Technical Whitepaper

AI Smoke and Fire Detection: Platform Architecture, Security and Compliance

Audience: Heads of IT, HSE, Risk and Operations Leadership
Document: Technical Whitepaper v1.0
Published: April 2026

UK BUILT	CLOUD OR EDGE	ZERO TRUST	SOC 2 / ISO 27001
Designed and engineered in the UK	Customer choice of deployment topology	TLS 1.3, AES-256, RBAC, 2FA	Aligned to enterprise controls

CONTENTS

Technical Whitepaper

This whitepaper describes the Woodcock AI platform from the camera to the alert, explains how the same inference core runs in three deployment topologies, and sets out the security controls that apply in every topology. It is written for technical buyers, HSE and risk leaders, and operations teams assessing the platform for pilot or rollout.

ABOUT THIS DOCUMENT

Facts cited are sourced from the Woodcock AI Edge Device Security Overview (v. August 2025) and from the Woodcock AI public product pages at woodcock.ai. Performance figures such as visual detection in under four seconds refer to design targets and typical observed behaviour on representative industrial scenes, not guaranteed SLAs. Exact commitments will be included in contract schedules.

SECTION 1

The Detection Gap

Industrial fires rarely begin with flame. They begin with heat, then smoulder, then smoke, and only in the final stage do they reach the thresholds that conventional detection is built to see. In waste, recycling, battery handling, logistics and manufacturing, that final stage is too late: by the time a beam detector or sprinkler head is triggered, the material has often been feeding the event for many minutes.

Why conventional detection misses the window

Heat and particulate detectors are calibrated for enclosed, low-variance rooms. On a tip floor, in a baler, in a yard or in a high-rack warehouse, the air is noisy: dust, exhaust, weather, vehicle movement and thermal variance all conspire to raise the activation threshold. Operators compensate by setting thresholds high, which trades false alarms for latency. Thermal cameras help but are expensive per camera and struggle with line of sight through smoke; beam detectors are discrete, not continuous; aspirating systems are excellent indoors but ill suited to open sites.

Visual detection changes the shape of the curve

Smoke is visible before heat is measurable and before flame is present. A camera that can recognise smoke, flame and the textures around them, and that does so continuously on every frame of every feed, collapses the detection window from minutes to seconds. Woodcock AI is designed against a target of visual detection in under four seconds of the earliest visible smoke or flame on representative industrial scenes. In combination with a relay output that can trigger a fire panel directly, this turns every existing CCTV camera into an additional line of early warning.

< 4 s	**3**	**1,200+**
target visual detection time from earliest visible smoke or flame	deployment topologies: our cloud, your cloud, or on-prem / edge only	object classes, including smoke, fire, human, vehicle, injury, leak and weapons

SECTION 2

Platform Overview

Woodcock AI is an AI video analytics platform that ingests live CCTV feeds over RTSP, runs continuous smoke and fire detection plus a broader set of object classes, and raises real time alerts through the customer's chosen channels. The same inference core is deployed in three topologies: managed in our UK cloud, managed inside the customer's own cloud account, or run entirely on an on-premises edge device that never leaves the customer site.

The platform is camera agnostic. Any CCTV camera that exposes an RTSP stream or is ONVIF compliant can be onboarded. Existing cameras, existing NVRs and existing cabling are reused; there is no rip and replace.

