



VisualCortex User Documentation

Unified, Real-Time, AI Driven Video Intelligence

Contents

1. Introduction

2. Video Sources

- 2.1 Introduction
 - 2.2 Live Video
 - 2.2.1 Technical Requirements
 - 2.2.2 Adding cameras via ONVIF scan
 - 2.2.3 Manually adding a Camera
 - 2.2.4 What's next?
 - 2.3 Archived Footage
 - 2.3.1 Virtual Cameras
 - 2.3.2 Technical Requirements
 - 2.3.3 What's next?
-

3. Concepts

- 3.1 Architecture
 - 3.1.1 Concepts
 - 3.2 Alerts
 - 3.2.1 Introduction
 - 3.2.2 Alert Types
 - 3.2.3 Notification Targets
 - 3.3 Credits
 - 3.3.1 Obtaining more credits
 - 3.4 Dashboards
 - 3.4.1 Creating charts
 - 3.5 Edge Clusters
 - 3.6 Event Definitions
 - 3.6.1 Directional Count
 - 3.6.2 Zones
 - 3.7 Folders
 - 3.8 Models
 - 3.8.1 Applying models to video
 - 3.8.2 Processing on Virtual Cameras
 - 3.9 Permissions
 - 3.9.1 Folder Permissions
 - 3.9.2 Roles
 - 3.9.3 Administrative
-

4. Investigations

- 4.1 Introduction
 - 4.1.1 Sources
 - 4.1.2 Models
 - 4.1.3 Searches
 - 4.1.4 Reports
- 4.2 Investigation Dashboard

4.2.1 Navigation

4.2.2 Sources

4.2.3 Models

4.2.4 Activity

4.2.5 Annotations

4.3 Bulk Import

4.3.1 Selecting or importing

4.3.2 Setting Video Information

4.4 Searches

4.4.1 Filters

4.4.2 Results

4.4.3 Detections

4.5 Report

5. Integrations

5.1 Authentication Providers

5.2 Event Integrations

5.3 Video Integrations

5.3.1 Integration types

5.3.2 Importing into VisualCortex

5.4 VixVizion

5.4.1 In Investigations

5.4.2 Searching un-enrolled faces

5.5 Video Management System Integrations

5.5.1 Functionality offered

5.5.2 Support Matrix

6. Administration

6.1 Code Transformations

6.1.1 Sample Use Cases

6.2 Edge Clusters

6.2.1 Technical Details

6.2.2 Management

6.2.3 Settings

6.2.4 Provisioning

6.3 Groups

6.4 IP Whitelisting

6.5 Organisations

6.6 Roles

6.7 Users

6.7.1 Provisioning

6.7.2 Groups

1. Introduction

Welcome to VisualCortex.



VisualCortex is a computer vision platform which makes it easy to administer and run computer vision models on your cameras and get real business insight out of them.

Common use cases include:

- People counting and analytics
- License Plate Recognition (ANPR / LPR)
- People and Vehicle attributes
- General object detection (planes, bags, boxes, luggage etc)

This documentation site outlines how to configure and use the VisualCortex platform.

2. Video Sources

Learn about how to add video into VisualCortex.

Introduction

VisualCortex is a computer vision platform, which processes video and turns it into data. Conseq...

Live Video

VisualCortex is right at home with live video. Live video can be processed within seconds, and ca...

Archived Footage

While VisualCortex works well with Live video sources, it also has first class support for Archived ...

2.1 Introduction

VisualCortex is a computer vision platform, which processes video and turns it into data. Consequently, video data is essential for getting value out of the VisualCortex platform.

Video data comes from two main sources:

- Live Camera Feeds
- Archived Footage

Camera feeds can be ingested either directly from cameras, or can be ingested via Video Management Systems (VMS), or Network Video Recorders (NVR).

VisualCortex can ingest archived footage from a number of sources. These are outlined in the Archived footage section of the documentation.

2.2 Live Video

VisualCortex is right at home with live video. Live video can be processed within seconds, and can trigger external services and send alerts in near real-time.

2.2.1 Technical Requirements

VisualCortex supports several video transport mechanisms. These include:

- Real-Time Streaming Protocol (RTSP)
- Real-Time Streaming Protocol Secure (RTSPS)
- Motion JPEG (MJPEG)
- HTTP Live Stream (HLS)

Within RTSP(S), VisualCortex supports the following protocols:

- H264
- HEVC / H265

VisualCortex supports a minimum resolution of 200x200, and a maximum resolution of 4K (3840x2160). VisualCortex supports from 1-30 frames per second (fps).

Your cameras must be accessible via network from the [Edge Cluster](#) you add them to.

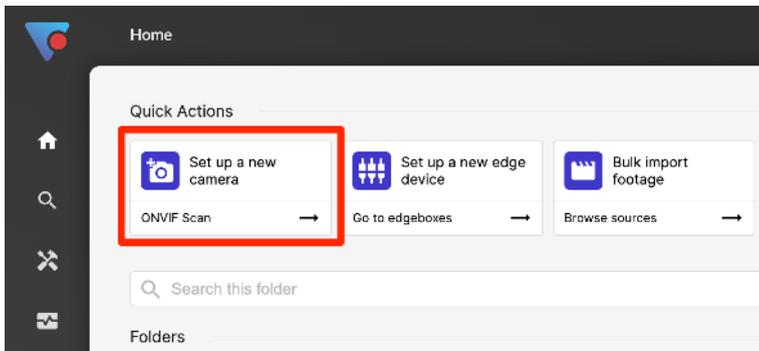
2.2.2 Adding cameras via ONVIF scan

The easiest way of getting live video into VisualCortex is via the ONVIF scanning feature. This allows you to add large numbers of cameras quickly and easily.

ONVIF is an open industry standard for IP-based physical security products. Most IP cameras, Video Management Systems (VMS) and Network Video Recorders (NVR) support ONVIF. Check with your vendor to see if ONVIF is supported

It's easy to get started with ONVIF scanning. The only prerequisite is to have an [Edge Cluster](#) which has access to the Camera/VMS/NVR you want to scan.

As an [Administrator](#), Simply click the *ONVIF Scan* icon at the top of the screen.

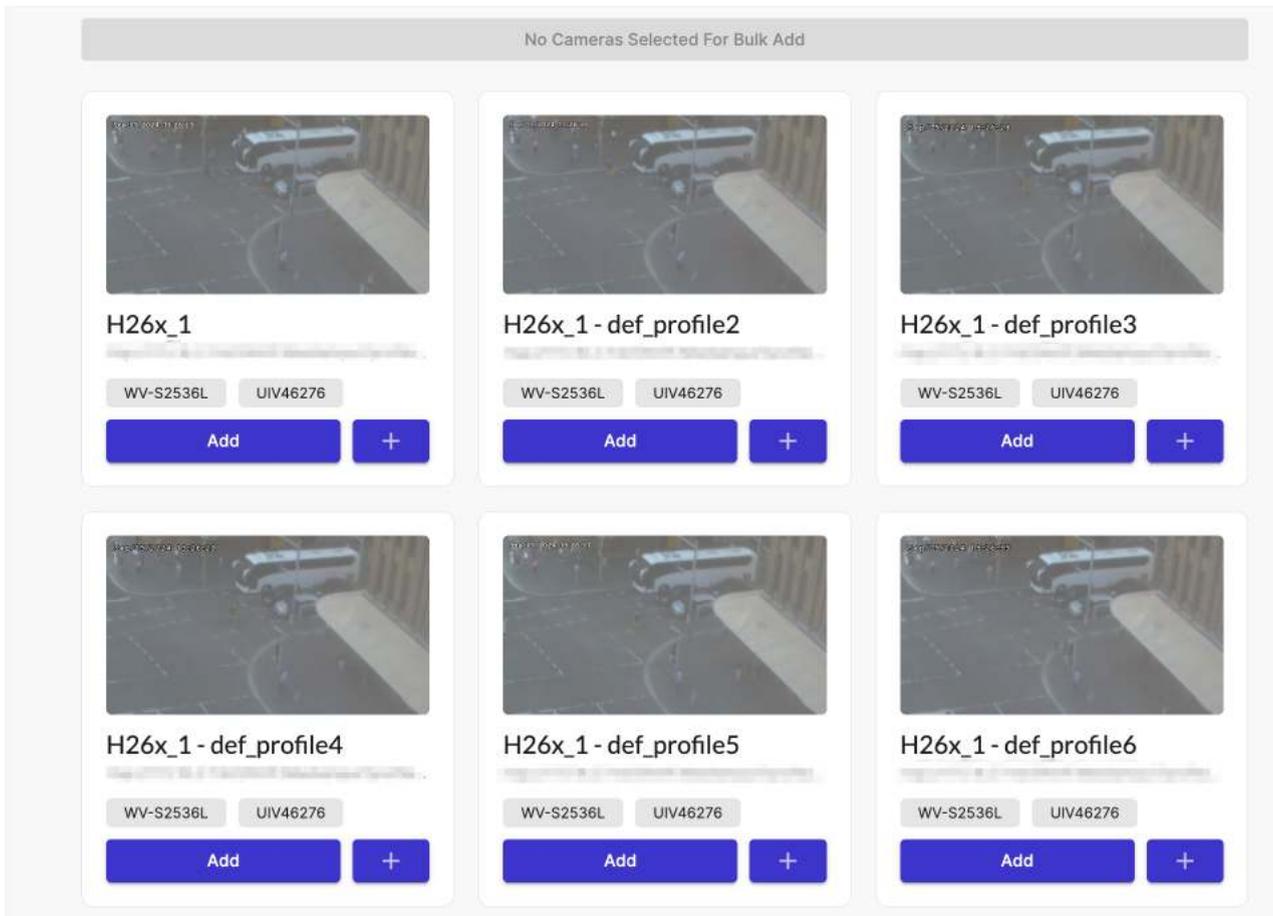


From here, you'll be presented with a wizard. The first step is to select an [Edge Cluster](#). This cluster *must* be able to access the live video source. Make sure appropriate firewall rules are in place to allow this.

In step 2, you'll get to fill out the information for an ONVIF scan. In the **IP Scan Range**, you can either put a single IP address if you know the IP of your Camera / VMS / NVR, or a range like `192.168.1.4-192.168.1.254`. Ranges are useful if you have multiple cameras on the same network which use the same credentials.

Make sure you put in your username and password, as well as the ONVIF scan port. This is usually `80`, but is often a separate port on VMSs or NVRs - check with your vendor documentation.

Once you've filled in the information, click **Scan**. VisualCortex will attempt to make an ONVIF connection, and you'll be greeted with a screen similar to the one below.



