

SentinelSuite

TECHNICAL & COMMERCIAL DATASHEET

Infrastructure security you run yourself — by Dognet Technologies

Vulnerabilities, assets, firewall, hardening and threat intelligence in **one integrated suite** that you deploy on your own hardware. The tools work alone or together: they exchange events, query each other and react according to rules that you decide.

SentinelCore	FireDog	CyberShepherd	IntelliDog	dog-agent
VULNERABILITIES AND RISK	FIREWALL AND TRAFFIC	HARDENING AND CORRELATION	THREAT INTELLIGENCE	ENDPOINT AGENT

Self-hosted · on-premise · **No SaaS** · no phone-home · **Air-gap** supported · **GPG + SHA256** signed artifacts · **MCP** server on every tool · **NIS2** · ISO/IEC 27001 · PCI DSS · NIST

Dognet Technologies S.r.l. · Via XXV Aprile 47, 24055 Cologno al Serio (BG), Italy
info@dognet.tech · +39 352 068 1747

October 2026 edition
sentinelsuite.dognet-technologies.online

The suite at a glance

Four tools and one lightweight agent: visibility, protection, detection, intelligence

Independent tools, one perimeter.

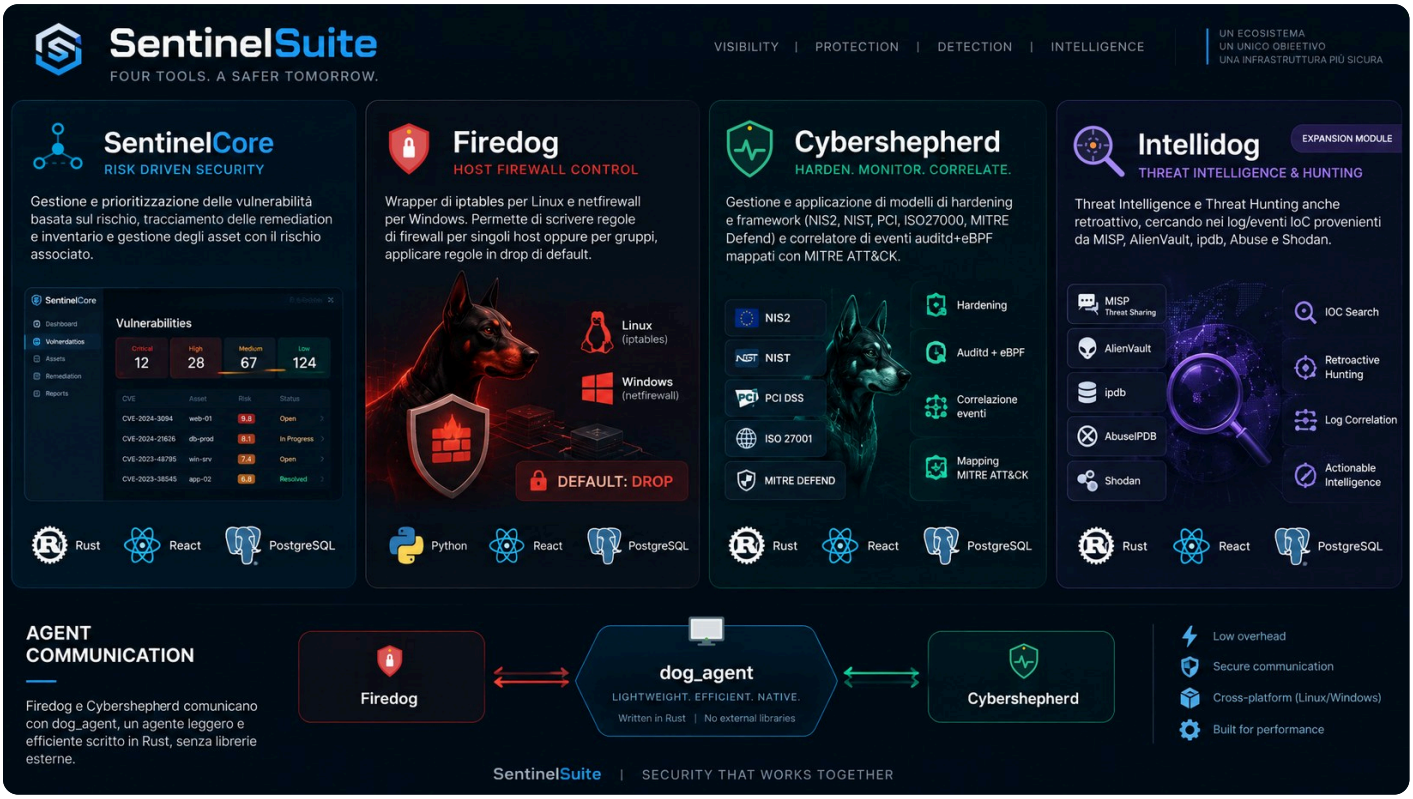
SentinelSuite is a cybersecurity toolchain for security teams and system integrators, shipped as signed appliances and packages that you install on your own hardware. No SaaS, no phone-home: data stays on-premise, including in regulated and air-gapped environments.

4 + 1
modular tools and a single agent

MCP
every tool can be driven by an AI agent

0
data sent outside

4
frameworks covered: NIS2, ISO, PCI, NIST



SentinelSuite architecture: each tool works on its own and integrates with the others; dog-agent connects the endpoints.

