeoIP/ASN enrichment, Crime Score calculation, MITRE ATT&CK tagging, passive DNS, WHOIS correlation	Processing in <50ms per indicator
L3 — Distribution Engine	WebSocket push (WCF Agent delta), REST API (pull), TAXII 2.1 collection server, Syslog forwarder, SIEM connectors	<200ms end-to-end
L4 — Enforcement & Feedback	166+ supported security platforms; block event sighting telemetry closes the feedback loop into L2	Real-time; auto OTA rule update

4. Threat Intelligence Pipeline

4.1 End-to-End Flow Diagram

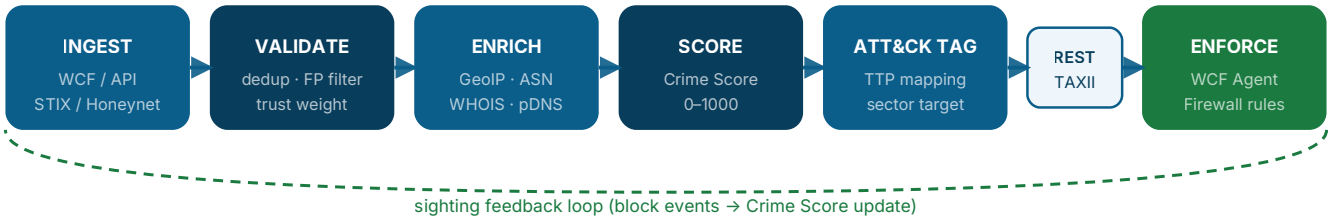

Fig 4. End-to-end intelligence pipeline with sighting feedback loop.

4.2 Indicator Types & Feed Scale

Indicator Type	Format	Feed Volume	Example
IPv4 / IPv6 Address	CIDR or host; RFC 4291	32,000,000+	`198.51.100.42`
Fully Qualified Domain	RFC 1123 FQDN	8,000,000+	`c2-panel.evil.example`
URL	RFC 3986	9,000,000+	`https://evil.example/dl/payload`
File Hash	MD5 · SHA-1 · SHA-256	26,000,000+ (signatures)	`d41d8cd98f00b204e9800...`

4.3 Crime Score: Mechanics

The **Crime Score** is a proprietary risk integer in the range **0–1000** assigned to every indicator. It is computed dynamically on each sighting or update event.

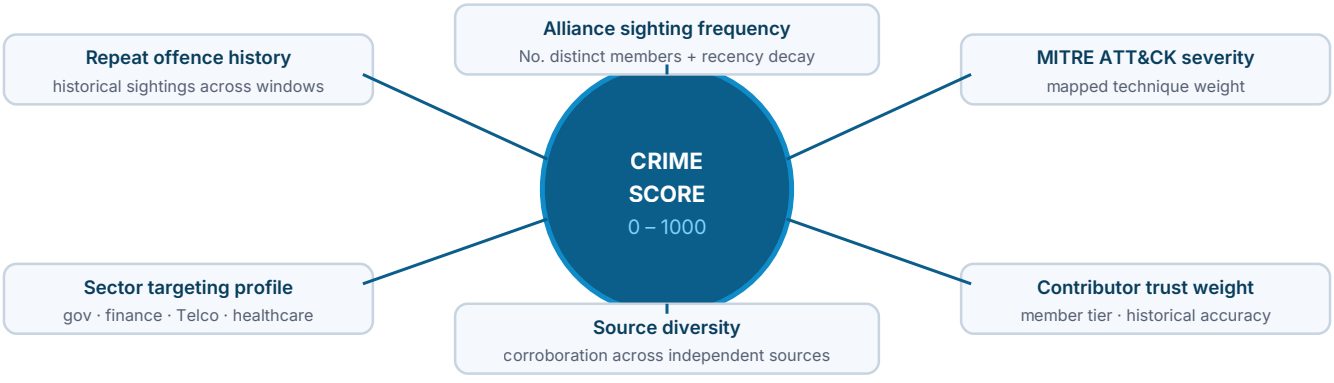

Fig 5. Crime Score composition — six weighted factors combined into a single 0–1000 risk value.

