

Decoy services

Twenty decoys to choose from in Settings. The chip's network sockets limit how many run at once (7 TCP + 2 UDP on a standard build); the smart-home set marked • runs by default.

PORT	SERVICE	WHAT ATTACKERS SEE	WHAT KELPIELITE RECORDS
TCP 23 •	Telnet	Fake BusyBox login and root shell	Logins, commands, downloads
TCP 2323 •	Telnet (alt)	Same as port 23 (IoT botnet port)	Same as port 23
TCP 22 •	SSH	OpenSSH 8.4 banner, then a tarpit	Client software, time wasted
TCP 21	FTP	"220 FTP server ready"	User names, passwords, commands
TCP 80 •	HTTP	Fake router login page (Boa)	Logins, paths, exploit attempts
TCP 8080	HTTP (alt)	Same as port 80	Same as port 80
TCP 554 •	RTSP	IP camera: "401 Unauthorized"	First bytes of every request
TCP 1883 •	MQTT	Smart-home message broker	First bytes of every request
TCP 5555 •	ADB	Android TV / streaming-box debug	First bytes of every request
TCP 6379	Redis	Database: "NOAUTH" reply	First bytes of every request
TCP 5900	VNC	Remote screen: RFB 3.8 banner	First bytes of every request
TCP 3389	RDP	Windows Remote Desktop	First bytes of every request
TCP 445	SMB	Windows file sharing	First bytes of every request
TCP 3306	MySQL	Database server	First bytes of every request
TCP 8291	Winbox	MikroTik router admin	First bytes of every request
UDP 161 •	SNMP	Network device management	Every packet, never answered
UDP 1900 •	SSDP	UPnP device discovery	Every packet, never answered
UDP 5060	SIP	Internet phone (VoIP)	Every packet, never answered
UDP 69	TFTP	Firmware file transfer	Every packet, never answered
UDP 137	NetBIOS	Windows name service	Every packet, never answered

Safe by design

- **Nothing ever runs.** The fake shell answers common commands (uname, id, ls, cat /etc/passwd) from a script, and every download "times out". Nothing an attacker sends is executed.
- **Never a weapon.** UDP decoys record but never reply, so KelpieLite cannot be used to amplify attacks on others.
- **Bounded.** Up to 8 sessions at once, each capped at 10 minutes; idle sessions close after 10 to 60 seconds.
- **Tempting by design.** Default passwords such as root / xc3511 open the fake shell to capture what a botnet does next. It can also advertise itself as a file server.

Detection rules: decoys

Every rule runs on the device itself. Anything that touches a decoy is suspicious by definition, so these rules decide how serious it is.

DETECTION	TRIGGER	SEVERITY
Port scan	3+ different decoys from one address within 60 s	High
Brute force	12 password guesses from one address within 60 s	High
Coordinated campaign	6+ different addresses probe the decoys within 60 s	High
Exploit attempt	Path traversal, Log4Shell, command injection or known router exploit paths	High
Fake-shell login	A default password, or a third try	High
Scanning tool	A tool names itself, or sends 2+ Nmap-style probes	Warning
Shell command	Every command typed into the fake shell	High
Password guess	Each user name and password tried	Warning
Contact	Any connection to or probe of a decoy	Info Warning

How alerts reach you

Severity	High: act now. Warning: have a look. Info: logged only, never alerted.
Channels	Discord and phone push (ntfy) for warnings and high; text messages (Twilio) for high only.
Minimum level	Warnings and high by default; adjustable in Settings.
Per source	One alert per source every 10 minutes by default, unless its next event is more serious.
Overall cap	At most 8 alerts every 5 minutes, so an attack cannot flood your phone.
Daily summary	Optional Discord report of the last 24 hours at a chosen time.
Worst offenders	A 30-day tally per address (by hardware address where known), shown as a Pareto chart.
Attack types	A second 30-day Pareto chart by type of attack, with what each means and what to do.
Switching off	Six detection groups, and each decoy, can be switched off in Settings.
Status light	Every warning and high event pushes it from green towards red; it fades back within about an hour.