FIGURE 1 End-to-end reference architecture

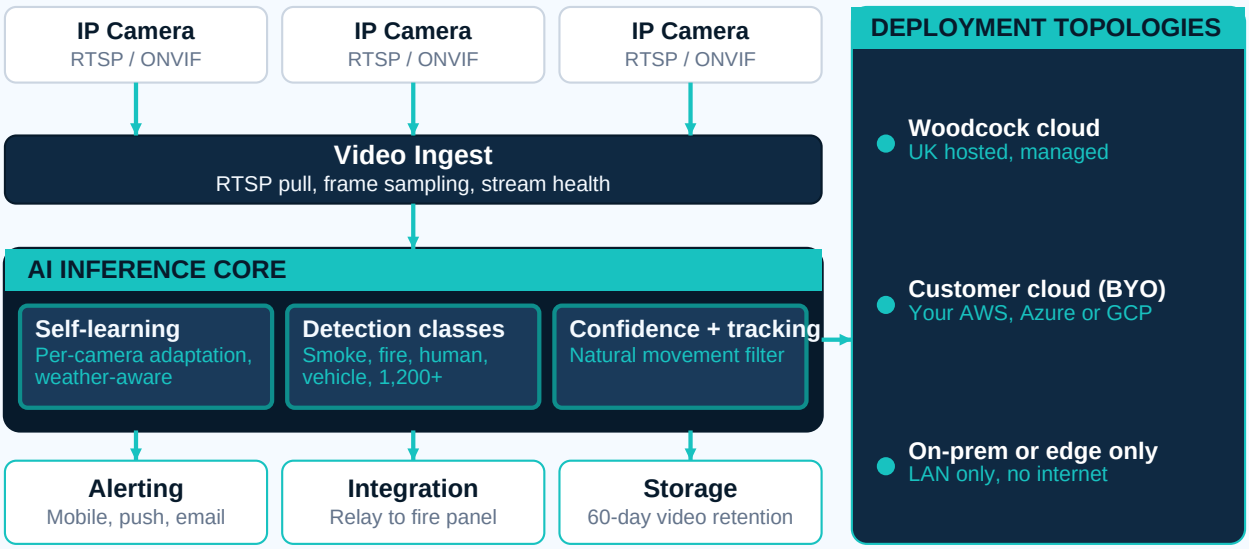

Data flows top to bottom in the main column. The same inference core runs in all three deployment topologies on the right.

Figure 1 shows the reference architecture. Read left to right: cameras publish RTSP streams into the ingest layer, frames are sampled and passed to the AI inference core, detections are emitted to the alerting, integration and storage layers, and the whole pipeline is packaged to run in whichever topology the customer chooses. Security controls described in Section 7 apply to every topology equally.

SECTION 3

Video Ingest and Processing Pipeline

The ingest layer is built around standard protocols so that it can work with the vast majority of CCTV camera manufacturers without bespoke drivers. Three principles guide the design: use the customer's existing cabling and cameras, fail closed on stream loss rather than silently stop detecting, and expose enough operational telemetry that IT teams can monitor health without access to raw footage.

3.1 Supported inputs

The platform accepts any camera that exposes an RTSP stream. ONVIF discovery is supported where enabled on the camera. Both H.264 and H.265 codecs are handled. Resolutions from CIF up to 4K are supported; in practice, 1080p at 10 to 15 frames per second is a good balance between scene detail and network load. Streams can be pulled directly from the camera, from a camera manufacturer's NVR, or from a VMS that re-exposes RTSP.

3.2 Stream lifecycle and health

Each stream runs inside an isolated worker. Workers detect stream drops, frame starvation, codec errors and clock drift, and attempt exponential-backoff reconnection. An operations dashboard surfaces per-camera uptime and last-frame timestamps without ever exposing the content of the video to the dashboard viewer. Stream loss that persists beyond a configured threshold raises a dedicated alert so that blind cameras are not silently tolerated.

3.3 Frame handling and privacy

Frames are decoded and scored in memory. In every topology, raw frames are not persisted by default; the system stores only a rolling buffer sized to support retrospective review of alerts, together with metadata about the detection (timestamp, class, confidence, bounding region). In the customer's preferred topology, video retention can be extended to a standard 60-day window to support investigation workflows.

3.4 Relay outputs

The edge appliance includes configurable relay outputs that can be wired directly into a fire panel, a siren, a shutter, a dust-suppression system or any other on-site safety system. This creates a hard, network-independent path from visual detection to physical response: even if site networking is degraded, a detection can still trigger on-site mitigation.

SECTION 4

Deployment Topologies

The same inference core runs in three topologies. The choice is a commercial and security decision, not a technical one: the detection logic, the model behaviour and the security controls are identical. This section summarises the trade-offs so that IT, HSE and procurement can align on the right option for each site.

