VCAcore Release Notes

Version: **2.4.0**

Release date: **12/11/2025**

VCAcore is an analytics engine developed by VCA Technology, available as a standalone application for Windows and Linux (VCAserver) and as the VCAsdk libraries which can be integrated into third party applications such as camera firmware or embedded into a VMS.

The release notes outline all the changes made to VCAcore in this release period. Changes and features described are valid for all platforms which VCAcore is distributed on (VCAserver and VCAsdk) unless otherwise stated.

Index

1. Release Summary
    1. Supported Products
    2. New Features
    3. Bug Fixes
2. Known Issues
3. Deep Learning Model History

4. Dependant Libraries/Applications
5. Test Environment

Release Summary

This is a major release, bringing a the new algorithms, new rules and a host of usability improvement and bug fixes.

Supported Products

- VCAserver (Ubuntu 22.04 x86, Ubuntu 20.04 ARM/Jetpack 6)
- VCAsdk (Windows 10 x86, Ubuntu 22.04 x86, Ubuntu 20.04 ARM/Jetpack 6)

New Features

Deep Learning Thermal Tracker (DLTT)



Further increasing the tracing options available; a Deep Learning Thermal Tracker has been developed to enable the classification and tracking of people and vehicles in thermal camera streams. The DLTT is included in the ProAI license

Face Attributes

A new face algorithm has been added to VCAcore. This algorithm performs a number of functions including the generation of:

- Age, Gender and Emotion metadata

- Face Covering data for use with a new Face Accessory Filter rule

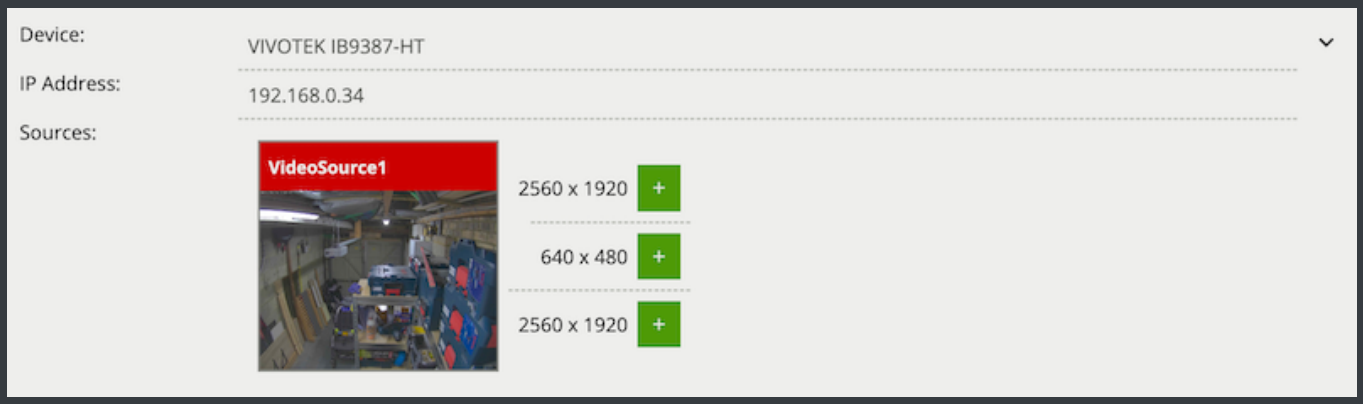
This feature has a new VCAface license

Smoke and Fire

A new Smoke and Fire algorithm has been added to VCAcore. This new tracker identifies and tracks Smoke and Fire as distinct objects that can be used with the extensive VCAcore rule set. This allows for both the detection, as well as absence, of smoke and fire objects, applicable to health and safety early warning systems but also industrial process monitoring functions.