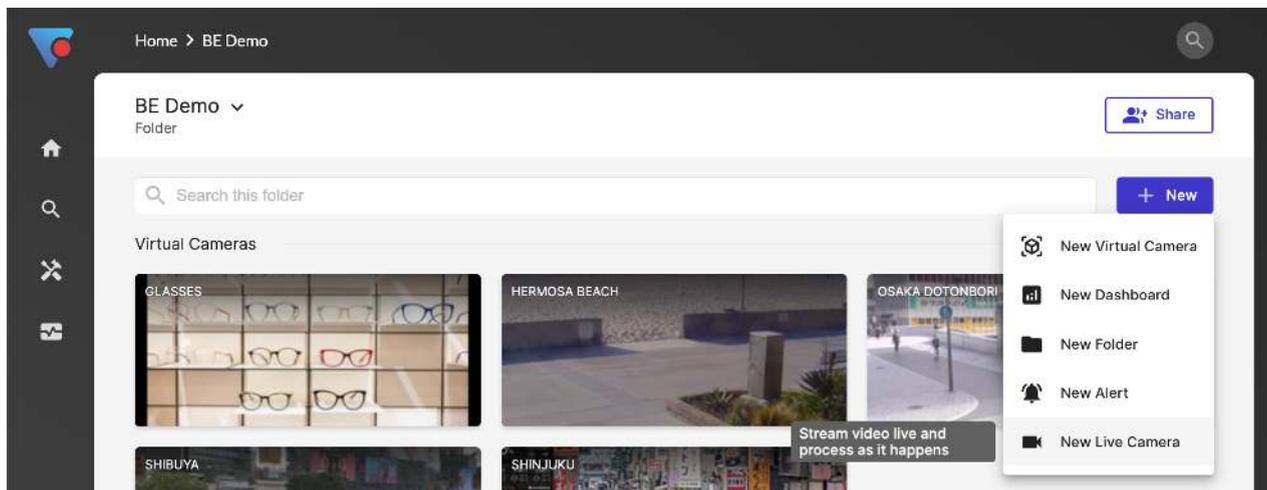
From here, you can either click **Add** to add a single Camera. VisualCortex will ask you what [folder](#) you want to add the camera to.

If you want to add many cameras to the same [folder](#) you can click the **+** icon to choose the cameras you want, then select the **Bulk Add Selected Cameras** button. You'll be prompted for a [folder](#) to all all the cameras to.

2.2.3 Manually adding a Camera

If your video device doesn't support ONVIF, you can add a camera manually to VisualCortex.

On the Landing Page or any [folder](#), simply click **New** and choose Live Camera.



You'll be asked to add a camera name, description, and specify the connection details of the camera.

2.2.4 What's next?

Once you have configured your camera, you can add [models](#) to it, and set up [alerting](#) and [dashboards](#).

2.3 Archived Footage

While VisualCortex works well with Live video sources, it also has first class support for Archived Footage. This essentially means video files.

Video files can come from a number of sources. These include:

- Uploading the video directly into VisualCortex
- Importing from one of the supported [integrations](#)
- Storage available to an [edge cluster](#)

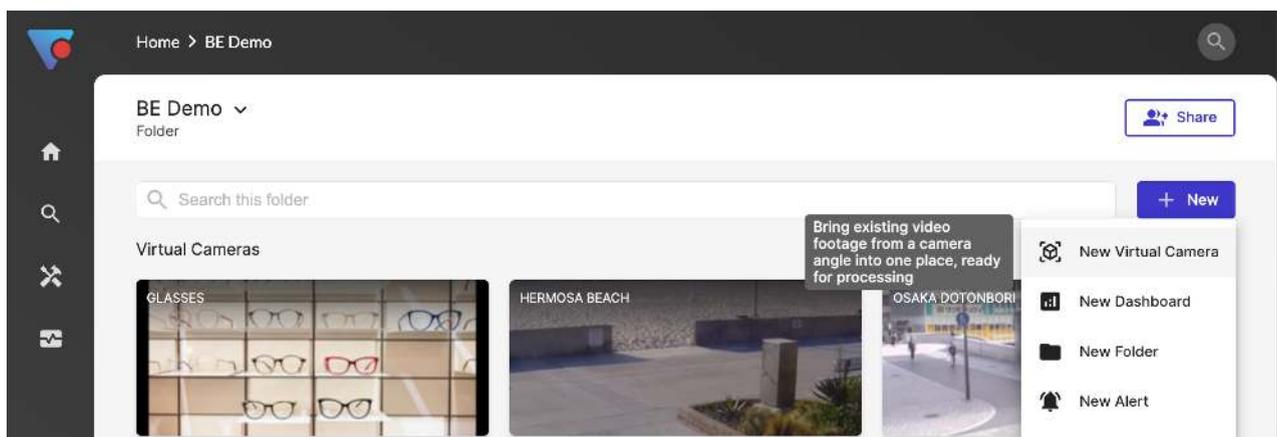
2.3.1 Virtual Cameras

Live footage belongs to a *virtual camera*. Think of a virtual camera the same way as a live camera, but instead of live video, it uses previously recorded footage.

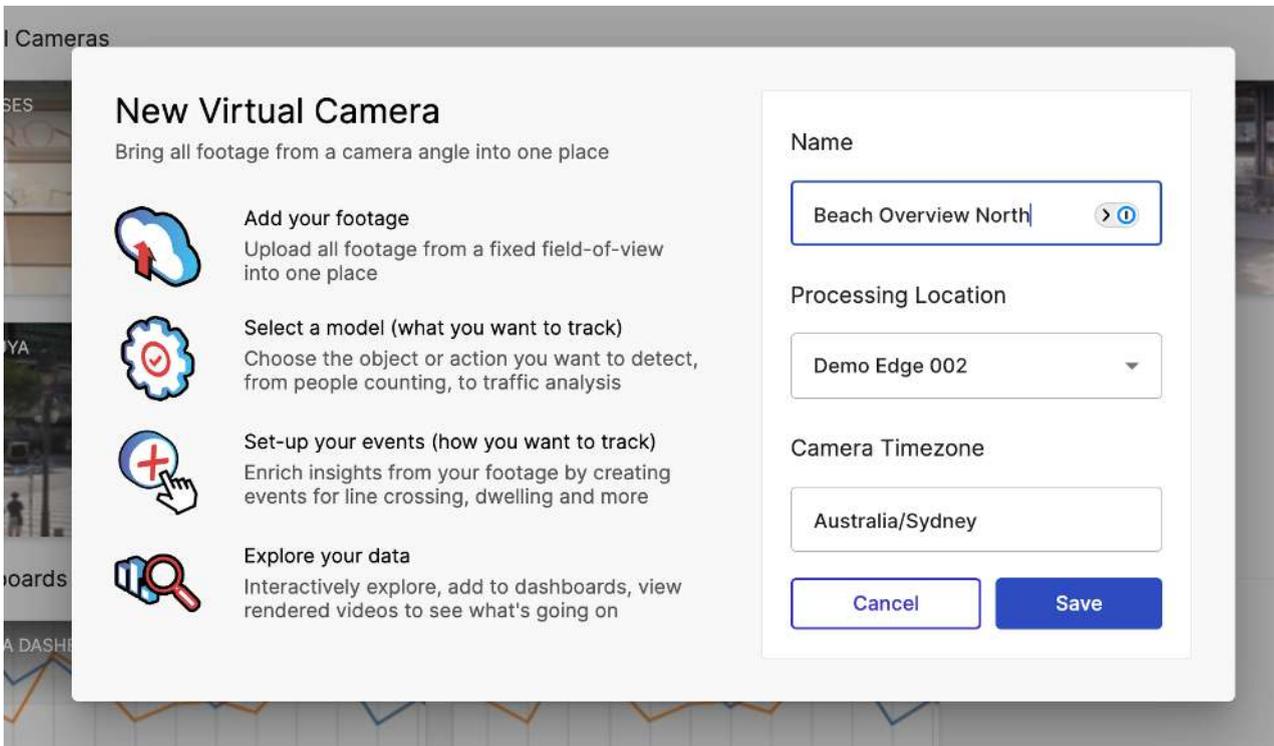
If you've got footage physically recorded from different cameras, they belong in different virtual cameras. If you've got multiple video files from the same camera, keep them in the same virtual camera.

2.3.1.1 Creating Virtual Cameras

On the Landing Page or any [folder](#), simply click **New** and choose Virtual Camera.

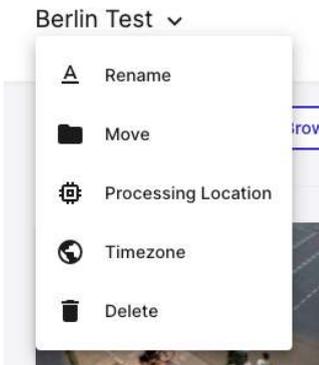


You'll be asked for information about the virtual camera, including a name, a timezone and a processing location. This processing location is an [edge cluster](#). Any [models](#) you add to this virtual camera will be processed on the selected cluster. You can change clusters after the virtual camera has been created.



Once you've created the virtual camera, you'll be prompted to either upload or import footage.

If at any time you wish to change the [Edge Cluster](#) which processes your footage, simply change the processing location on the Virtual Camera.



2.3.2 Technical Requirements

VisualCortex supports a wide range of video files. In terms of container formats, we support:

- MP4 / M4V
- MKV
- AVI
- MPEG Transport stream (TS)

In terms of codecs, VisualCortex supports:

- MPEG-2 Video
- H.264
- H.265 / HEVC
- AV1
- VP9

If your video format is unsupported, you can [Contact Us](#) and we may be able to add support.

2.3.3 What's next?

Once you have configured your camera, you can add [models](#) to it, and set up some [dashboards](#).

3. Concepts

Learn about key VisualCortex Concepts.

Architecture

VisualCortex is a highly scalable, Enterprise grade application.

Alerts

Introduction

Credits

Every installation of VisualCortex comes with Credits. Credits are needed to process Virtual Cam...

Dashboards

While VisualCortex provides comprehensive Integrations for data analysis using third party Busin...

Edge Clusters

Edge clusters (also referred to as Edge Boxes or Edge Devices) are where VisualCortex runs its C...

Event Definitions

By themselves, Computer Vision models output metadata about what has been detected. While s...

Folders

VisualCortex uses the concept of folders for organisation. While not strictly required to use Visual...

Models

At a high level, VisualCortex allows you to configure, run and monitor Computer Vision Models.

Permissions

Folder Permissions

3.1 Architecture

VisualCortex is a highly scalable, Enterprise grade application.

The application itself can be run in the following configurations:

- Utilising a Cloud Provider (AWS, GCP, Azure etc)
- Utilising a Virtual or Physical Machine on-premises
- On-premises in an Air-gapped environment

3.1.1 Concepts

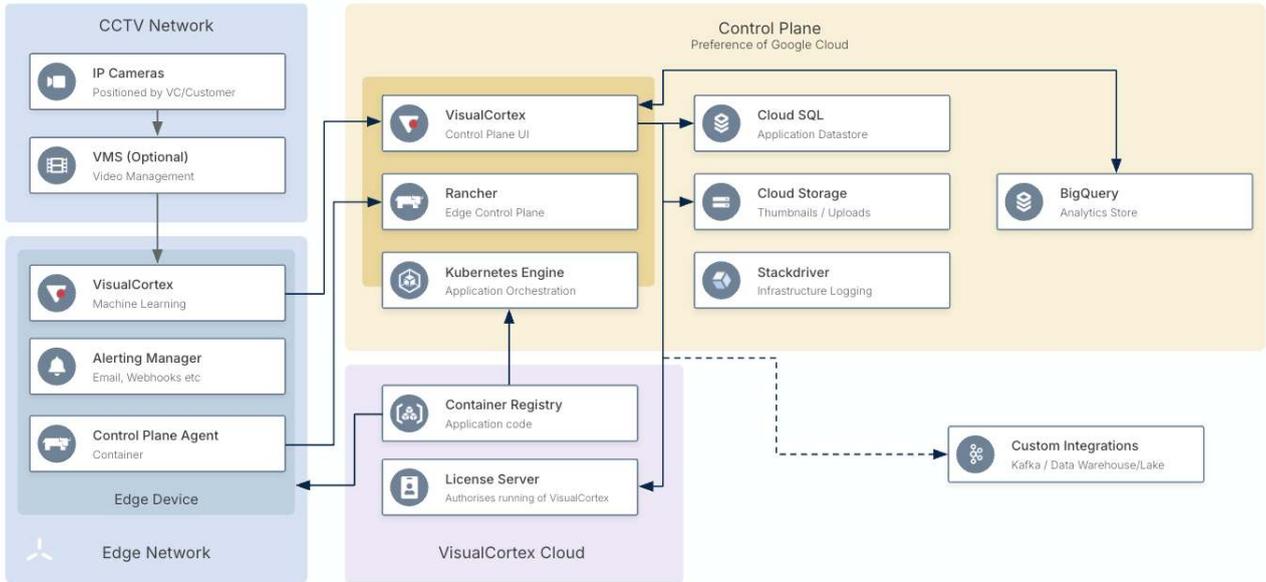
VisualCortex is made up of two main deployable components. These are:

- The Control Plane, which houses the platform configuration and reporting interface
- [Edge Clusters](#)

The platform interface has to be accessible by [Edge Clusters](#) and users wishing to interact with the platform. This *typically* means hosting this component in a central location accessible to users - the most common being a cloud tenancy.