	Woodcock Cloud	Customer Cloud (BYO)	On-premises / Edge Only
Where inference runs	Woodcock UK cloud	Customer's AWS, Azure or GCP	Edge appliance on site
Internet dependency	Outbound only, site-to-site VPN	Outbound only, inside tenant	None required
Management	Fully managed by Woodcock AI	Joint: customer infra, our app	Private VPN or local admin
Data residency	UK hosted	Customer choice	Customer site only
Best for	Multi-site rollouts with standard IT	Regulated enterprises with a cloud policy	Air-gapped, remote or constrained sites
Update path	Continuous, managed	Scheduled with customer change windows	Signed updates via VPN or physical
Relay outputs	Via edge gateway	Via edge gateway	Direct on the edge appliance

4.1 Topology selection guidance

Most multi-site customers start in the Woodcock cloud for speed of rollout, then move select sites to edge only where bandwidth, latency or air gap rules require it. Customers already operating a strict cloud policy typically choose BYO cloud so that footage and metadata remain inside their existing security boundary. Mixed estates are supported: one site can be edge only while another is in our cloud, under a single management plane and with unified alerting.

SECTION 5

Detection Intelligence

The detection layer is where Woodcock AI differs most from first-generation video analytics. Rather than a single fixed model pushed to every camera, Woodcock AI deploys a self-learning system that adapts to each individual camera's scene and is continuously retrained to separate true smoke and fire from the ordinary visual noise of industrial environments.

5.1 Self-learning per camera

Every camera sees a different scene: a different angle, lens, lighting profile, weather exposure, dust load and traffic pattern. A model that is excellent on a baler will produce false alarms on a tip floor if it has not learned that particular visual context. Woodcock AI pioneered self-learning artificial intelligence for smoke and fire detection, with the system adapting to each individual camera environment and constantly retraining itself to understand the true presence of smoke or fire. It learns to recognise natural movement and adapts to changes in weather that would otherwise produce confusion.

5.2 Detection classes

The primary classes are smoke and fire. Beyond those, the platform also detects a broader set of safety and security events, including human, vehicle, injury, leak, weapons and more than 1,200 additional object classes. Customers can enable classes selectively per camera so that, for example, a tip floor camera is tuned for smoke and fire while a gatehouse camera is tuned for vehicle and human.

5.3 False alarm management

False alarms are the single biggest reason security operations teams disable analytics. Woodcock AI manages them at three layers: at the model, through per-camera adaptation; at the scene, through geometric masks, time-of-day rules and weather aware filters; and at the operator layer, through a feedback loop where operators mark events as true or false and the system uses this feedback to refine detection over time.

SECTION 5 (continued)

Detection Intelligence

5.4 Performance framing

Woodcock AI is designed against a target of visual detection in under four seconds of the earliest visible smoke or flame on representative industrial scenes. This is stated as a design target and typical observed behaviour rather than a universal SLA: real performance is scene dependent. Factors that move the figure up or down include camera resolution and frame rate, field-of-view coverage of the critical area, lighting, weather, occlusion and how long the system has had to adapt to that camera. A pilot measures this on the customer's real feeds before any commercial commitment is made.

5.5 Model update path

Model updates are delivered as signed artefacts and applied through the same change-control process as application updates. In cloud topologies, updates are continuous and managed by Woodcock AI; in customer cloud topologies, they are applied on the customer's change windows; on edge-only deployments, they are delivered through the private VPN or, where internet access is forbidden, as physically signed artefacts verified on the device. In all cases the operator has visibility of which model version is running on which camera at which time.

5.6 Explainability and operator trust

Every detection produced by the platform carries a timestamp, a class label, a confidence score and a bounding region on the frame. Operators can step forward and backward through the event clip to see exactly what the system reacted to, then confirm or dismiss the event. This transparency is essential for the HSE workflow and is also what makes the feedback loop described in 5.3 possible.

DETECTION CLASSES AT A GLANCE

Primary safety classes. Smoke. Fire.

Safety and ops. Human, vehicle, injury, leak, weapons.

Extended catalogue. More than 1,200 object classes, enabled selectively per camera.

SECTION 6

Alerting and Integration

A detection is only useful when it reaches the right person, on the right channel, in the right amount of time. Woodcock AI fans alerts out across several channels in parallel so that a detection is never waiting on a single device to wake up.

6.1 Mobile application

A native mobile application is available for iOS and Android. User sessions to the application run over end-to-end TLS 1.3 encryption, with device and user authentication. The mobile app can connect directly over the local network when the wider internet is unavailable, with the same authentication and encryption controls applied, preserving the zero-trust posture even in isolated environments.