Component	What it covers	Version	Edition	Stack
SentinelCore	Asset inventory, vulnerabilities prioritised on real risk, SLAs, remediation plans, reports	v1.2.0	Community (open source) + PRO	Rust · React · PostgreSQL
FireDog	Centralised firewall for host fleets: rules, default-DROP policy, blocked traffic, threat score	v1.0.1	Community (open source)	Django · React · PostgreSQL
CyberShepherd	Hardening by models and frameworks, continuous compliance, event correlation on MITRE ATT&CK	—	PRO	Rust · React · PostgreSQL
IntelliDog	Threat intelligence and retroactive threat hunting across the events of the whole suite	—	Expansion module · PRO	CyberShepherd module
dog-agent	Single endpoint agent: telemetry, heartbeat and command execution for all the tools	v1.1.1	Included	Rust · static musl binary

Included for Dognet customers: access to SentinelSuite licences is part of the Essential tier of the maintenance contract. Each tool can be adopted on its own: start with the one you need and add the others when your process is ready.

SentinelCore v1.2.0

Vulnerability management driven by real-world risk

AVAILABLE | COMMUNITY + PRO

WHAT IT IS

All your vulnerabilities, one workflow.

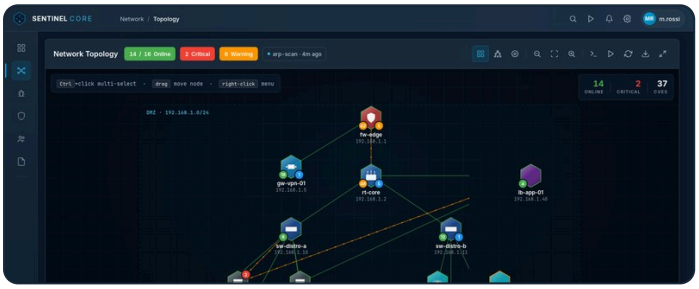
SentinelCore is the suite's vulnerability management platform, built for in-house security teams. It gathers in one place the vulnerabilities of your infrastructure, whether found by the built-in network scan or imported from third-party scanners. It gives each one a **real risk score**, not just the headline CVSS, routes it to the right people, tracks resolution against deadlines (SLAs) and produces the reports you need for audits and compliance.

A single view
Assets, services and vulnerabilities in one console.

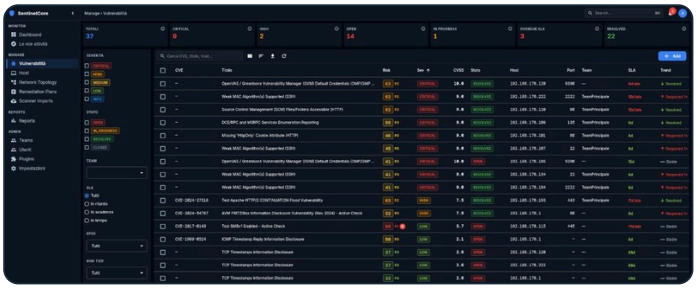
Real risk
Priority from context, not from CVSS alone.

Effective teams
An owner and a deadline for every item.

Reports ready
Evidence for audits and management.



Network Topology: discovered hosts with their associated risk (real screenshot)



Vulnerability management with CVSS, EPSS, status and SLA (real screenshot)

THE OPERATING WORKFLOW

From discovery to compliance, as a continuous process

1 Discover
Built-in network scan (ARP-scan, nmap): hosts, ports, services.

2 Import
Ingest results from 10 third-party scanners or the OpenVAS connector.

3 Score
Compute risk from CVSS, EPSS, KEV, exposure and business impact.

4 Assign
Route to teams and people with priority and SLA deadline.

5 Remediate
Remediation plans, comments, progress status.

6 Verify
Confirm closure with new scans or updated data.

7 Report
Reports for audits and management, in PDF, CSV and JSON.

What SentinelCore does

- **Rebuilds the network** as a map: every host appears in the topology with ports, services and risk.
- **Centralises vulnerabilities** from different scanners and ties them to the same asset.
- **Reprioritises dynamically**: when a CVE enters the KEV catalogue or its EPSS changes, its priority changes.
- **Assigns and enforces deadlines**, with signals at 75%, 90% and 100% of the SLA window.
- **Notifies on your team's channels**: email, Slack, Telegram, Teams, webhook, PagerDuty, OpsGenie.
- **Syncs with JIRA** in both directions (status, priority, custom fields).
- **Records everything** in an audit log usable for PCI DSS, ISO 27001, SOC 2 and HIPAA.
- **Can be driven by an AI agent** through a read/write MCP server.

What SentinelCore does not do

- **It is not a vulnerability scanner** in the strict sense: beyond network and port discovery, vulnerabilities come from external scanners or the OpenVAS connector.
- **It is not multi-tenant:** one installation serves one organisation.
- **It has no built-in high availability or clustering:** it runs on a single host with a dedicated PostgreSQL.
- **It neither requires nor installs agents** on endpoints: visibility comes from network scans and imported data.
- **It does not patch automatically:** remediation plans track the work, they do not carry it out for the team.
- **It has no built-in LDAP/SSO:** authentication is local, with optional 2FA.

SentinelCore technical sheet

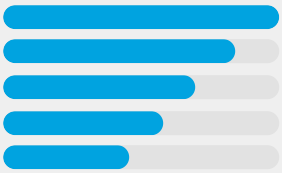
Risk engine, SLAs, integrations, architecture and distribution

V1.2.0

RISK ENGINE

A composite score, not a fixed number

The risk of each vulnerability comes from five factors and is updated whenever one of them changes.



Illustrative bars: the factors contribute to the score, they do not represent the actual weights. An **override** channel pushes zero-days, ransomware-related CVEs and CISA KEV entries to the top; EPSS is refreshed daily, NVD enriches CVE metadata.

SLA

Severity becomes a deadline

1 d	7 d	30 d	90 d
CRITICAL	HIGH	MEDIUM	LOW

Each tier has its own resolution window. The system flags the approaching deadline at **75%** and **90%** and the breach at **100%**; throttling and quiet hours keep the notification channel clean.