[Edge Clusters](#) are more likely to be in more protected environments - either part of an existing CCTV network, or at least able to connect to any required [live sources](#). Edge clusters also need to be able to communicate **outbound** to the platform configuration and reporting interface; for most customers, this is achieved with outbound internet access. Edge Clusters can successfully work behind NAT Gateways, through VPNs, and can also work with Corporate Proxies, provided that those proxies have support for [Web sockets](#).

An overview of how different components interact in a Google Cloud environment is below. Differing cloud vendors utilise different components with similar functionality.

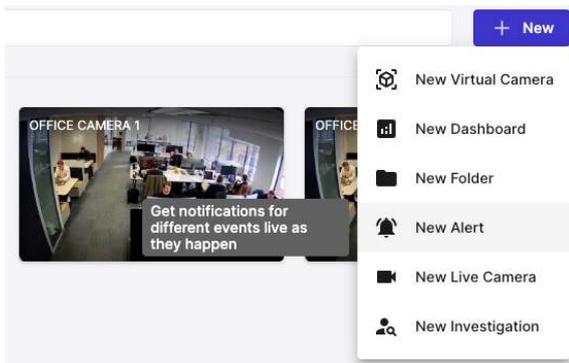


3.2 Alerts

3.2.1 Introduction

VisualCortex has a real-time alerting capability which allows you to get alerts related to [Live Cameras](#). Alerting is specified in a natural language style, and is very easy to configure. Most alerts work in conjunction with [Event Definitions](#).

You can create alerts in [Folders](#) by clicking the **New** button and selecting *Alert*.



Alerts themselves can be *Active* or *Disabled*, and can have multiple *Times* in which they can fire (eg. Days or times during the week). Alerts also have the concept of a *Recovery Time*, which ensures you don't get alert fatigue when multiple alerts are firing unnecessarily.

Example: If you set up an exclusion zone where *"More than 1 person is in the zone for more than 5 seconds"*, if a person moves *out* of the zone and then *back into* the zone within the defined *Recovery Time* time period, a separate alert will *not* trigger.

Queue time ▾

Alert

Share

Rule **Active** ▾Across **1 dwell area** ▾Alert if **1** or more objects are in area for longer than **5 seconds** ▾Timing **+**

Australia/Sydney

No time filters, this can fire at all times

Notification Targets **+**

Notify once

Email Ben
EmailWebhook from the...
WebhookMilestone
MilestoneWebhook from the...
WebhookSMS Ben
SMS

Alerts

Recovery Time: 30 seconds

May 8 2023 - 3:55:06 PM



Ended after 59 seconds

May 8 2023 - 3:33:36 PM



Ended after 20 minutes 30 seconds

May 8 2023 - 3:09:06 PM



Ended after 24 minutes

Every alert triggered has an *Alert Log*. This log shows you when an alert started, what endpoints were notified successfully (or unsuccessfully), and when the alert was resolved.

Alert Log		
	Alert started	5th May 2023 - 03:55:06
	Email Ben	03:55:07
	Webhook from the Cloud	03:55:07
	SMS Ben	03:55:07
	Milestone	03:55:07
	Webhook from the Edge - Unknown Error	03:55:12
	Alert ended	5th May 2023 - 03:56:05

3.2.2 Alert Types

VisualCortex has multiple alert types for a variety of use cases.

- **Queue Time** - This allows you to alert if an object has been in a [Zone](#) for a given period of time. Can also be used for malicious dwelling, safety use cases and intrusion detection
- **Simple Line Cross** - This allows you to alert if an object has crossed a [Directional Line](#)
- **Line Cross Capacity** - This allows you alert on a running total across multiple [Directional Lines](#), which can optionally be across multiple cameras. Typically used to see if an area is at capacity.
- **Average Queue Length** - While *Queue Time* alerts work great for a single instance of an object (eg. 1 person standing in a queue for > 5 minutes), Average queue time allows you to make alerts for *average* queue times (eg. The average wait time is over 5 minutes).

3.2.3 Notification Targets

VisualCortex supports several different alert notification targets. These are outlined below:

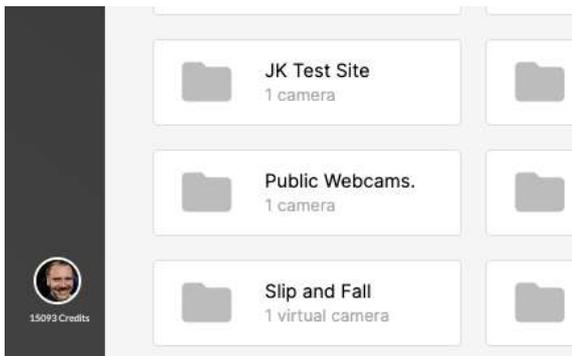
- **Email** - This sends an email to the specified address. The [Control Plane](#) must have internet access. Emails are sent through VisualCortex servers.
- **SMS** - This sends an SMS to the specified Phone number. The [Control Plane](#) must have internet access. SMS's are sent through VisualCortex servers.
- **Webhook (Control Plane)** - This sends a [Webhook](#) from the [Control Plane](#)
- **Webhook (Edge)** - This sends a [Webhook](#) from the [Edge Cluster](#). This type of alert can survive temporary loss of connectivity with the VisualCortex Control Plane.
- **Milestone** - This sends an analytic event to a [Milestone VMS](#) instance. This type of alert can survive temporary loss of connectivity with the VisualCortex Control Plane.
- **Geutebruck** - This sends an analytic event to a [Geutebruck VMS](#) instance. This type of alert can survive temporary loss of connectivity with the VisualCortex Control Plane.

3.3 Credits

Every installation of VisualCortex comes with Credits. Credits are needed to process [Virtual Camera footage](#). Different Computer Vision Models require different amounts of credits. The total amount of credits required for processing Virtual Camera footage depends on the models themselves, and the length of the video footage being processed.

Credits are not required for Live cameras.

You can see your amount of credits in the bottom left hand side of the screen at any time.



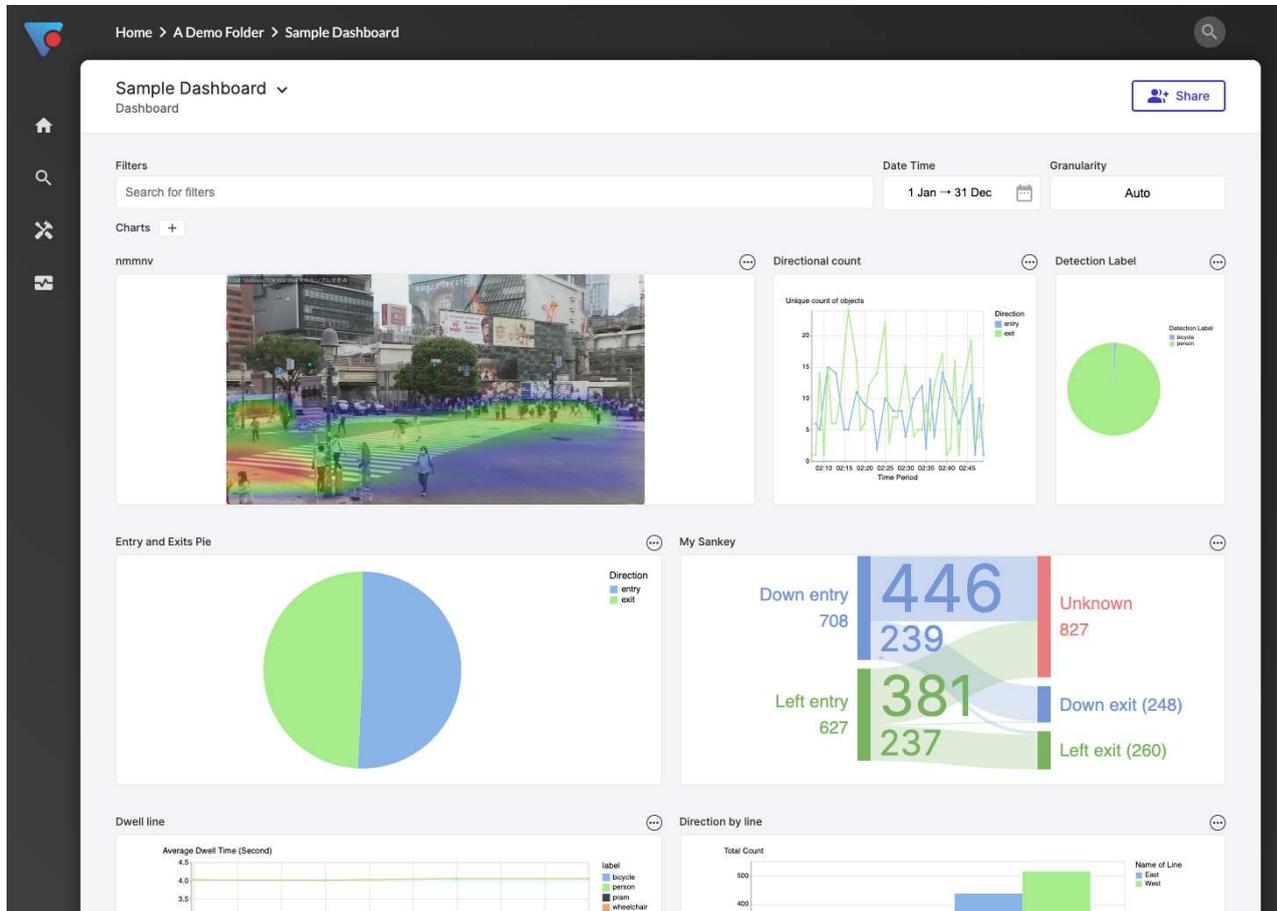
Credits are deducted immediately when submitting Virtual Camera Footage jobs. If the job fails for any reason, credits will be automatically refunded.