The feature has a new VCAsmokeFire license

ONVIF Camera Detection



Onvif discovery has been added to the Edit Sources page; allowing users discover ONVIF compatible devices on the local network. If the correct credentials are supplied then the provided profiles of that device can also be added directly as an RTSP source

Prebuilt Models

When using Nvidia GPUs, an engine building process is run that optimises our deep learning algorithms to that specific hardware, leading to increased channel capacity. This optimisation step can take a long time depending on the hardware. The engine builder will now download these engines, where possible, from a VCA repository for some of our common platforms. Allowing the engines to be downloaded, as opposed to built, reduces the engine builder run time to a few minutes. Please note that the hardware must have access to the internet to take advantage of this.

Camera GPS Position and Calibrated Position



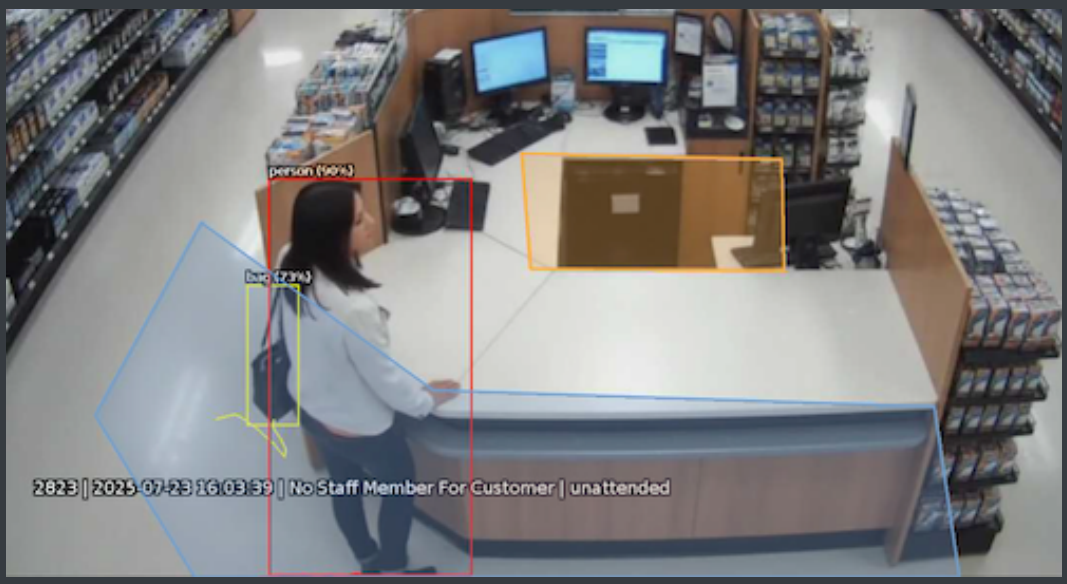
VCAcore has, for a long time, supported the notion of calibration in camera views, providing estimated speed and size of objects tracked using the trackers. This release extends the metadata generated by this algorithm by including both a relative position from the camera (Calibrated Position), but also, with the provision of a camera GPS position and orientation, an extrapolated GPS position for a given tracked object.

Occupancy, Unattended and Absence Rules

This release introduces a number of new rule concepts, some to speed up deployment, some boiling complex logical rule graphs down to a single rule.

Occupancy is a combined presence, object filter and counter in a single rule. Allowing for very quick configuration for queue monitoring applications, car parking or EV charging sites.

Unattended provides a one rule configuration to generate events when one zone is occupied whilst another is not. Targeting retail cash registers or information desks, this rule allows for real time notification for when a customer is waiting and there is no staff member to serve.



Lastly, Absence, which is a simple rule that allows for events when a given zone has not had an object of a specified class present within it for a defined duration.

Zone Mask and Colour Events

Zone mask is a new zone property allowing the VCA rules engine to use a zone as a tracked object. This allows the generation of a colour signature for any defined area and use that, via a Zone Colour rule, within the logical rules. This allows for static monitoring of areas for colour changes like traffic lights or railway crossings and use those light changes to enable or disable given rule configurations.