ROLES AND ACCESS

Three roles (**admin**, **team leader**, **user**), TOTP 2FA, remotely revocable sessions, JWT authentication in an `httpOnly` cookie with a CSRF token on write operations.

IMPORTERS · 10 SCANNERS

Qualys · Nessus · Burp Suite · OpenVAS / GVM · Nexpose / InsightVM · OWASP ZAP · Nmap · Nikto · Grype · Trivy

NOTIFICATIONS · 7 CHANNELS

Email · Slack · Telegram · Microsoft Teams · Webhook · PagerDuty · OpsGenie

INTEGRATIONS AND FEEDS

Bidirectional JIRA · NVD · Daily EPSS · CISA KEV · Offline CVE feed · Plugins

ARCHITECTURE



Background workers: assignment, automatic rescan, critical alerts, EPSS and KEV updates, NVD enrichment, JIRA sync, OpenVAS polling, SLA checks, plan progress, report generation, weekly report, notification digest, log retention.

MCP SERVER

Single endpoint `POST /api/mcp` · JSON-RPC 2.0

Read: `list_vulnerabilities` `get_vulnerability` `get_risk_summary`

Write: `assign_vulnerability` `update_vulnerability_status` `accept_risk` `trigger_scan`

Every write tool is bound to the scope of the API key that calls it. No external gateway.

MINIMUM REQUIREMENTS

System	Debian 12 / 13 · x86_64
CPU	4 vCPU · x86-64-v2
RAM	4 GB · 8 GB above 500 hosts
Disk	20 GB SSD
Network	1 NIC · L2 access to targets
Console	HTTPS :443

RELEASE FORMATS

.tar.gz	Debian installer · 16 MB
.ova	VirtualBox / VMware · 678 MB
.qcow2	KVM / Proxmox · 702 MB
Source	GitHub · AGPL-3.0 licence

Every artifact is GPG-signed and published with a detached SHA256 checksum.

INSTALLATION AND UPGRADE

`install.sh` sets up the database, nginx with HTTPS, nmap, arp-scan and the systemd service, and creates the first administrator with no prompts.

`upgrade.sh` upgrades in place: database backup, migrations, health check and automatic restore if a migration fails.

Documentation: a complete user manual, in Italian (installation, roles and permissions, network and assets, vulnerabilities and risk, reports and remediation), is available online and as a PDF on the suite website.

FireDog v1.0.1

Centralised firewall management for fleets of Linux hosts

AVAILABLE · COMMUNITY · OPEN SOURCE

WHAT IT IS

One panel, every firewall.

FireDog is the centralised firewall management platform for fleets of Linux hosts. Instead of configuring **iptables** by hand on every server, a single web panel, the **master**, lets you define rules and distribute them to all managed hosts, watch their blocked traffic and detected threats, and react from one place.

THE COMPONENTS

Master

a dedicated host or VM

Django backend and React frontend: policy editor, real-time traffic dashboard, REST API and MCP server. It receives heartbeats, statistics and threats from the agents over WebSocket.

Firewall tools

on every target, in `/opt/sentinelsuite/firedog`

`firewall-manager` (rules CLI), `traffic-analyzer` (PCAP analysis with threat scoring) and `firewall-init.sh` (DROP policy and anti-attack protections).

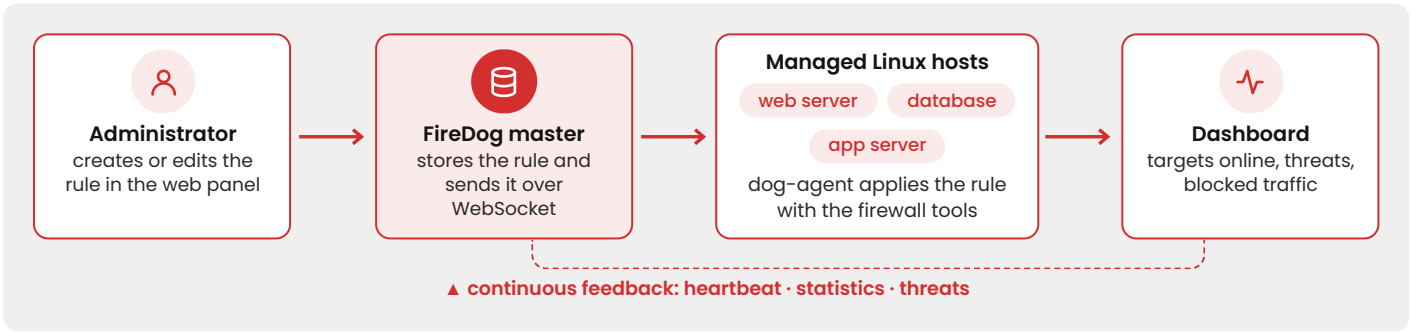
dog-agent

on every target, `/usr/bin/dog-agent`

Static Rust binary shared with the rest of the suite: it authenticates the target to the master with a two-phase pairing, sends telemetry and executes the rule commands it receives.

OPERATING WORKFLOW

From the panel to protection, in real time



1 Definition
The administrator creates or edits rules in the panel, per host or per network interface.

2 Distribution
The master stores the rule and sends it to the agents. If an agent is offline, the rule waits for reconciliation.