3.3.1 Obtaining more credits

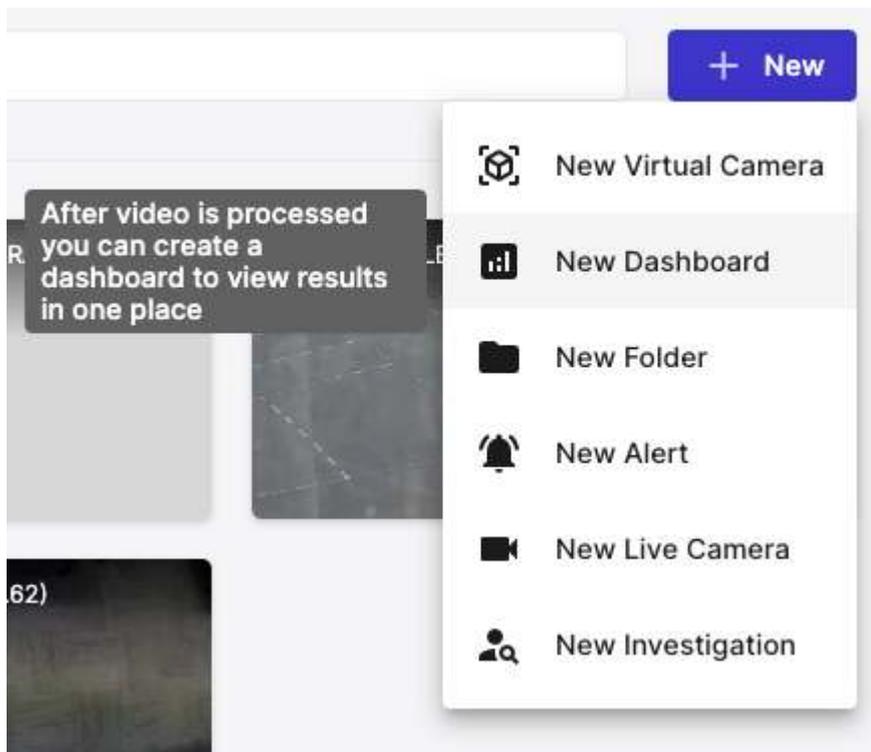
To add credits to your installation, reach out to [VisualCortex Support](#).

3.4 Dashboards

While VisualCortex provides comprehensive [Integrations](#) for data analysis using third party [Business Intelligence](#) tools, there is an ability to create Dashboards in the platform.



To create a dashboard, simply click **New** in a folder, and select **New Dashboard**. Dashboards can have their own [Permissions](#), so you can grant access to them without granting access to the underlying cameras.



Dashboards work with both [Live Cameras](#) and [Virtual Cameras](#). If data is present in the VisualCortex Control Plane analytic data store, it can be reported on here.

3.4.1 Creating charts

Dashboards are made up of *Charts*. You can add charts by clicking on the **[+]** icon in the Dashboard itself. You'll be presented with a chart selection page similar to this:

1 Select Chart 2 Data Sources 3 Configure

Camera

Heatmap



Directional Count

Line



Table



Directional Flow Sankey



Pie chart



Grouped bar chart



Dwell

Line



Dwell stacked bar chart

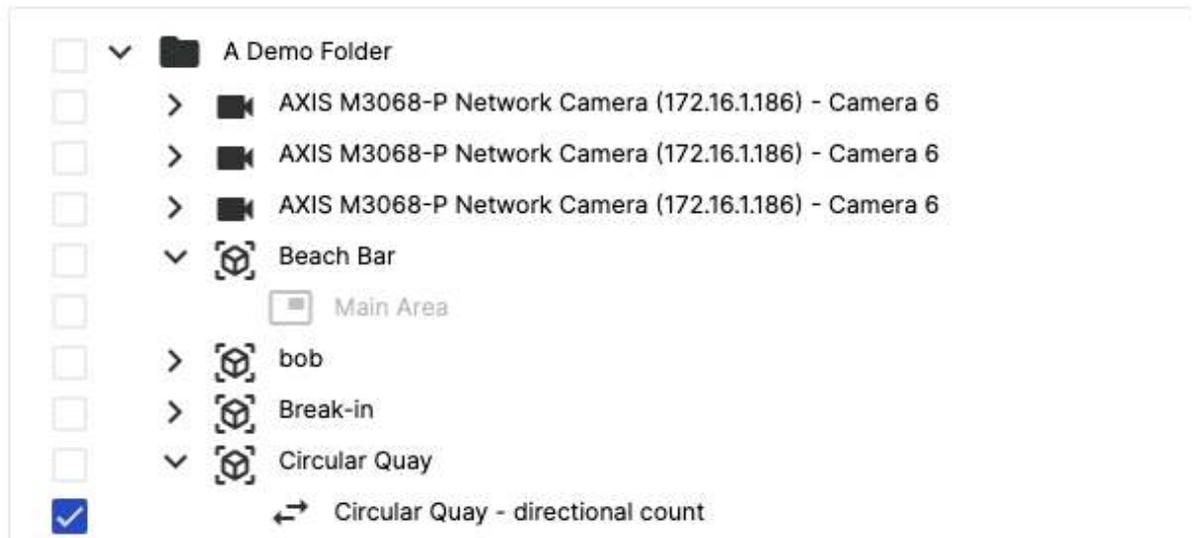


Histogram Heatmap chart



Choose a chart suitable for what you're wanting to report on. Every chart will ask you to select a data source. This is typically one or more [Event Definitions](#), but could also be a camera.

1 Select Chart 2 Data Sources 3 Configure



Lastly, you'll get the chance to configure any options that a specific chart needs. This differs per type of chart.

1 Select Chart 2 Data Sources 3 Configure

Options

Give your chart a name *

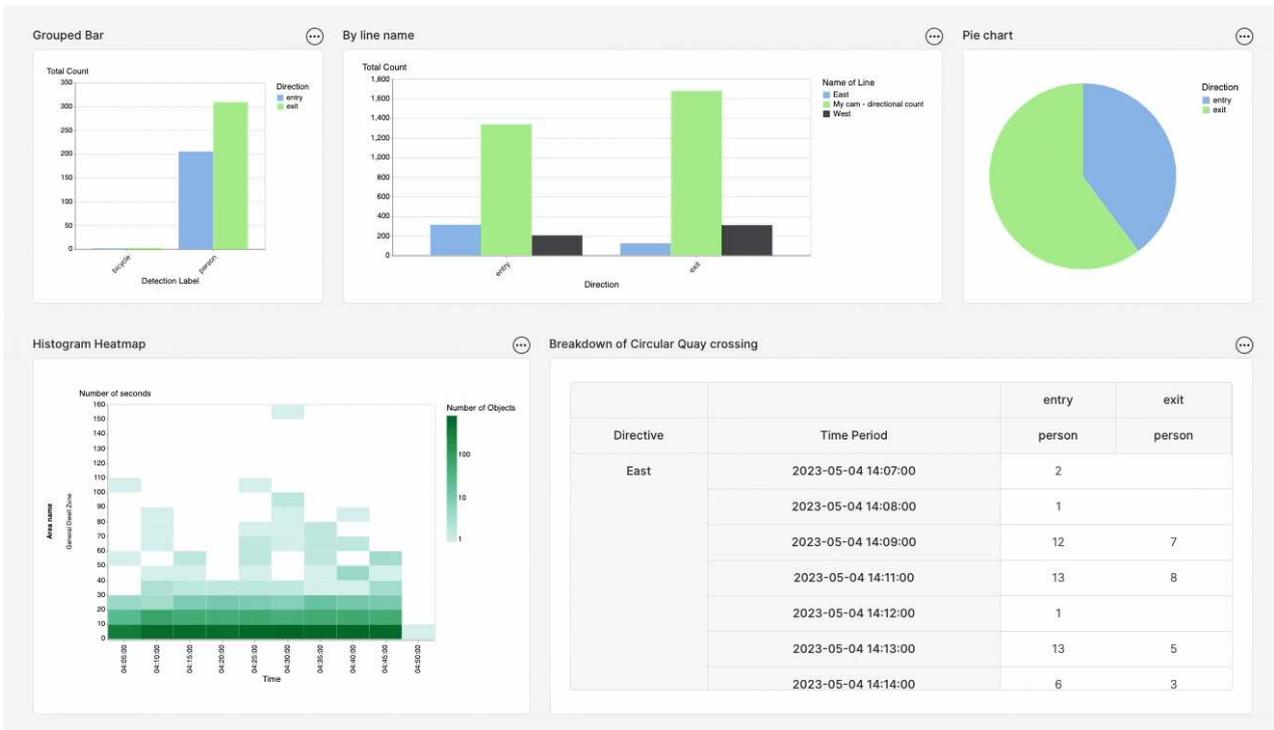
Select pivot options *

Duplicate behavior *

Filters

Back

After charts have been created, you can easily move and resize them. Move by dragging the chart. Resize by dragging the resize icon in the bottom right-hand corner of the chart.



You can filter an entire dashboard by using the filter controls at the top of the Dashboard. There is also a **Granularity** drop-down, which allows you to specify how granular you want your reports (eg. Minutely, hourly, daily etc.)

3.5 Edge Clusters

Edge clusters (also referred to as *Edge Boxes* or *Edge Devices*) are where VisualCortex runs its [Computer Vision Models](#). Edge clusters can process both Archived Footage from [Virtual Cameras](#) and footage from [Live Cameras](#).

Edge Clusters can either belong to a single [Organisation](#), or be shared across Organisations.

Every [Live Camera](#) or [Virtual Camera](#) in VisualCortex is allocated to an Edge Cluster. VisualCortex will automatically distribute live cameras across the available GPUs in the Edge Cluster.

For more information on administering Edge Clusters, see the [Administration](#) section.

3.6 Event Definitions

By themselves, Computer Vision models output metadata about what has been detected. While somewhat useful, it's difficult to make business decisions (either through analytics or via alerting) without *grouping* or *classifying* this metadata. This is what Event Definitions do.

Event Definitions can be applied to both [Live Cameras](#) and [Virtual Cameras](#). You'll find Event Definitions on both. Simply click the **[+]** icon.

Kabukicho Live Channel II ▾
Virtual Camera

[+ Add Video](#) [Browse 1 Video](#) [Process Video](#) [Explore results](#)

Camera info



- ☰ 1 complete
- 🎬 1 video uploaded
- ⊕ Created one year ago

Event Definitions **+** 



3.6.1 Directional Count

Directional counts are useful for understanding if objects are crossing a certain threshold. It's useful for use cases like people/vehicle counting and intrusion zones.

Directional counts can have multiple points - simply click the point icon, or double click on the image to add another point. They also have directions, and both *Entry* and *Exit* directions can have additional tags if desired.

Edit directional count

Count objects that cross this line with the direction of travel



* Double click to add a new point

Name

Entry names

Exit names

3.6.2 Zones

Zones are useful for understanding the presence of an object in an area, or the length of time in which an object stays in an area. Zones fulfil a variety of use cases including safety, queue management, intrusion detection, weapon detection etc.

Like Directional counts, zones can have multiple points - simply click the point icon, or double click on the image to add another point.