Detection rules: Wi-Fi

KelpieLite listens to your Wi-Fi channel (2.4 GHz) and checks the header of every frame it hears. Contents stay encrypted and unread.

DETECTION	TRIGGER	SEVERITY
Deauth attack	8+ disconnects in 10 s (3+ at KelpieLite; re-sends count once), checked against your AP's numbers and signal	High if forged Info if genuine
Evil twin	An unapproved access point using your network name	High if open Info if mesh-like
Cloned access point	Your AP's address on two channels, or from two radios	High
MAC spoofing	Two radios using one device's address at once	High
Fake channel switch	A "moving channel" notice that fails your AP's checks	High if forged Info if genuine
Handshake attack	A log-off message, or a flood of handshake messages	High (log-off) Warning (flood)
Sleep attack	Fake "going to sleep" frames from another radio	Warning
Karma attack	One AP answering for 5+ network names within 60 s	High
Fake-AP flood	15+ brand-new access points within 30 s	Warning
Join flood	30+ join requests in 10 s (WPA3 SAE floods named)	Warning
Wi-Fi scanning	40+ probe requests from one device in 10 s	Warning
Pwnagotchi nearby	A password-handshake collecting gadget's beacon	Warning
New device	Unknown device joins, after a 10-min learning period	Warning
Router impersonation	The router's hardware address changes (ARP spoofing)	High
Network scan	One device talks to 15+ different devices within 30 s	Warning
Unusual behavior	Machine-learning score above the threshold (page 4)	Warning High at 2x

AI and machine learning

Two complementary models run on the chip itself, in about 10.5 KB of memory. They read frame headers only: traffic is never decrypted, and nothing leaves the device.

MODEL 1 · LEARNS EACH DEVICE

Per-device baseline

Every 30 seconds each device is measured on 9 traffic features, read from frame headers only:

Frames sent	Frames rece...	Data sent
Data receiv...	Retransmiss...	Tiny-frame ...
Devices con...	Join / prob...	Broadcasts

Learning: 20 active windows (about 10 minutes of activity) per device.

Score: the two most unusual features, as weighted deviations on a log scale.

Thresholds: Cautious 4.0 · Balanced 3.0 · Sensitive 2.5.

Adapts: a change lasting 10+ windows becomes the new normal.

Explains: every flag names the features, e.g. "broadcasts 256 (usually about 5)".

MODEL 2 · WATCHES THE WHOLE NETWORK

Half-Space Trees

A streaming anomaly detector (Tan, Ting and Liu, 2011) shared by all devices. It spots a device that behaves unlike the rest, with no training data.

Structure: 12 random trees, depth 6.

Learning: rolling windows of 64 device-windows.

Flags: about the lowest 3% of scores.

Role: confirms borderline baseline scores (at least half the threshold).

Memory: fixed 8.8 KB.

Risk score: 0 to 100 per device, halving every 12 hours; touching a decoy adds 25.

MEASURED ON SIMULATED TRAFFIC

8 simulated home devices, about 197 device-hours: the loudest normal 30-second window scored 2.24, so the default threshold of 3.0 raised no false alarms. A hijacked smart bulb scored 6.6 and a broadcast storm 4.2 (both caught); a phone's burst of 40 probe requests scored 2.9 (caught on Sensitive). Real networks vary; the threshold is adjustable.

Platform

PROCESSOR

ESP32-S3, dual-core 240 MHz

DASHBOARD

Phone web app, password-locked

POWER

5 V USB

DECOYS AT ONCE

7 TCP + 2 UDP, chosen of 20

WI-FI

802.11 b/g/n, 2.4 GHz

ALERTS

Discord, push, SMS; daily report

SIZE

70.0 × 34.2 × 21.2 mm

30-DAY HISTORY

About 23 KB, two Pareto charts

DETECTION MEMORY

About 25 KB (models 10.5 KB)

STATUS LIGHT

RGB LED, green to red

HARDENING

800,000 fuzzed decoy sessions

TUNING

6 switchable detection groups