Score Band	Classification	Recommended Action
700–1000	High confidence	Automated block — production enforcement
400–699	Medium confidence	Alert + log; conditional block or investigation

Score Band	Classification	Recommended Action
190–399	Low confidence	Alliance baseline — log and monitor
0–189	Informational	Observation only; do not block

Default WCF Agent blocking threshold: 700. Alliance baseline threshold: ≥190. Both configurable per deployment.

5. The Contribution Model

5.1 Contribution Flow

Members contribute indicators automatically via the WCF Agent (passive observation mode) or manually via the Contribution API. Every accepted submission earns **OFA Coins** credited to the member account and immediately strengthens the alliance corpus.

Fig 6. Contribution flow: detect → submit → validate → distribute → all members auto-block.

CONTRIBUTION API PAYLOAD

```
POST /api/v2/ips  Authorization: Bearer <api-key>

[
  {
    "source":     "edge-fw-01",
    "ip":         "203.0.113.99",
    "confidence": 0.87,
    "notes":      "SSH brute-force observed",
    "tags":       "brute-force,ssh"
  }
]

// Response
{ "accepted": 1, "coins_credited": 3, "status": "queued" }
```

5.2 Privacy & GDPR Guarantees

Guarantee	Implementation
Anonymous contribution	Other members receive only updated Crime Scores — never the contributing member's identity or organisation name.
No personal data in indicators	IP addresses, FQDNs, URLs, and hashes are technical infrastructure identifiers, not personal data under UK GDPR.
On-premise data control	The WCF Agent and F3D Agent run on member infrastructure. The member controls what is submitted and can pause or delete at any time.
GDPR Article 6 basis	Legitimate interests (cybersecurity and network security). Processing is limited to the minimum necessary for threat defence.
Right to delete	Members may request deletion of their contributed indicators at any time via the API or support channel.

5.3 OFA Coins

Each accepted indicator submission credits **OFA Coins** to the submitting member's account. Coins are redeemable against platform subscription costs, additional API quota, or partner services. The Coins mechanism incentivises high-quality contributions and rewards members who maintain good historical accuracy (contributor trust weight).

---

6. DeceptionGrid Honeynet

DeceptionGrid is OneFirewall's globally distributed honeynet: geographically spread high-fidelity honeypot nodes designed to attract real attacker activity and convert it into actionable threat intelligence. Findings are fed directly into the Layer 2 processing engine, enriching the alliance corpus with TTPs and attacker fingerprints that would not appear in member perimeter telemetry.

Node architecture

Each DeceptionGrid node simulates multiple service categories simultaneously to maximise attacker engagement:

Category	Simulated Services / Ports
Network & Remote Access	SSH (22), Telnet (23), RDP (3389), OpenVPN/IPsec gateways
Web & API	HTTP/HTTPS (80/443), REST APIs, WebSocket endpoints
IoT & OT Protocols	Modbus (502), MQTT (1883), BACnet (47808), UPnP/SSDP, Zigbee (simulated)
Databases & File Services	MySQL, PostgreSQL, MongoDB, Redis, Cassandra, ElasticSearch (9200), FTP/SFTP, SMB/CIFS (445), NFS (2049)
DevOps & Cloud	Docker API (2375), Kubernetes API/Kubelet (10250), Jenkins (8080), GitLab (8929)
Auth, Email & Messaging	SMTP (25), IMAP (143), POP3 (110), LDAP/LDAPS (389/636), Kerberos (88), OAuth/OIDC endpoints

Intelligence extraction pipeline

1. **Lure & Engage**

   Nodes respond to scans and probing with realistic service banners and behaviours, drawing attackers into extended interaction.

2. **Record & Correlate**

   Full session recording, packet capture, and real-time logging. Activity is correlated at scale to identify patterns, tooling, and tradecraft.

3. **Extract Intelligence**

   Attacker behaviour is converted into IOCs (IPs, domains, hashes), TTPs, and attacker fingerprints — MITRE ATT&CK mapped.