3 Execution
dog-agent receives the command, applies the rule through the firewall tools and watches the traffic.

4 Feedback
Blocked traffic is saved as PCAP, analysed and shown in real time with its threat score.

What FireDog does

- **Manages rules from one place** and distributes them in real time; no SSH access to the targets is involved.
- **Applies a default-DROP policy** on INPUT and OUTPUT: only what has been authorised gets through.
- **Protects against SYN flood, port scans and SSH brute force**, with configurable thresholds and ban duration.
- **Supports hosts with several network interfaces**: per-NIC rules and dedicated rx/tx counters.
- **Captures and analyses blocked traffic** in PCAP format and assigns a threat score.
- **Records every write** in an audit log with user, action and values.
- **Exposes an MCP server** for authorised AI agents, read and write.

What FireDog does not do

- **It does not manage Windows hosts yet:** the current version is for Linux systems; support arrives with the next release.
- **It does not replace network equipment:** it is not a router and does not handle NAT, VPN or load balancing.
- **It is not a full IDS/IPS:** it relies on iptables and on the analysis of blocked traffic.
- **It does not analyse allowed traffic:** it works on the PCAP of blocked traffic only.
- **It does not configure application services** such as web servers or databases.
- **It does not decide rules on its own:** actions are driven by the administrator or by authorised AI agents over MCP.

With the next release: `firewall-manager` ported to PowerShell and dog-agent built for Windows 11 and Windows Server; support for Oracle Linux and Arch Linux; integrations between the tools of the suite.

FireDog technical sheet

Protections, traffic capture, MCP server, requirements and installation

V1.0.1

BUILT-IN PROTECTIONS

Default deny, then the exceptions

- **DROP** policy on INPUT and OUTPUT
- **SYN flood:** limit of 50 connections per second (burst 100), configurable
- **Port scan:** detection and blocking
- **SSH brute force:** configurable threshold and window (20 attempts in 60 seconds by default)
- Temporary or permanent **ban** through a persistent ipset, managed with `--list-bans` and `--unban`

TRAFFIC AND THREATS

See what gets rejected

- Separate INPUT and OUTPUT logs in **PCAP** via ulogd2
- Automatic retention with logrotate
- `traffic-analyzer` gives every source a **threat score** with the reasons behind it
- Statistics and threats streamed to the master dashboard
- List of the public remote peers observed for each target

TARGET CONFIGURATION

`/etc/firewall/firedog.conf`

- **Monitored interfaces** reported to the master
- **Always-open ports** ahead of the DROP policy
- **SSH protection:** attempts, window, ban duration
- Created at first install and never overwritten by upgrades
- Per-NIC rules with `-i` / `-o`

MCP SERVER · POST /api/mcp · JSON-RPC 2.0

16 tools to query and govern the firewall from an AI agent

Targets

`list_targets`
`get_target`
`list_interfaces`
`get_policy_summary`

Rules

`list_rules`
`get_rule`
`create_rule` ADMIN
`delete_rule` ADMIN

Threats and blocks

`list_threats` · `get_threat`
`resolve_threat` ADMIN
`list_blocked_ips`
`block_ip` · `unblock_ip`
ADMIN

Traffic

`list_traffic_stats`
`list_network_flows`

Volume snapshots and observed remote peers.

The MCP key is generated in **Settings › API Keys MCP** and inherits the role and permissions of the user who created it: tools marked Admin only work with an administrator's key. Every write over MCP lands in the audit log with the `MCP <tool>` prefix, so it can be told apart from changes made in the interface.

REQUIREMENTS · MASTER

System	Debian 12+ · Ubuntu 22.04+
Runtime	Python 3.11–3.13 · Node 20
Services	PostgreSQL · Redis · nginx
RAM	2 GB · 4 GB recommended
Ports	TCP 80 / 443 via nginx

REQUIREMENTS · TARGET

System	Debian / Ubuntu · openSUSE / SLES
Firewall	iptables + ulogd2
Path	`/opt/sentinelsuite/firedog`
Agent	dog-agent · .deb / .rpm
Startup	systemd

DISTRIBUTION

Source	GitHub · `stabile` branch
Licence	AGPL-3.0
Master	`INSTALL.md` guide · systemd units
Target	`get-firedog.sh` script
Manual	in Italian, online

target — installing the firewall tools

```
# download the installer, inspect it, then run it
curl -fsSL https://raw.githubusercontent.com/Dognet-Technologies/
  firedog/stabile/firedog-package/get-firedog.sh -o get-firedog.sh
less get-firedog.sh
sudo bash get-firedog.sh # installs and activates the DROP policy
sudo bash get-firedog.sh --skip-init # tools only, activate later
```

TWO-PHASE PAIRING

1. **API key.** The agent presents the key generated in Settings › API Keys Agent.
2. **Identity.** The master checks the fingerprint `SHA512(ip + hostname + mac)` against the target record.
3. **Channel up.** Heartbeats, statistics, threats and rule commands flow from here.

DROP policy without locking yourself out. In the configuration file you declare the ports that must stay open at startup, so activating the DROP policy does not cut off access to the host, and the interfaces to monitor on hosts with several network cards. With `--skip-init` you install the tools and activate the firewall later.

CyberShepherd

Hardening, continuous compliance and host event correlation

PRO EDITION

WHAT IT IS

Harden. Monitor. Correlate.

CyberShepherd secures Linux hosts and then checks that they stay secure. It applies hardening models aligned with the reference frameworks, measures compliance over time and correlates system events, mapping them onto the **MITRE ATT&CK** matrix.

In the PRO edition it is also the governing point of the suite: this is where you define the playbooks and the policy for the automated responses that SentinelCore and FireDog apply.

REFERENCE FRAMEWORKS

NIS2 · NIST · ISO/IEC 27001 · PCI DSS · CIS · MITRE ATT&CK · MITRE D3FEND

Every hardening control carries a reference to the frameworks it satisfies; every correlated event carries the ATT&CK tactic and the D3FEND defensive technique that mitigates it.