6.2 Push and email notifications

Real-time alerts to mobile devices are delivered through Apple Push Notification Service (APNS) and Firebase Cloud Messaging (FCM) over secure, standards-based channels. Scheduled or critical alerts are sent via an ISO-27001-certified email provider with TLS 1.3 transport security. Recipient lists, severity filters and quiet hours are configurable per site and per camera group.

6.3 Integration with existing safety systems

Edge appliances expose built-in relay outputs that can trigger a fire panel or any other safety system instantly, providing a network-independent path from visual detection to physical response. The platform also exposes event webhooks and a REST API so that alerts can be forwarded into existing ARC, SOC, CAFM or ticketing workflows. Customers who already run a VMS typically keep it as their viewing front end, with Woodcock AI as the analytics and alerting layer behind it.

6.4 Video retention

A 60-day rolling backup of event video supports retrospective review, HSE investigation and insurer reporting. Retention, access and deletion are governed by the controls described in Sections 7 and 8.

SECTION 7

Security Architecture

The platform is engineered for enterprise-grade security using a defence-in-depth strategy with zero-trust access control, TLS 1.3 encryption and continuous monitoring to ensure data confidentiality, integrity and availability. It is aligned with SOC 2 Type II and ISO 27001 standards and employs role based access control, tamper-proof audit trails and annual penetration testing. Every control described in this section applies equally to the edge appliance and to the cloud components.

7.1 Principles

Three principles shape the security model. **Multi-layered protection**: network, application and data-layer safeguards work in tandem to block, detect and contain threats. **Zero trust**: every request, internal or external, is authenticated and authorised; implicit trust is never granted. **Least privilege**: only essential services run on each device, reducing the attack surface.

7.2 Connectivity

Remote administration runs over SSH through a private VPN tunnel with RBAC and mandatory two-factor authentication. Mobile app user sessions run over HTTPS with end-to-end TLS 1.3 and device plus user authentication. Outbound cloud AI communication uses encrypted traffic routed through a site-to-site VPN, requiring no inbound firewall exceptions. Push notifications use APNS and FCM; email uses an ISO-27001-certified provider with TLS 1.3.

7.3 Data protection

All traffic leaves the device wrapped in TLS 1.3 with modern cipher suites (ECDHE with AES-GCM-256). System drives and embedded databases are protected with AES-256 disk encryption. Decryption keys are hardware-bound; drives remain unreadable outside the device. USB and other maintenance ports are disabled by default and require a signed authorisation token to enable.

7.4 Identity and access

Role-based access control ensures users see only what they need. Two-factor authentication is enforced for privileged roles. Passwords and tokens are stored using PBKDF2-SHA-512 with per-user salt.

SECTION 8

Data Handling, Privacy and Compliance

CCTV footage is sensitive personal data. The platform is built so that the customer remains the data controller and Woodcock AI is a processor acting under instruction. Each topology described in Section 4 produces a different data-flow map; the controls in this section apply to all of them.

8.1 Residency and topology

In the Woodcock cloud topology, data is hosted in the United Kingdom. In the customer cloud topology, data remains entirely inside the customer's own tenant in the region of their choice. In the edge-only topology, footage and metadata never leave the customer's site. Customers can select the topology per site according to their regulatory and commercial requirements.

8.2 Third-party assurance

All cloud and notification providers undergo annual SOC 2 Type II and ISO 27001 audits. Outbound access from Woodcock AI components is restricted to pre-approved domains only; there is no unrestricted internet egress. Container-level network segmentation contains potential breaches.

8.3 Audit trails

Every system event and user action is hashed, time-stamped and appended to an immutable log store. Logs are encrypted and off-loaded to a write-once cloud bucket with lifecycle management. Internal log reviews run weekly, with quarterly internal security audits.

8.4 Customer controls

Customers retain the ability to request extracts, restrict access by role, set per-camera retention policies and remove footage on request. Woodcock AI can provide a Data Protection Impact Assessment template and a standard Data Processing Addendum to support the customer's UK GDPR and DPA 2018 obligations. Detailed audit reports and penetration test results are available to the customer's security consultants under NDA.

SECTION 9

Resilience, Monitoring and Recovery

Detection has to work on the worst day, not the best. The platform is engineered so that each of its components has a defined failure mode, a defined observation path and a defined recovery path.