4. **Feed into Alliance**

   Extracted intelligence enters the Layer 2 processing pipeline, contributing Crime Score updates and new indicators to the shared corpus.

> **Unique value:** DeceptionGrid captures attacker infrastructure before it targets real member environments. Indicators sourced from honeypots carry high sighting confidence because they represent confirmed malicious intent.

7. WCF Agent — Technical Reference

7.1 Deployment Architecture

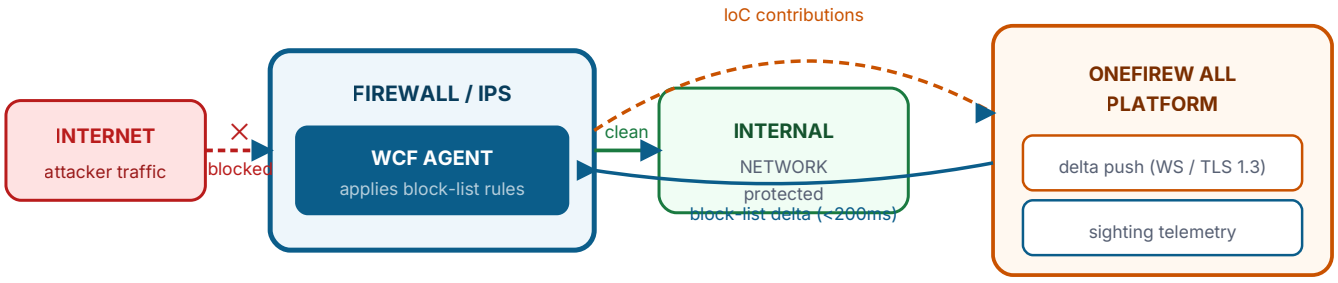

Fig 7. WCF Agent deployment: sits on or alongside the firewall, receives delta block-lists, contributes sightings.

7.2 Deployment Modes

Mode	Behaviour	Use Case
Inline (enforce)	Applies block rules to live firewall policy in real time. Full prevention mode.	Production perimeter
Monitor (passive)	Receives feed, logs matches, generates alerts. No blocking.	PoV testing · SIEM enrichment
Hybrid	Score ≥700 → auto-block; 190–699 → alert only. Threshold configurable.	Risk-controlled production
SIEM-only	No firewall integration. Feed delivered as STIX/TAXII stream to SIEM only.	SIEM detection rules only

7.3 System Requirements

Parameter	Minimum	Recommended
OS	Linux kernel 4.x+	Ubuntu 22.04 LTS or Debian 12
CPU	1 vCPU	2 vCPU
RAM	512 MB	1 GB
Disk	500 MB	2 GB (local log retention)
Network outbound	HTTPS port 443 to `api.onefirewall.com`	Dedicated management VLAN
Auth	API key (Bearer token)	API key + IP allowlist + mutual TLS (optional)
Updates	Auto OTA	Auto OTA with rollback

8. STIX 2.1 / TAXII 2.1 & API Reference

Base URL: `https://api.onefirewall.com/v2` · Auth: `Authorization: Bearer <api-key>`

Method	Endpoint	Description
`GET`	`/indicators`	Indicator feed. Params: `?since=` (ISO8601 delta), `?min_score=`, `?type=ip|fqdn|url|hash`, `?limit=` (max 10k)
`GET`	`/indicators/{id}`	Single indicator with full enrichment and STIX 2.1 object
`POST`	`/indicators`	Submit indicators (JSON array or STIX 2.1 bundle). Returns receipt with per-indicator acceptance status.
`GET`	`/feed/stix`	Full STIX 2.1 bundle export. Supports `?since=` for incremental.
`GET`	`/taxii/collections`	TAXII 2.1 collections discovery endpoint
`GET`	`/taxii/collections/{id}/objects`	TAXII 2.1 objects — returns STIX bundle from collection
`POST`	`/sightings`	Submit sighting events (block events). Feeds back into Crime Score.
`GET`	`/health`	Platform health: sync lag, queue depth, node status