THE OPERATING WORKFLOW

From the model to proof of compliance

1. **Profile** — Pick the model that fits the host's role and the framework to meet.
2. **Harden** — Apply services, auditd rules and sysctl parameters through dog-agent.
3. **Collect** — Receive audit events, metrics and control status from the hosts.
4. **Correlate** — Turn raw events into security events mapped onto ATT&CK.
5. **Alert** — Notify violations and correlations by email, Slack and Discord.
6. **Verify** — Repeat the compliance scans and keep the history.

Model-based hardening

- Declarative templates applied by the agent
- System hardening, IAM, encryption
- Network security, logging and monitoring
- CIS, ISO 27001, NIST, PCI audit baselines

Continuous compliance

- Scheduled and on-demand scans
- More than 16 behavioural control policies
- Automatic violation tracking
- History on time series (InfluxDB)

Event correlation

- Frequency and sequence over time
- Real user identity and session
- Process lineage, files, network
- Criticality of the asset involved

WHAT CORRELATION RECOGNISES · MITRE ATT&CK TACTICS

Credential access — brute force, access to credential files

Privilege escalation — privilege jumps attributed to the real user

Execution — suspicious executions, reverse shells

Persistence — changes to cron, systemd, SSH keys

Lateral movement — access chains between hosts

Command & control — beaconing to the same destination

Exfiltration — anomalous outbound transfers

Defense impairment — sensor tampering or silence

Primary source: **auditd**; **eBPF** sensor as an optional plugin, off by default. Tactic schema: MITRE ATT&CK v19.

What CyberShepherd does

- **Applies hardening** repeatably across services, audit and kernel.
- **Measures compliance** with the frameworks and keeps its history.
- **Correlates host events** and places them on the ATT&CK matrix.
- **Extends through plugins**, including those for third-party IDS/IPS.
- **Exposes an MCP server** for alerts, correlations, compliance and hardening.
- **Governs the integrations** of the suite: playbooks, policy, approvals, audit.

What CyberShepherd does not do

- **It does not replace IDS/IPS,** perimeter firewalls or routers: it integrates them through plugins.
- **It is not a general-purpose SIEM:** it focuses on the hosts managed by the suite.
- **It takes no intrusive action on its own:** they go through the policy and the required approval.
- **It does not clone hosts:** the customer creates and declares test and pre-production environments.

IntelliDog

Threat intelligence and threat hunting, including retroactive

dognet TECHNOLOGIES

EXPANSION MODULE · PRO

WHAT IT IS

What the world knows, searched in your own house.

IntelliDog is CyberShepherd's threat intelligence module. It collects indicators of compromise from recognised sources and looks for them in the data the suite already holds: FireDog traffic, SentinelCore vulnerabilities, CyberShepherd host events.

The search also works backwards: when an indicator becomes known today, IntelliDog checks whether it already passed through your systems in the previous weeks.

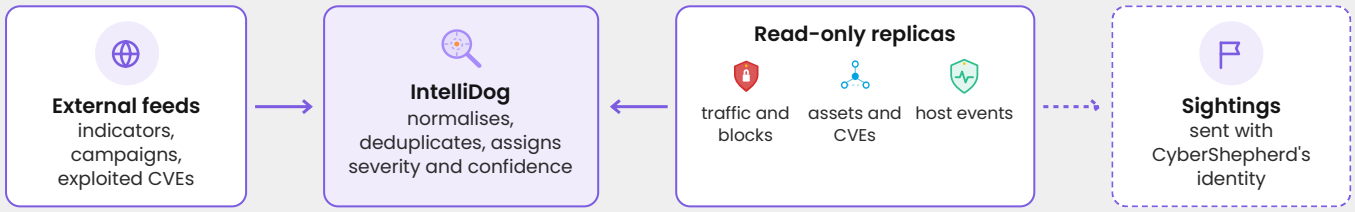
INTELLIGENCE SOURCES

- **MISP** threat and campaign sharing
- **AlienVault OTX** pulses and indicators
- **AbuseIPDB** IP reputation
- **Shodan** Internet exposure

Plus the CISA KEV catalogue and custom feeds in CSV and JSON format.

THE DATA FLOW

From indicators to sightings, without touching the tools in production



IntelliDog reads PostgreSQL copies of the tools' data: even heavy searches do not slow down the systems that are doing the work.

Hunting through history

Once a week it searches for indicators across the history of all the tools. When a strong signal arrives, such as a CVE added to the KEV catalogue or an active campaign reported in MISP, a targeted search on that single item starts right away.

Six ways to correlate

- **IP addresses** in firewall traffic
- **Domains** contacted (with an integrated IDS)
- **CVEs** present on assets
- **Hashes** of files known to be malicious
- **Patterns** in logs
- Anomalous **behaviour**

Weighted indicators

Every indicator has a type, severity, confidence level, age and TLP marking. Weight matters: many IPs belong to shared infrastructure such as CDNs and hosting, and blocking them blindly would break legitimate services.

WHAT YOU GET IN THE CONSOLE

- **Overview** feed status and recent detections
- **Feeds** sources, updates, outcomes
- **Indicators** search and filters on IoCs
- **Detections** acknowledge and close
- **Threat hunting** custom searches through history

What IntelliDog does

- **Ingests indicators** continuously from the configured intelligence sources.
- **Searches for them in your data,** including data collected before the indicator was known.
- **Classifies the hosts involved:** vulnerable, in contact with the attacker, or both.
- **Changes the weight of vulnerabilities:** a medium CVE exploited in an active campaign becomes urgent.
- **Proposes countermeasures** to the other tools, such as block rules for FireDog.
- **Uses open standards:** STIX and MISP for intelligence, TLP for sharing.