9.1 Backups

Differential backups run hourly; full snapshots run daily. Backups inherit TLS 1.3 during transfer and AES-256-GCM in storage. Disaster-recovery playbooks are exercised bi-annually to verify RPO and RTO targets.

9.2 Monitoring and vulnerability management

Authenticated vulnerability scans run weekly. CVEs with a score of 7.0 or higher trigger immediate patching. Independent security firms conduct black-box and white-box penetration tests annually; critical findings are remediated within thirty days. These practices align with recognised critical and high severity vulnerability management standards.

9.3 Offline and local-network operation

The mobile app can connect directly to the edge appliance over the local network when the wider internet is unavailable. All of the same authentication and encryption controls apply, preserving the zero-trust posture even in isolated environments. The edge appliance itself operates independently of internet connectivity; with its built-in relay outputs it can trigger fire panels or other safety systems even during a complete wide-area network outage.

9.4 Capacity and scaling

Each edge appliance is sized against a target camera count. Within the cloud and customer-cloud topologies, the inference layer scales horizontally so that adding sites or cameras does not require customers to re-architect their deployment. Woodcock AI monitors per-camera health, per-stream latency and per-site throughput and raises capacity events proactively.

SECTION 10

Operational Commitments

This section summarises the operational expectations customers can set with Woodcock AI at the point of pilot and rollout. Exact figures are set per customer in the contract schedules; the figures here are typical.

Pilot duration	Typically 30 to 90 days on a defined camera set
Install effort per site	Network and power only; existing CCTV reused
Onboarding support	Joint call with IT and HSE; per-camera class tuning
Update cadence	Continuous in our cloud; change-window aligned elsewhere
Critical CVE remediation	Patched within 30 days of disclosure
Backup RPO target	Hourly differential, daily full
Third-party assurance	Annual SOC 2 Type II and ISO 27001 provider audits
Audit evidence	Available under NDA to customer security teams
Exit and portability	Export of events, metadata and event video on request

10.1 From pilot to production

A pilot is defined around a specific risk question at a specific site: typically a tip floor, a baler, a battery storage area or a high-rack aisle. Success criteria are agreed up front (time to detect, false alarm rate, operator feedback) and measured on the customer's real footage. At the end of the pilot, the decision to scale is based on observed results, not on forecast metrics.

10.2 Rollout across sites

Multi-site customers typically roll out in waves, with a standard site pack that bundles install, camera audit, class tuning and operator training. A single management plane provides per-site, per-camera health and detection telemetry so that central HSE, security and IT teams have a unified view without touching raw footage.

REFERENCES

Sources, Contact and Next Steps

Primary sources for facts cited

1. Woodcock AI Ltd., *Edge Device Security Overview: A Comprehensive Security Framework for Enterprise Deployment*, 6 August 2025. Security controls, encryption, access, audit, backup and vulnerability management claims in Sections 6, 7, 8 and 9 derive from this document, which the customer may request under NDA in full.

2. Woodcock AI public product pages at woodcock.ai, including the AI Smoke and Fire Detection and Smoke and Fire Detection Camera pages, and the product specification page. Product capability claims (self-learning per camera, under-four-seconds visual detection target, 60-day video backup, 1,200-plus object classes, relay outputs, cloud or edge deployment, UK build) derive from these pages.

3. Woodcock AI internal design targets for visual detection latency. Figures in this document are framed as design targets and typical observed behaviour rather than contractual SLAs; contractual commitments are made in customer-specific schedules.

Independent and contextual reading

Recent academic and industry work on vision-based fire and smoke detection reinforces the case for a CCTV-based approach alongside traditional detection, including research covered in the International Fire and Safety Journal on AI systems that detect fires via standard security cameras before alarms sound, and survey work on advancements in fire and smoke detection for indoor and outdoor surveillance feeds in ScienceDirect (2024). These are provided as context only; claims about the Woodcock AI platform in this document are sourced from primary references 1 and 2.

Next steps

A 60-minute technical call with your IT, HSE and security leads is the fastest way to pressure-test this architecture against your estate. We will walk through the topology you prefer, identify the first pilot site, and agree the success metrics that matter to you.

Email. hello@woodcock.ai **Web.** woodcock.ai