STIX 2.1 indicator object (abridged)

```
{
  "type": "indicator",           "spec_version": "2.1",
  "id": "indicator--<uuid>",    "pattern_type": "stix",
  "pattern": "[ipv4-addr:value = '198.51.100.42']",
  "valid_from": "2026-02-25T14:00:00Z",
  "confidence": 85,
  "labels": ["botnet","c2","ransomware-infrastructure"],
  "extensions": {
    "x-onefirewall": {
      "crime_score": 823,
      "sector_targets": ["finance","government"],
      "mitre_techniques": ["T1071.001","T1102"]
    }
  }
}
```

Quick-start: delta pull and submit

```
# Pull indicators updated in the last hour with Crime Score ≥700
curl -s "https://api.onefirewall.com/v2/indicators?since=2026-02-25T14:00:00Z&min_score=700&limit=5000" \
     -H "Authorization: Bearer YOUR_API_KEY"

# Submit a new IoC
curl -s -X POST "https://api.onefirewall.com/v2/indicators" \
     -H "Authorization: Bearer YOUR_API_KEY" \
     -H "Content-Type: application/json" \
     -d '[{"type":"ip","value":"198.51.100.42","confidence":0.9,"tags":"c2,ransomware"}]'

# WCF Agent install (Linux)
```

```
curl -sSL https://install.onefirewall.com/wcf-agent | sudo bash
sudo wcf-agent configure --api-key YOUR_API_KEY --mode inline --min-score 700
sudo systemctl enable --now wcf-agent
```

9. Integration Matrix

The WCF Agent ships native modules for 166+ platforms. Representative selection below; full list at docs.onefirewall.com.

Vendor	Platform	Method	Layer
Check Point	NGFW R80.x/R81.x · SecureXL	WCF Agent (SecureXL API)	L3/L4
Fortinet	FortiGate FortiOS 6.x/7.x	WCF Agent (FortiGate REST API)	L3/L4/L7
Forcepoint	NGFW · Web Security · SMC	WCF Agent (SMC API)	L3/L4
Cisco	ASA · FTD · Firepower · Meraki MX	WCF Agent (FMC REST / Dashboard API)	L3/L4
Palo Alto	PAN-OS · Panorama · XSOAR	WCF Agent (XML API) · TAXII 2.1	L3/L4/L7
Juniper	SRX Series · Junos	WCF Agent (Junos RPC)	L3/L4
Sophos	XG · XGS · Firewall OS 18+	WCF Agent (Sophos API)	L3/L4/L7
Netgate / pfSense	pfSense CE/Plus · OPNsense	WCF Agent + OneDevice module	L3/L4
SonicWall	NSA · TZ Series · SonicOS 7	WCF Agent (SonicOS API)	L3/L4
Microsoft	Azure Sentinel · Defender XDR	TAXII 2.1 / STIX 2.1	SIEM
Splunk	Splunk SIEM · SOAR	TAXII 2.1 / STIX 2.1	SIEM
IBM	QRadar SIEM	TAXII 2.1 / STIX 2.1	SIEM
Elastic	Elastic SIEM / Security	TAXII 2.1 / STIX 2.1	SIEM
CrowdStrike	Falcon XDR	STIX / Threat Intel API	XDR
SentinelOne	Singularity XDR	STIX / Threat Intel API	XDR
AWS	WAF · CloudFront · Network Firewall	WCF Agent (AWS APIs)	L7 / Cloud
Google Cloud	GCP Cloud Armor · Network Firewall	WCF Agent (GCP APIs)	L7 / Cloud
Cloudflare	WAF · Gateway · Access	WCF Agent (Cloudflare API)	L7 / CDN
HAProxy / ModSecurity	Open-source reverse proxy / WAF	WCF Agent (config injection)	L7
Windows	Windows Defender / Firewall	WCF Agent (Windows API)	Endpoint

10. Product Ecosystem

All products are powered by the same Alliance CTI layer. They represent enforcement surfaces across different network points.