What IntelliDog does not do

- **It does not work on its own:** it is a CyberShepherd module and requires the PRO edition.
- **It does not act directly on systems:** it reports. Responses go through the policy defined in CyberShepherd.
- **It has no connection settings of its own:** it talks to the other tools through CyberShepherd.
- **It does not write to the tools' databases:** it works on read-only replicas.
- **It does not replace the analyst:** a contact without a vulnerability is still a case to investigate.

dog-agent v1.1.1 · MCP

One agent on the endpoints, one language to govern the suite

AVAILABLE · INCLUDED

DOG-AGENT · WHAT IT IS

Small, quiet, honest.

A single static Rust binary for the whole suite. You install it once on the host and it reports to the products in use, one or several at the same time, over an authenticated WebSocket channel. Each product needs just one `[[targets]]` block in `/etc/dog-agent/agent.conf`, with its own API key.

<12 MB resident memory	**1.8** MB .deb package
0 runtime dependencies	**1** agent for all the tools

Inventory

Installed packages and versions on every run, to match against CVE feeds and policy.

Ports and services

Listening sockets and owning processes: the host's own view, to cross-check with scans and rules.

Configuration drift

Hashes of watched files and flags on changes between one check-in and the next.

Commands and telemetry

Heartbeat, statistics and threats to the master; execution of the rules it receives.

REQUIREMENTS

System	Linux · x86_64 / arm64
Kernel	4.19+ · no libc dependency
Memory	< 12 MB · 32 MB peak when scanning
Disk	40 MB + local buffer
Privileges	unprivileged where possible

PACKAGES

.deb	Debian / Ubuntu · 1.8 MB
.rpm	openSUSE / SLES / RHEL · 1.9 MB
Build	static musl · TLS with rustls
Service	systemd unit included
Source	GitHub · SHA256 checksum
Coming	Windows 11 · Windows Server

ENROLMENT

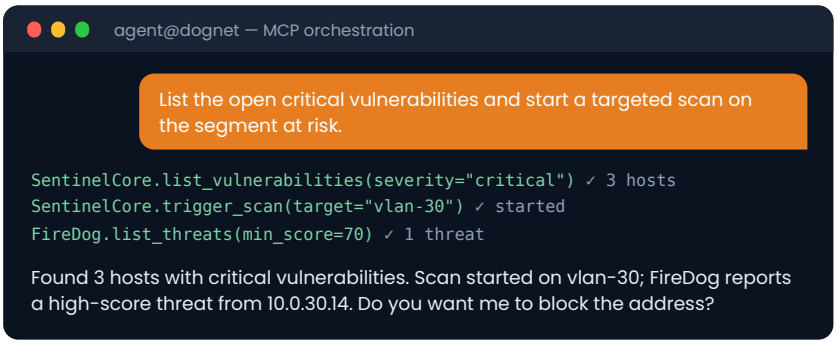
1. **Install** the package: the service stays idle until configured.
2. **Configure** one `[[targets]]` block per product.
3. **Start:** the two-phase pairing runs by itself on first connection.

MANAGEMENT BY AI AGENT · MODEL CONTEXT PROTOCOL

Not just dashboards: orchestration.

Every product in the suite exposes its own **MCP** server: the whole infrastructure can be driven in natural language by an AI agent, such as Claude, running inside the customer's perimeter. You can use one agent per tool or a single one for the whole suite.

- **Same contract on every tool:** endpoint `POST /api/mcp`, JSON-RPC 2.0
- **Dedicated keys**, separate from endpoint-agent keys and tool-to-tool keys
- **Permissions inherited** from the key: read for all, write only where authorised
- **Traceability:** every write over MCP is recorded in the audit log
- **No data leaves:** consistent with NIS2 and air-gapped environments



Illustrative example of orchestration by agent

MCP server	Read	Write
SentinelCore	vulnerabilities, detail, risk summary	assignment, status change, risk acceptance, scan start
FireDog	targets, interfaces, rules, threats, blocked IPs, traffic, network flows, policy summary	rule creation and deletion, IP block and unblock, threat resolution
CyberShepherd	targets, metrics, alerts, correlations, compliance scans, hardening templates	alert acknowledgement, hardening execution

Integration between the tools

Pairing, event exchange and reception rules: how the tools recognise and talk to each other

Alone they see a piece. Together they see the case.

Integrating means connecting the tools to each other at configuration time, exchanging the keys they use to recognise one another. Once connected they send events, ask each other for information and react according to rules, the **playbooks**, enabled one by one. Without this configuration every tool works on its own.

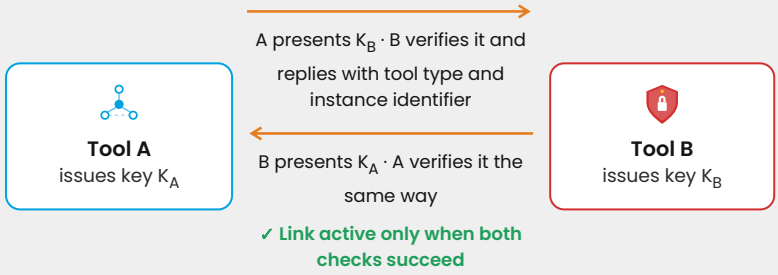
Events
"I detected this." An observed fact, carrying only the sender's own data.

Questions
"Is this IP in your inventory?" Each tool answers about its own data.

Playbooks
Event, questions, conditions, response. Enabled individually.

1 · PAIRING

Two tools are connected only when each has verified the other's key



✓ Link active only when both checks succeed

The administrator enters the other tool's URL and key in each tool, under Settings › Connected tools. One direction alone does not complete the handshake.