Edit zone

Count objects in the zone, along with their duration within it



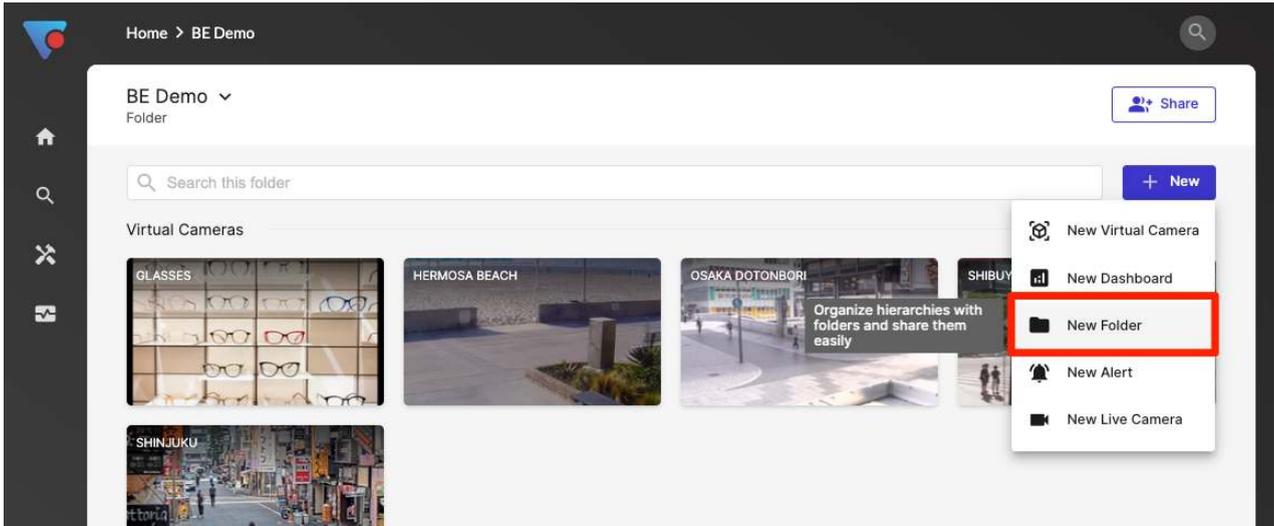
* Double click to add a new point

Name

3.7 Folders

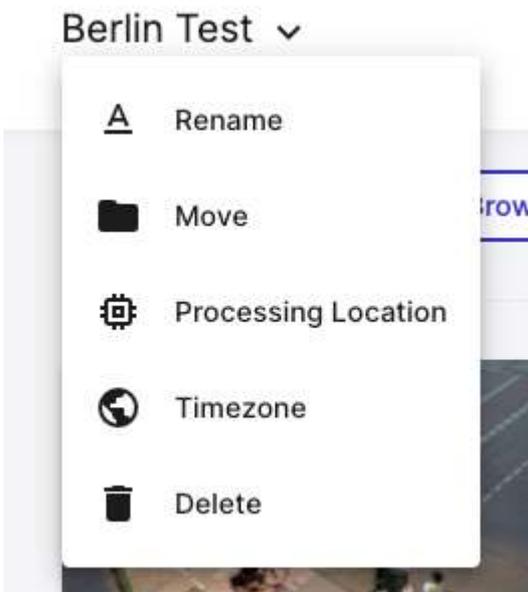
VisualCortex uses the concept of folders for organisation. While not strictly required to use VisualCortex, Folders facilitate easy grouping of entities like [Live](#) or [Virtual Cameras](#), [Dashboards](#) and [Alerts](#).

To create a folder, click on the **New** button, then **New Folder**. You can create folders on the main landing page, or inside other folders.

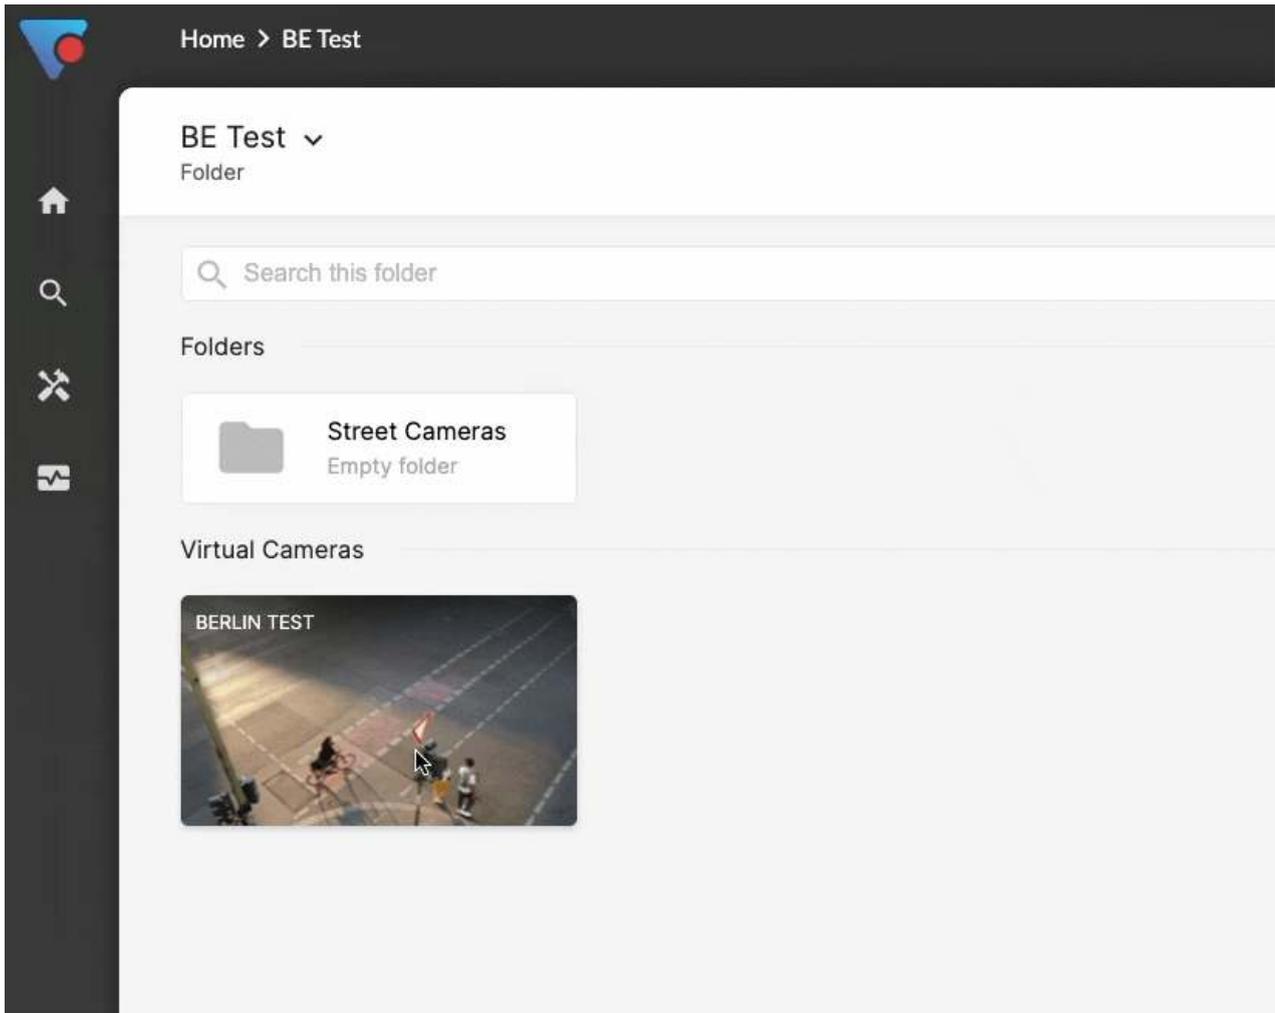


Folders are also easily shared; For more information, please see the section on [Permissions](#).

You can move any entity in a folder to another folder easily. This can either be done in the entity settings, or by clicking **Move**. You'll be presented with a dialog asking you where you'd like to move the folder.



You can also drag and drop tiles from the folder browser into other folders.

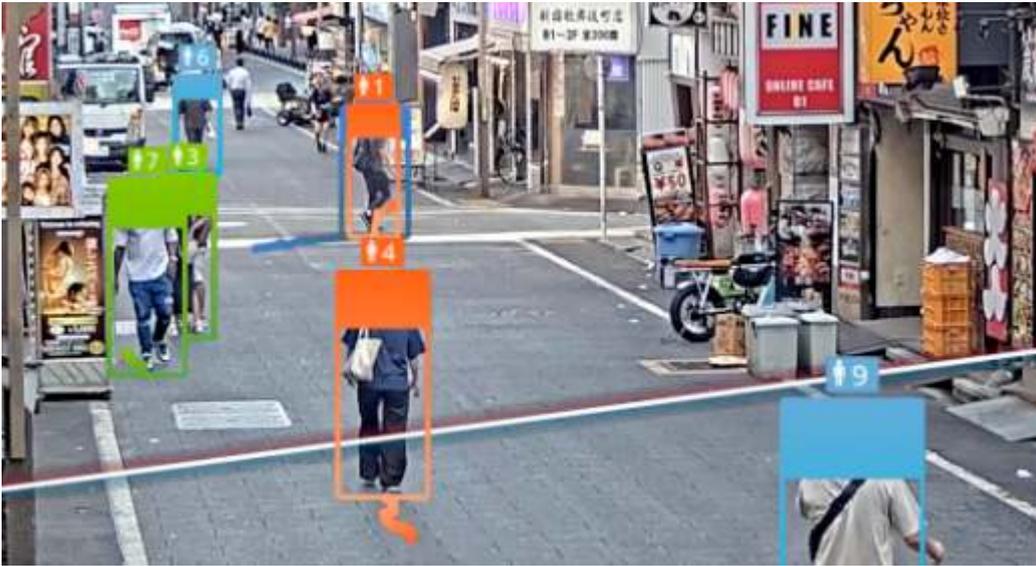


3.8 Models

At a high level, VisualCortex allows you to configure, run and monitor Computer Vision *Models*.

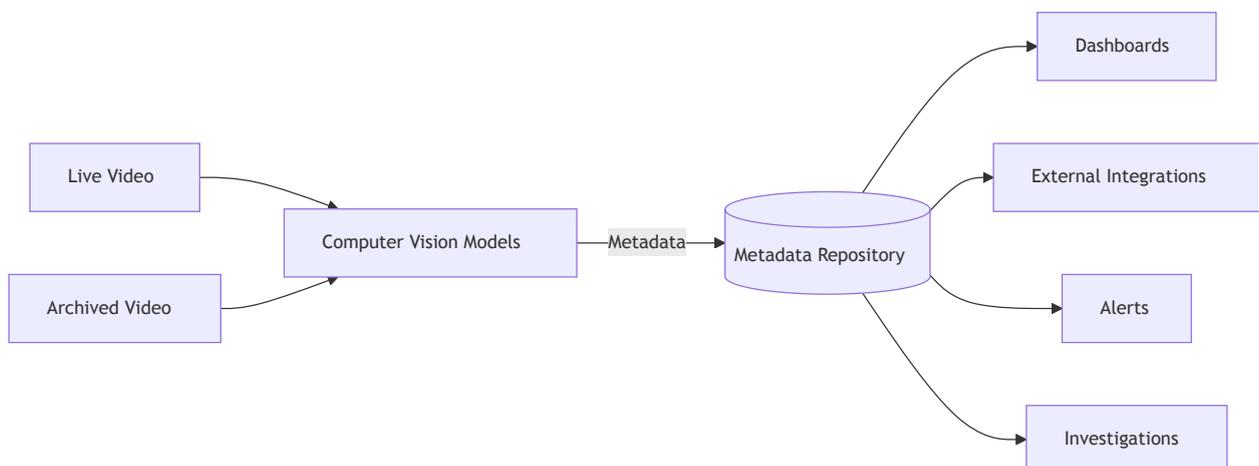
Computer vision models can come from a variety of sources. These include:

- Models part of VisualCortex
- Extra models created by VisualCortex
- Models developed by VisualCortex Partners
- Private models developed users of VisualCortex



Above: A sample visualisation of VisualCortex's *Pedestrian* model, which comes with VisualCortex.