Host Time settings for ABP range

The ABP range of hardware now supports the configuration of the local time settings including, setting manual time, time zone and NTP server.

Evaluation License Expiry



When using an evaluation license, when the time remaining on that license drops below `14 days` a warning message will now be displayed in the channel preview.

In addition, the License Source can now generate an event when a Eval license has less than `7 days` remaining, which can be used to trigger any configured actions.

RTSPs Support

RTSPs support is now provided for both RTSP sources as well as when streaming out of VCAcore, combined with the inclusion of HTTPs support to the web server in a previous release, the end to end pipeline of camera to metadata generation and delivery can be fully encrypted.

Bug Fixes

This release contains fixes for the following issues:

- Fixed missing hard hat detections on some scenes with DLAD
- Fixed issues with tracker jumping to unrelated objects
- Fixed issue where some relevant settings were not shown in UI for some trackers
- Fixed intermittent bug with ISO8601 timestamp conversion
- Fixed Manager restart function being unavailable when status is failed
- Fixed issue with multiple engine builder instances conflicting with each other
- Fixed erroneous warnings in logs
- Improved stationary object detection
- Improved reporting of inference backend failures
- Fixed alignment issues in UI
- Fixed hardware info showing incorrectly on Jetson devices
- Reduced false positives in directional crossing algorithm
- Fixed issue with Stopped rule only working for a single object in the zone
- Fixed memory leak in ONVIF media service
- Added missing authenticaton check in ONVIF media service
- Fixed Not rule false negative
- Fixed issue with Colour Signature not working for DLST

- Fixed metadata bug causing duplicate counts with Stopped + Counter rules
- Display language names in native language rather than all in English
- Fixed RTSP streaming issue with certain sources
- Fixed web server intermittently getting stuck on shutdown
- Improved zone manipulation UI for small zones
- Fixed text rendering bug in BIA with long event names
- Use CUDA graphs with TensorRT inference backend to improve efficiency
- Added X-Content-Type-Options header to http responses
- Improved manager service logging when child process exits unexpectedly
- Fixed bugs with Linux installer command-line options
- Fixed bug when switching between local and cloud licensing
- Fixed bug with single-character Action bodies
- Fixed issue where certain file sources would not stream
- Fixed UI themeing issues
- Fixed issue where pipeline would occasionally fail to restart after crop change

Known Issues

Outlined below are known issues that are under consideration by the development team.

- With the release of v2.3.2 for Jetpack 6 it is advised that due to the fixes made for the networking, the ABP Pack software should be reinstalled.
- With the release of v2.3.0 there has been a **breaking change** to the event `type` strings. Previously these were `Hands Up`, `AggressiveBehaviour` and `Stopped`, these have been renamed to `hands-up`, `aggressive-behaviour` and `stopped` respectively, in line with the rest of the event type strings
- With the release of v2.2.1 there has been a **breaking changes** to the VCAsdk and REST API. The `Hands Up` observable `threshold` property has been renamed to `confidence_threshold` in line with the rest of the observables

Deep Learning Model History

Below are the deep learning models included in this release.

Algorithm	Version	Comments
DLOT	v4.20.1	Included data from all available dataset sources
DLPT	v2.1.1	New architecture
DLTT	v1.1.5	First Release
DLST	v2.2.6	Reworked training approach, additional data
DLFT	v1.1.3	First release
DLF	v2.18.0	Added thermal data and animal as background
Fall	v1.1.6	Redefined fall in dataset to reduce sitting FP
Aggressive	v1.1.12	Tuned for sensitivity
HOI	v2.1.1	Added person class
REID	v4.0.2	Smaller model architecture
DLAD (torso)	v1.0.59	Added data
DLAD (head)	v2.0.66	Added data

For support queries and documentation please see VCA Technology support page.

VCA Technology Ltd
15, Riverside Court
Beaufort Park
Chepstow
Monmouthshire
NP16 5UH
United Kingdom

www.vcatechnology.com