LINK STATES

Not configured	URL or key missing
Pending	only one direction verified
Active	events enabled
Error	key rejected, peer unreachable or certificate changed
Revoked	key deleted on either side

Issued keys: stored only as hashes.

Received keys: encrypted at rest, outside the database.

One key, one peer: not valid if another tool presents it.

HTTPS required. Keys grant no permissions.

2 · EVENT CONTRACT

One envelope for everyone, and reception rules that trust nobody

event — the common envelope

```
schema_version  version; unknown = rejected
event_id  UUID; duplicates are dropped
event_type  e.g. threat.detected
source_tool  the sending tool
timestamp  instant in UTC
case_id  the case it belongs to (optional)
payload  data validated against the schema
```

Sender from the key, not the envelope
The presented key is what counts. If `source_tool` does not match, the event is dropped.

Each tool emits only its own types
Assets and vulnerabilities, network and threats, hardening and detection. Anything else is dropped and logged.

Facts, not orders
Even a block request is a proposal: the receiver submits it to its own policy.

External data is never trusted
Hostnames, user-agents and pcap content are marked as untrusted before they reach an AI agent.

Delivery and load under control. If a peer is unreachable the event stays in a bounded queue, retried at increasing intervals; writes that originate from external data are rate-limited. In PRO every event is copied to CyberShepherd, which keeps it for review and for playbook simulation.

Automated responses, under control

Capability policy, AI agent, promotion pipeline, playbooks and audit

3 · CAPABILITY POLICY

For every event → action → tool combination, three possible states

Off
The event is shown. No automated response.

Safe
The tool carries out a deterministic action by itself, mostly monitoring: an observing iptables rule, an auditd rule, a rate limit. No AI agent involved.

Action PRO ONLY
The action is intrusive and goes through an AI agent over MCP, with the chosen progression mode. An action cannot be downgraded to safe.

In the PRO edition the policy is defined in CyberShepherd's Integrations page (Events, Playbooks, Approvals, Audit tabs), distributed to the connected tools with an increasing version number and applied locally by each. If CyberShepherd is unreachable, the last received policy stays valid. For **critical assets** no action starts without a person's approval.

4 · ACTIONS AND THE AI AGENT

Typed tools only, never free-form commands: five controls on every call

Enablement
The tool exists only if the capability is in the Action state.

Validation
Parameters outside the schema: call rejected.

Scope
Only the hosts in the event and their declared clones.

Untrusted data
Event fields are kept apart from instructions.

Traceability
Source event, agent and outcome for every call.

5 · PROMOTION PIPELINE

Clones first, then production: three stages, three gates

G1 **Test** the fix is tried on the test clone → G2 **Pre-production** confirmed on an aligned environment → G3 **Production** applied only after two positive outcomes

Mode	G1 · start	G2 · test → pre-prod	G3 · pre-prod → production
Supervised	Person	Person	Person
Semi-automatic	Person	AI agent	Person
Automatic	Automatic	AI agent	AI agent, person notified

Declared rollback
Before starting. If a stage fails it is rolled back and the pipeline stops.

Timeout per stage
Where the AI agent decides, expiry counts as a negative outcome.

Veto checks
Agent heartbeat and critical services: they can only block, never approve.

Critical assets
The gate to production always stays in a person's hands.

6 · PLAYBOOKS

Declarative, never code

Trigger · Questions · Conditions · Response

Created with a guided builder of menus, each filtered by the previous one. A playbook combines existing blocks and adds no capabilities. In PRO you can **simulate** it on stored events before enabling it.

7 · ENVIRONMENTS AND CLONES

Declared by the customer

Each production host can have a test clone and a pre-production clone. The customer creates them and keeps them aligned; the suite clones nothing. A host that declares no environment is treated as **production**, that is, with maximum caution.

8 · AUDIT

Who can do what, and what was done

Each tool records in a log that cannot be edited from the interface: pairings and revocations, policy versions, MCP calls, promotions, rollbacks, approvers, dropped events and every time human validation is switched off, with author and date.

The scenarios: when the tools work together

From "resolved" to "closed": threat intelligence changes the weight of a vulnerability

SCENARIO 1

From indicator of compromise to coordinated remediation

IntelliDog
THREAT INTELLIGENCE

A CVE on your servers is now exploited in an **active campaign**, and some hosts have already contacted its **command and control servers**.

- **SentinelCore** — Raises the risk, notifies the owning teams, tries the fix on the test and pre-production clones.
- **FireDog** — Blocks the addresses involved according to indicator confidence and age, and captures their traffic from then on.
- **CyberShepherd** — Checks the AppArmor profile of the exposed service and enables audit rules on the hosts at risk.

HOSTS TRIAGED

- CVE + contact → incident
- Contact only → investigate
- CVE only → priority remediation

1. **Hunt** — Weekly search through history, or targeted at the first strong signal.
2. **Risk** — The "medium" CVE becomes urgent: priority and plan to the teams.
3. **Contacts** — Indicators are cross-checked against observed traffic.
4. **Containment** — Weighted blocking: no blind cuts on CDNs and hosting.
5. **Surveillance** — Targeted audit, including from public Sigma rules.
6. **One case** — Everything tied to the campaign; on closure temporary rules expire.

Monitoring responses. They block nothing and can start by themselves: observation rules, traffic capture, audit.

Intrusive responses. They go through AI agents that use predefined tools only, are tried on the clones and, unless the customer chooses otherwise, require a person's approval.

SCENARIO 2

Resolved is not closed

Resolved
The fix is installed. This is verified by comparing the host's packages and updates with the vendors' official advisories, in **OVAL** and **CSAF** format.

→

THREE CLOSURE CHECKS

- **Hardening** — controls re-run on the host
- **Exposure** — the service is no longer exposed as before
- **Persistence** — targeted hunt to rule it out

→

Closed
Only if all three checks are negative. The outcome is recorded with date and source, as **audit-ready evidence**; temporary rules expire.