Computer vision models take in video feeds and produce metadata. This metadata can then be used for realtime operations (integrations, alerts etc.), or for reporting.



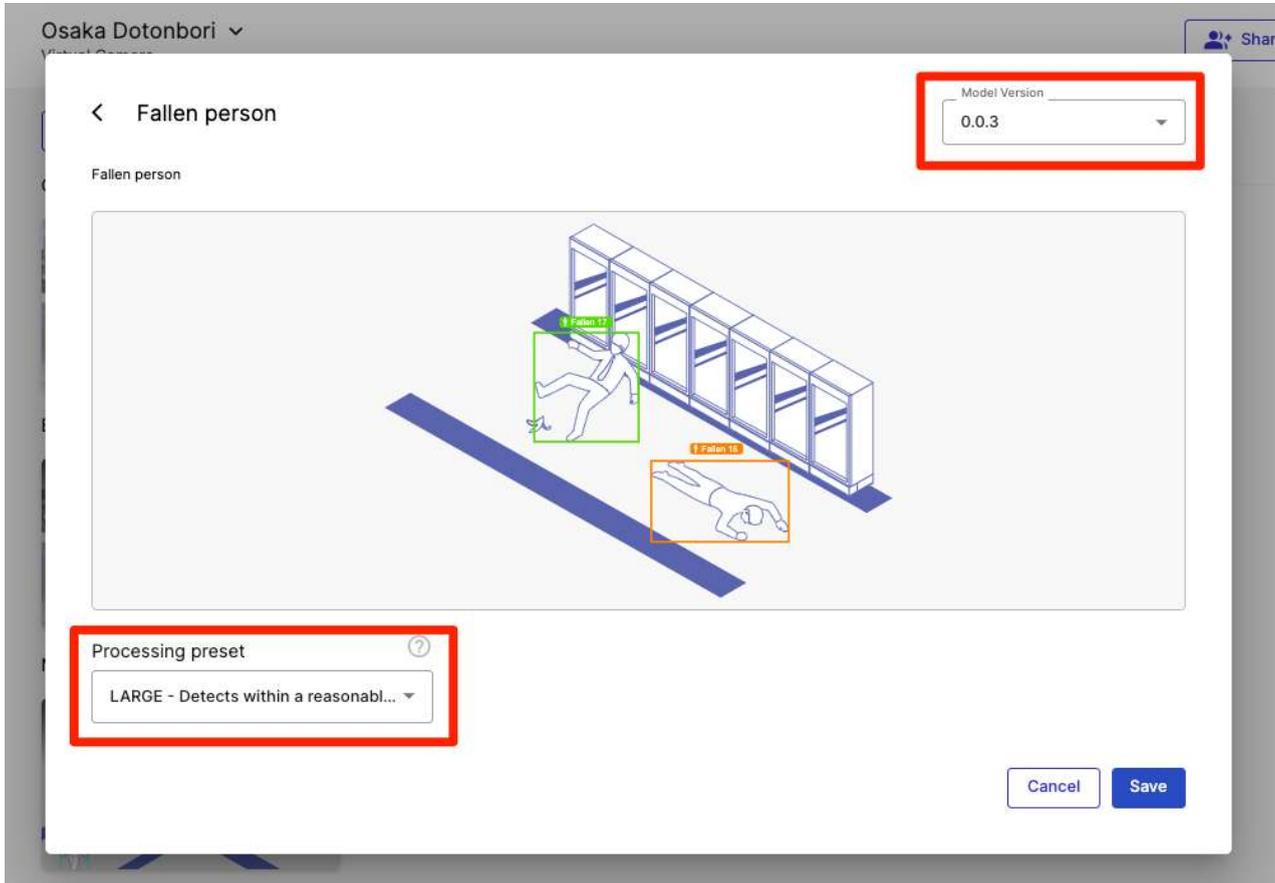
3.8.1 Applying models to video

Generally speaking, adding a Computer Vision model to a [Virtual Camera](#) or a [Live Camera](#) works the same way. Some models are applicable *only* to Live sources, while others are applicable *only* to Virtual sources, but generally speaking the same models work with both.

In the Camera overview, you'll be able to see the models which are currently running (or set to run) for that camera - whether or not it's live or virtual. To add a model, simply click the **[+]** icon next to models in the Camera overview section.

Different models have different settings. Some may require you to define a zone, provide metadata, confidence levels etc. Others have no settings - it depends on the requirements of the model selected.

Every model also has a **Version**, which is available in the top right-hand side of the model definition dialog, as well as **Processing Preset**. Processing presets allow you to select the right *optimisation* of the model for your given camera. Generally speaking, this affects internal settings that the model author has defined; This typically involves settings such as *frames per second* and *model resolution*.

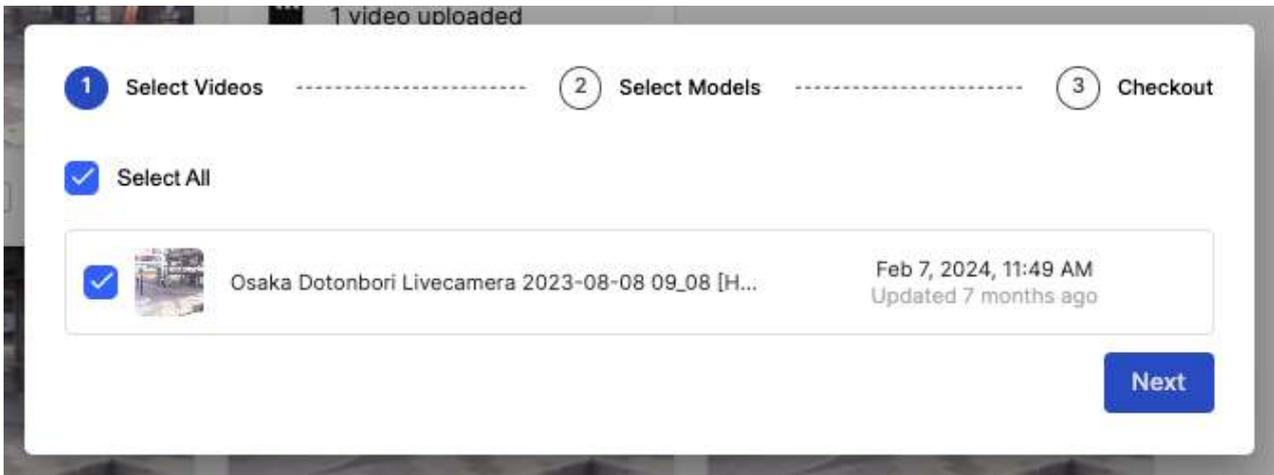


Once a model has been added to a camera, it's either *available for processing* (if the camera is a [Virtual Camera](#)), or immediately applied if it's a [Live Camera](#).

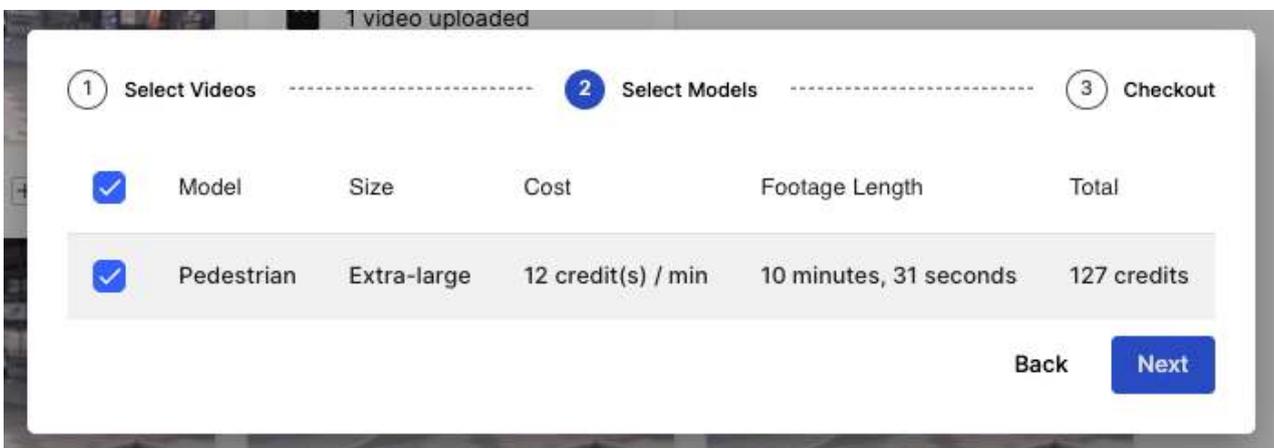
3.8.2 Processing on Virtual Cameras

Once a computer vision model is applied to a [Virtual Camera](#), you can elect to *Process* the footage associated with that camera. Simply click **Process** in the Virtual Camera overview screen, and walk through the wizard.

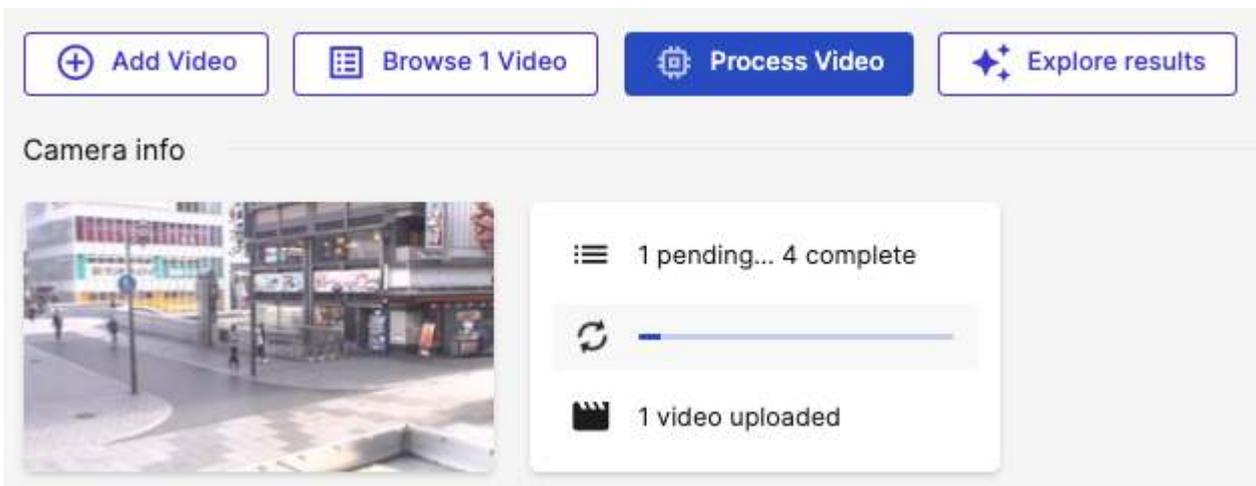
You'll first be prompted to choose the footage you want to process.



Next, you'll be asked which models you want to process on this footage. Each model takes a certain amount of [Credits](#) per second.