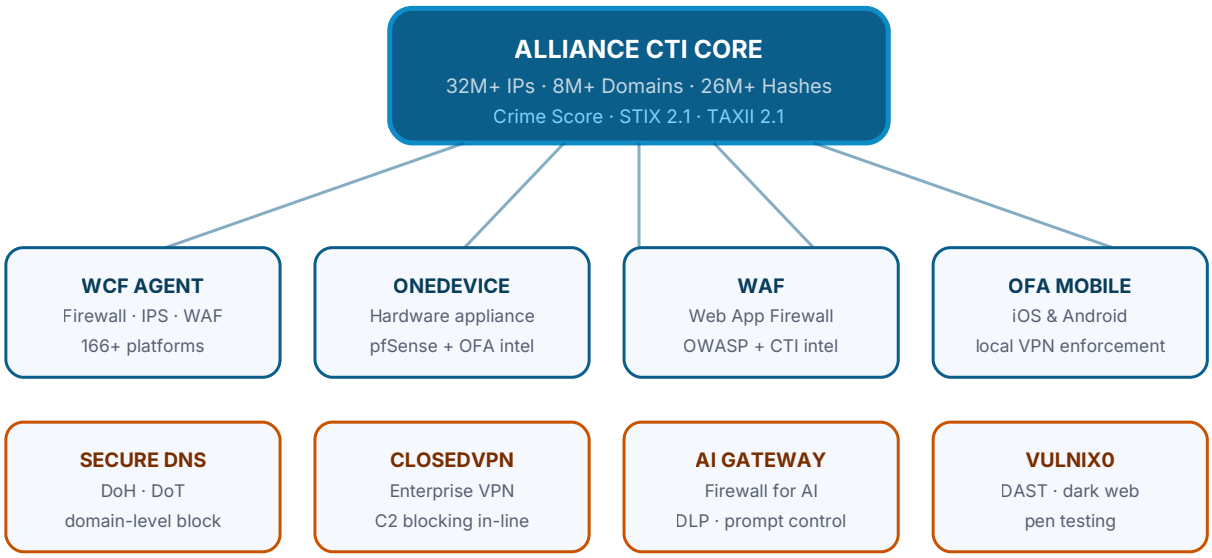

Fig 8. OneFirewall product ecosystem — all enforcement surfaces connected to the single Alliance CTI core.

11. Security & Compliance

Domain	Detail
UK GDPR / DPA 2018	Compliant. The platform processes only technical threat indicators (IP, FQDN, URL, hash) — no personal data in the CTI corpus. Member contact data handled separately under Article 6(1)(b/f).
Cyber Essentials	Certified (current). Covers boundary firewalls, secure configuration, access control, malware protection, and patch management.
ISO/IEC 27001	Processes aligned with ISO 27001 information security management practices. OneDevice appliance is compliance-ready.
STIX / TAXII	STIX 2.1 indicator objects; TAXII 2.1 server. Native integration with all STIX/TAXII-capable SIEM/SOAR platforms.
Transport security	TLS 1.3 for all API and WCF Agent channels. Certificates from public CA. Mutual TLS available for WCF Agent channels.
Authentication	API: Bearer token (API key). Portal: username/password + MFA. WCF Agent: API key (+ optional mTLS).
Contribution anonymity	The submitting member is never disclosed. Other members receive only updated Crime Scores.
Data residency	Platform infrastructure UK-hosted. Cross-border transfers use contractual safeguards per UK GDPR Chapter V.

Domain	Detail
Vulnerability programme	Responsible disclosure via security@onefirewall.com. Status page: status.onefirewall.com.

> **Cyber Essentials Certified:** Certificate details and scope available upon request. Contact support@onefirewall.com with subject line "Cyber Essentials Certificate Request".

OneFirewall Alliance LTD · 5 Greenwich View Place, London E14 9NN, UK · Co. No. 11150273 · VAT GB305742714 · DUNS 223612138 · onefirewall.com · support@onefirewall.com · +44(0)2038078020

Classification: Restricted — Authorised Recipients Only | **Document:** CTI as a Centralized Collaborative SOC | **Version:** 3.0 | **Date:** February 2026