THE MATURITY LADDER

1. Scan and patch in CVSS order
2. Priority on real risk: KEV, EPSS, asset. **SentinelCore already works here.**
3. Priorities updated by intelligence and a coordinated response across network, host and vulnerabilities
4. Verified closure, with evidence retained

BUILT ON OPEN STANDARDS

STIX · MISP · OVAL · CSAF · Sigma · MITRE ATT&CK

Data that is verifiable, interoperable and not tied to a proprietary format. The scenarios take you to steps 3 and 4: they require solid processes, are optional and are enabled playbook by playbook.

In the works

Where the suite is heading: what we have already planned to build

The suite is growing. Here is the direction.

SentinelSuite is a living project. Beyond what this datasheet describes, we have already scheduled four lines of development, designed for those who govern security, those who install it and those who use it every day.

A · FOR CISOS AND TECH MANAGEMENT

AlphaSentinel

A standalone executable for those who need the situation under control without opening the console of each tool.

- **Overall view** of the situation, with the most significant events up front
- **Mostly read-only activity:** you observe, with no risk of changing things
- **Targeted alerts** on what matters to decision-makers
- **Optional veto** on the actions that call for it
- And much more

B · FOR THOSE WHO INSTALL

The whole suite in a single release

Today every tool has its own virtual machine. It will be possible to have the entire SentinelSuite in just one.

- **A single Debian 13 VM** with all the tools of the suite
- **qcow2** and **ova** formats, as for current releases
- Fewer machines to create, update and watch
- The right choice for a quick start or for compact environments

C · FOR EXISTING INFRASTRUCTURES

More and more plugins

Deeper integration with third parties, to extend what you already have in your network instead of replacing it.

- New connectors to the tools you already use
- Data that today stays locked inside one product becomes **usable** by the whole suite
- And becomes **actionable**: it feeds events, playbooks and responses

D · FOR THOSE ON THE MOVE

Android and iOS apps

The suite within reach, even away from your desk.

- App for **Android**
- App for **iOS**

USING THIRD-PARTY TOOLS THAT ARE UNSUPPORTED OR LITTLE KNOWN?

Ask us for the integration.

If you want to use SentinelSuite and you want, or need, to keep using the tools you have already adopted in your network, but the right connectors or plugins are missing, ask us to build them.

Dognet Technologies is first and foremost a service provider that decided to release tools for people who, like us, do security. We know that adaptation, flexibility and integration are essential: building a component on request is not a problem for us.

Tell us what you use

info@dognet.tech
Tel. +39 352 068 1747

This page describes development directions that are already planned. Features and release arrangements may evolve.

Editions, compliance and distribution

Open Community, complete PRO: from technology to evidence

EDITIONS

What is included	Community	PRO
SentinelCore · FireDog · dog-agent	✓ **open source**	✓
MCP server on every tool	✓	✓
CyberShepherd	—	✓
IntelliDog (expansion module)	—	✓
Tools that can be connected	SentinelCore ↔ FireDog	all
Automated safe responses	✓	✓
Action responses with AI agent and pipeline	—	✓
Who defines the policy	each tool, locally	CyberShepherd
Playbook simulation · "closed" state	—	✓ **optional**

DISTRIBUTION

.ova · .qcow2 · .tar.gz · .deb · .rpm · git source

GPG-signed artifacts with **SHA256** checksums: verify before you install. Self-hosted deployment, even without Internet access.

Community: SentinelCore and FireDog are available, free and open source, from the website and GitHub.
PRO: licences for SentinelCore PRO and for CyberShepherd + IntelliDog are purchased and renewed on the customer portal www.sentinelsuite.online, with a personal login.

Coming up. The next release of FireDog and dog-agent brings support for Windows, Oracle Linux and Arch Linux and the integrations between the tools. CyberShepherd is at an advanced stage of development: at release it will include a read/write MCP server and the PRO integrations.

COMPLIANCE

How the suite supports NIS2, ISO/IEC 27001 and NIST CSF

SentinelCore

- Inventory of assets — **A.5.9**
- Management of technical vulnerabilities — **A.8.8**
- Risk assessment per asset — **ISO 6.1 / NIS2**

FireDog

- Network security and segregation — **A.8.20–8.22**
- Traffic monitoring — **A.8.16**
- Perimeter protection — **NIS2 Annex 1**

CyberShepherd

- Secure configuration — **A.8.9**
- Logging and event correlation — **A.8.15–8.16**
- Endpoint telemetry with dog-agent — **NIS2 Annex 1**

IntelliDog

- Threat intelligence — **A.5.7**
- Intelligence in analysis — NIST CSF **ID.RA-02**, **DE.AE-07**

With the PRO integrations: NIS2 art. 21(2) points b, e, f · ISO/IEC 27001:2022 A.5.7, A.8.8, A.8.16 · NIST CSF 2.0 ID.RA-01, ID.RA-02, ID.RA-05, ID.RA-06, DE.AE-07, RS.MI. A tool alone does not make you compliant: it provides the process and the evidence on which compliance is demonstrated.



LET'S TALK

Let's build your security perimeter.

From perimeter assessment to compliance, with the combination of tools and licences that best fits your organisation.

info@dognet.tech
Tel. +39 352 068 1747
Customer portal:
www.sentinelsuite.online
sentinelsuite.dognet-technologies.online
github.com/Dognet-Technologies

Dognet Technologies S.r.l. · Via XXV Aprile 47, 24055 Cologno al Serio (BG), Italy · VAT IT04867480164 · REA BG-495176 · SentinelSuite is a product of Dognet Technologies.