Once you've confirmed your selection, VisualCortex will schedule the processing on the [Edge Cluster](#) associated with your Virtual Camera. You'll be presented with a progress bar, and notified via email once the processing has been completed.



Once complete, you'll be able to view data in [Dashboards](#) and *Visualise* small portions of your video.

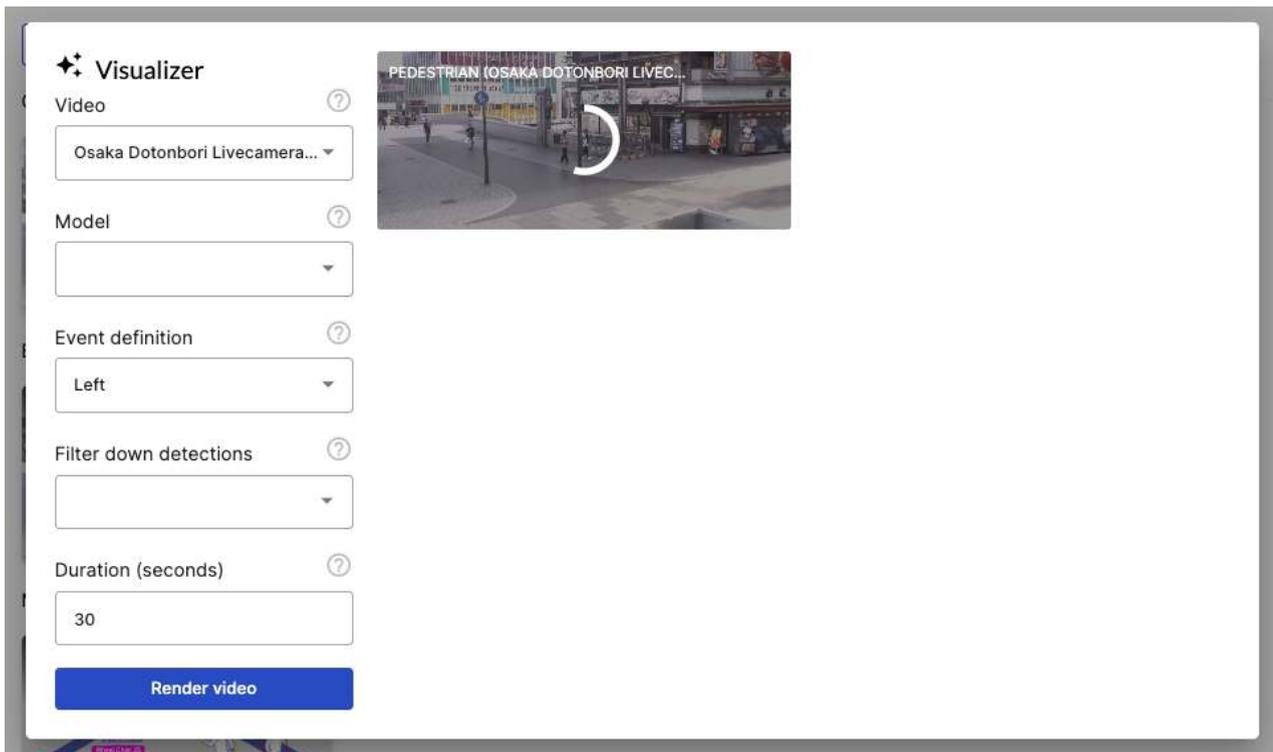
3.8.2.1 Visualising video

Sometimes it's helpful to *visualise* the metadata coming out of a computer vision model. This is particularly useful if you're trying to understand how well the model is performing or whether or not the *preset* is appropriate for the video.

VisualCortex allows you to visualise small segments of Virtual Camera Footage. In order to do this, you'll need to have:

- Footage added to the Virtual Camera
- Processed the footage through a Computer Vision Model
- Added in at least one [Event Definition](#)

To create a visualisation, simply press the **Explore Results** button. You'll then be asked to choose a Model, [Event Definition](#) and a duration from 1-600 seconds. Once you've selected what you need, press **Render Video**. The video render job will be scheduled on the Virtual Camera's [Edge Cluster](#) and you'll see the progress in the main window.



Once visualisation is complete, you can click on the thumbnail to view the appropriate video. Visualisation videos persist for only 5 minutes. If you want to save the video for later, simply choose the option in the video player to download.

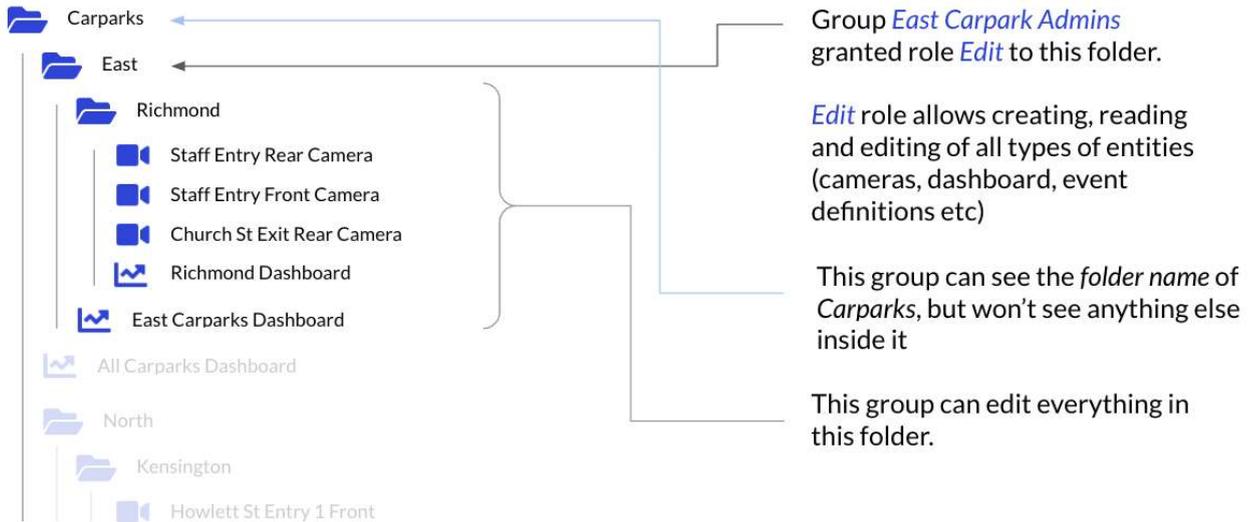
➕ Add Video | 📄 Browse 1 Video | 🎬 Process Video | 🔍 Explore results

The screenshot displays the VisualCortex web application interface. At the top, there is a navigation bar with four buttons: "Add Video", "Browse 1 Video", "Process Video", and "Explore results". The main area is a video player showing a street scene with several people walking. Blue bounding boxes are overlaid on the video, each labeled with a person count and "person", such as "6 - person", "5 - person", "8 - person", and "3 - person". In the bottom left corner of the video player, there is a pause icon and a progress indicator showing "0:09 / 0:29". On the right side, there is a sidebar with the "VISUALCORTEX" logo at the top. Below the logo are two buttons: "Entry" (blue) and "Exit" (red). Underneath these buttons is a counter showing "1" next to a person icon and "0" next to another person icon. At the bottom of the sidebar, there are three menu items: "Download" with a download icon, "Playback speed" with a play icon, and "Picture-in-picture" with a window icon.

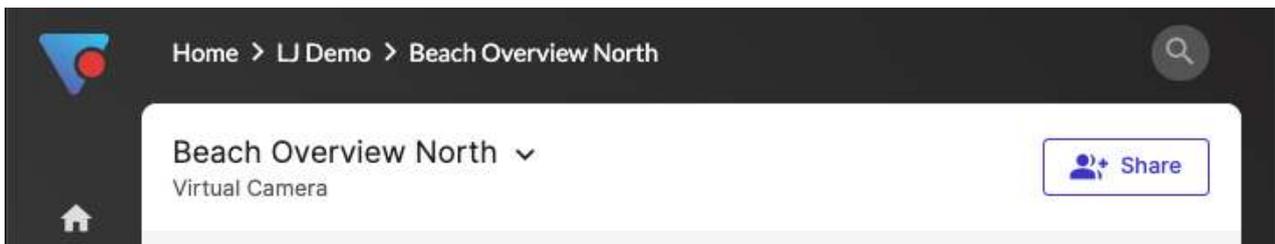
3.9 Permissions

3.9.1 Folder Permissions

VisualCortex utilises a simple *waterfall permissions model*. What this means is that every sub-[folder](#) or entity within a folder will *inherit* the permissions of its parent.



You can grant permissions on a [Folder](#), [Live camera](#), [Virtual Camera](#) or [Dashboard](#) by utilising the **Share** button. This button is displayed on every component within VisualCortex which can be shared.



When sharing, you must share to a [group](#).

3.9.2 Roles

VisualCortex has comprehensive *role* based permissions control (RBAC). When sharing things within VisualCortex, you can choose not only *who* to share with, but what *role* you want the shared group to have. For more information on this, see the [Roles](#) page.

3.9.3 Administrative

In addition to the configurable [Roles](#) specified in an [Organisation](#), VisualCortex has two Administrative roles, which are attached to an organisation [User](#).

These roles are:

- Organisation Administrator
- Super Administrator

3.9.3.1 Organisation Administrator

An Organisation Administrator can view, edit and delete **everything** within an organisation. In addition to this, they can:

- Administer [Edge Clusters](#)
- Administer [Integrations](#)
- Administer [Users](#) within their organisation
- Administer [Groups](#) within their organisation
- Administer [Code Transformations](#) within their organisation

3.9.3.2 Super Administrator

A Super Administrator can do everything an Organisation Administrator can, but across *all* [Organisations](#). In addition to this, Super Administrators can:

- Log into any [Organisation](#) configured in the Platform
- Set up [IP Whitelisting](#)
- Add and administer [Edge Clusters](#) which are *shared* across organisations

4. Investigations

VisualCortex Investigations allow deep analysis and insights into data, giving a full video intelligence experience. The Investigations module makes searching for people, vehicles, faces, and more, easy.

Introduction

An investigation is your one stop shop for footage, reports, and searches when investigating activ...

Investigation Dashboard

Investigations Home Page

Bulk Import

Bulk import allows you upload or import many files at once into an Investigation. Bulk importing c...

Searches

An Investigation Search is the place where users of VisualCortex can look specific people, faces, l...

Report

The Investigation report takes all annotations made in any Investigation Searches, and puts all of ...

4.1 Introduction

An investigation is your one stop shop for footage, reports, and searches when investigating activity. Investigations allow you to filter through thousands of detections across hundreds of hours of video footage quickly and easily.

Work collaboratively with a team, or by yourself. Tag detections and video snippets with commentary, and add to a simple report. Like everything else in VisualCortex, Investigations have [Permissions](#), and every action is strictly audited.

An Investigation acts as a container for Sources, Models, Searches and Reports:



4.1.1 Sources

Sources are your [Video Sources](#), whether they be [Live Video](#) from live cameras, or [Archived footage](#). Easily upload video from your browser, or import from Video Management Systems or [Cloud Storage](#).

4.1.2 Models

Each investigation is associated with multiple [Models](#). Examples might include [Face Detection](#), Person Attributes, License Plate Recognition, Fallen people, Weapons and more. Simply add a model to an investigation in order to ensure that your Sources are processed using those models.

Note: It's helpful to separate out your sources into different investigations with different models. eg. Don't run License Plate Recognition on sources which will only have pedestrian foot traffic.

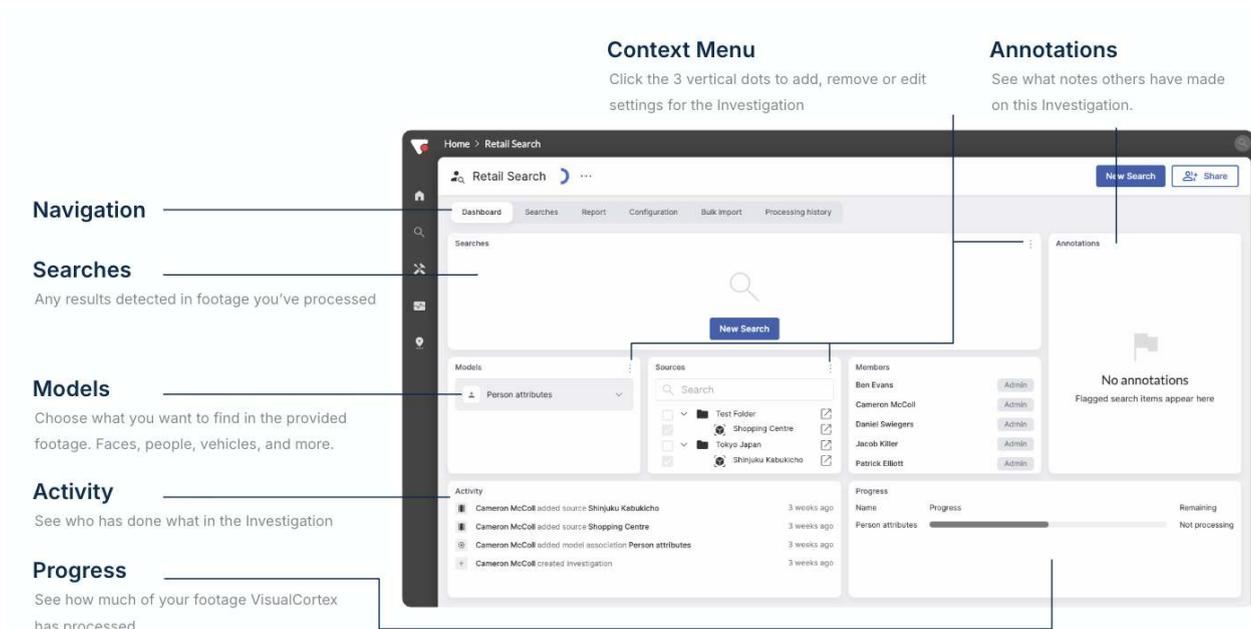
4.1.3 Searches

Searches are a core part of an Investigation. Specify *what* you're looking for in your sources, and easily filter through thousands of results. The Search page is very powerful, allowing you to view thumbnails, video snippets, detected objects in context and more. Filter by sources, Geolocation, tags and more. Add in annotations to prepare your report. Searches are also collaborative, allowing multiple people to search at the same time in the same Investigation.

4.1.4 Reports

Reports are ultimately what get produced as a result of performing an investigation. Reports show all annotated detections in an investigation made by users. Free-form notes can also be added, so that a comprehensive timeline of events can be built up. See maps, detections, video snippets and more. Easily export to PDF.

4.2 Investigation Dashboard



4.2.1 Navigation

Across the top of an investigation, there is a navigation bar, which has the items:

- Dashboard - The main home page for the Investigation
- Searches - A more detailed view of all Investigation searches
- Configuration - A more detailed view of all sources, as well as models
- Report - The Reporting Page for the Investigation
- Bulk Import - An easy way to import footage into the Investigation
- Processing History - A list of all jobs scheduled for all pieces of footage in the investigation

4.2.2 Sources

Sources are [Video Sources](#), whether they be [Live Video](#) from live cameras, or [Archived footage](#). With Archived Footage, you can either [Bulk Upload](#) footage, or add an existing [Virtual Camera](#).

Note: Any Archived Footage added directly into an Investigation will have the same permissions as the Investigation itself.

4.2.3 Models

The models panel shows you all models which are currently enabled on this Investigation. If you add a model, it will process this model against *all* video in the Investigation Sources. For [Live cameras](#) contained in the Investigation, adding a Model to the Investigation will add the same Model to the Live camera itself.

4.2.4 Activity

All interaction in an Investigation is logged. Any user in the system who changes any part of an Investigation is logged here in the Activity panel. Items in the activity panel are clickable and the links will take you to the relevant part of VisualCortex.

4.2.5 Annotations

Annotations are the basis of your [Report](#). Any time a user adds, edits or removes an annotation on a detection, or adds a note to a report, the log of that action appears in the annotation box.

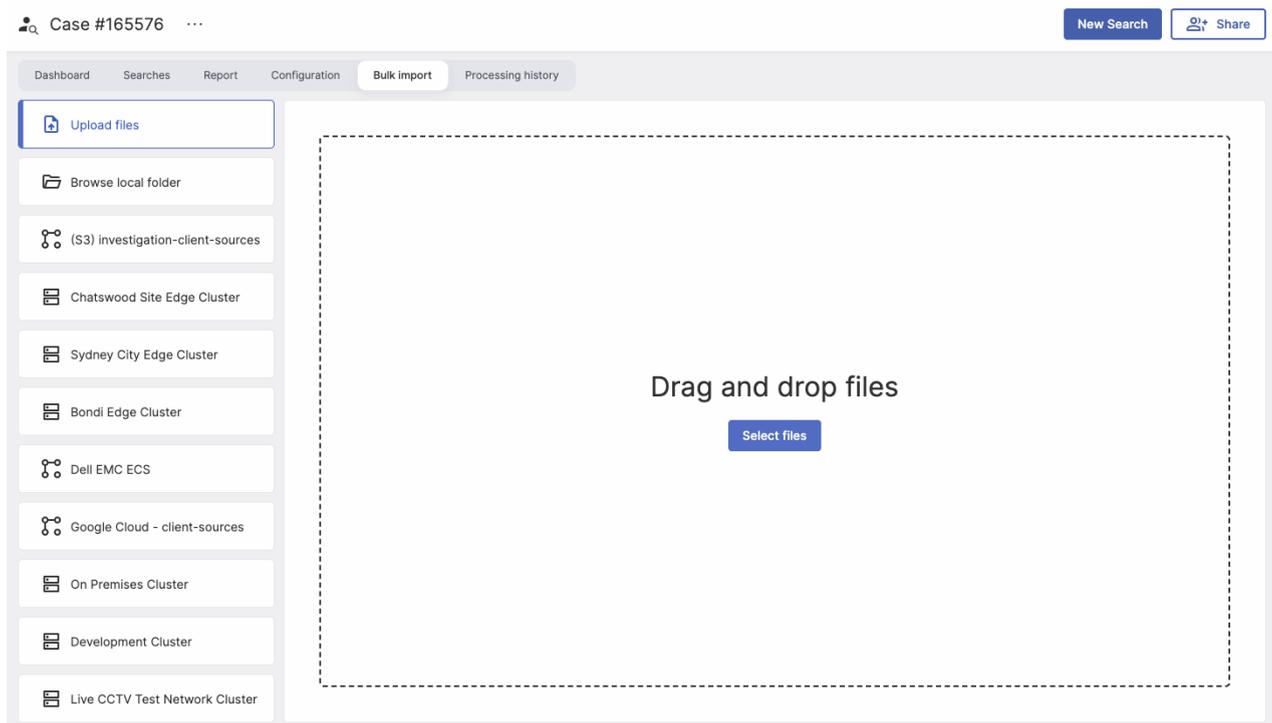
4.3 Bulk Import

Bulk import allows you upload or import many files at once into an Investigation. Bulk importing can work with:

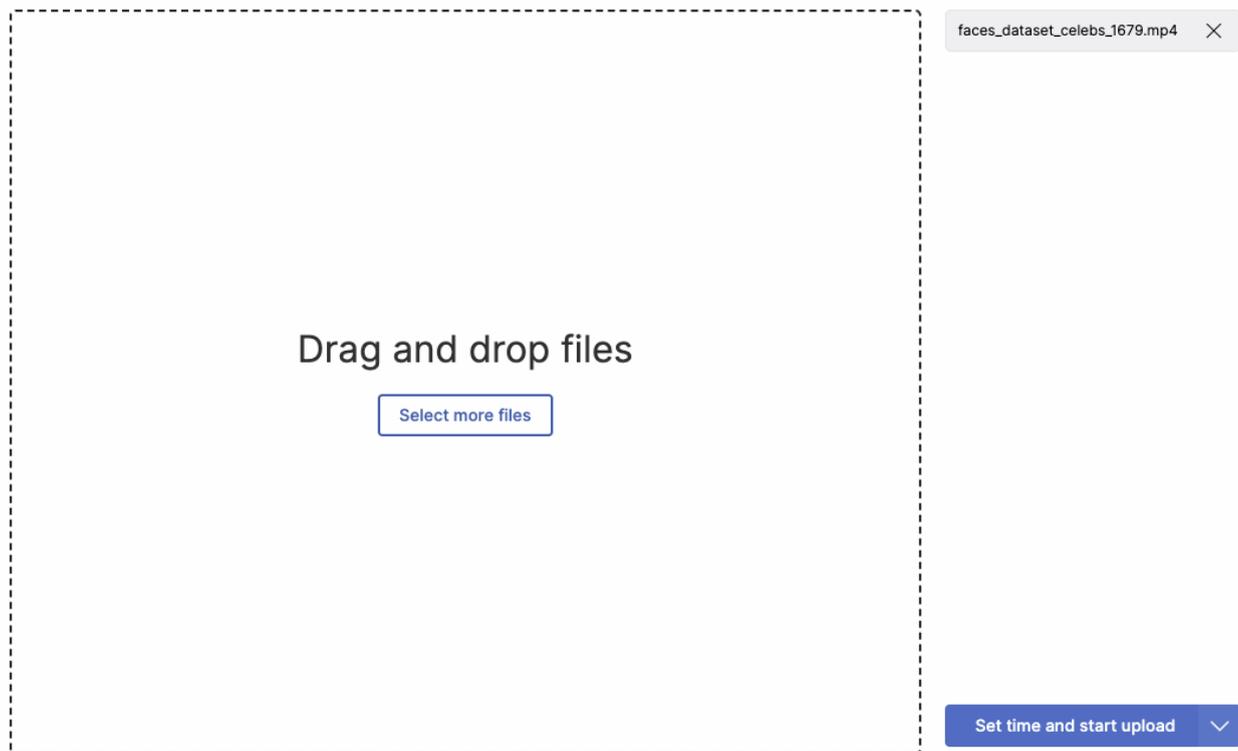
- Files and folders uploaded from a Browser
- Files and folders present in a [Video Integration](#) (eg S3, Google Cloud Storage)
- Files and folders present on an [Edge Cluster](#)
 - Files should be present in the `/media` directory on all nodes

4.3.1 Selecting or importing

The first stage of bulk importing is to actually give VisualCortex access to the relevant files. The easiest and recommended way of getting video data into VisualCortex is to Upload files or folders from your web browser. Simply drag and drop files into the dotted area on the screen, or choose *Select Files* to select files using a file dialog box.

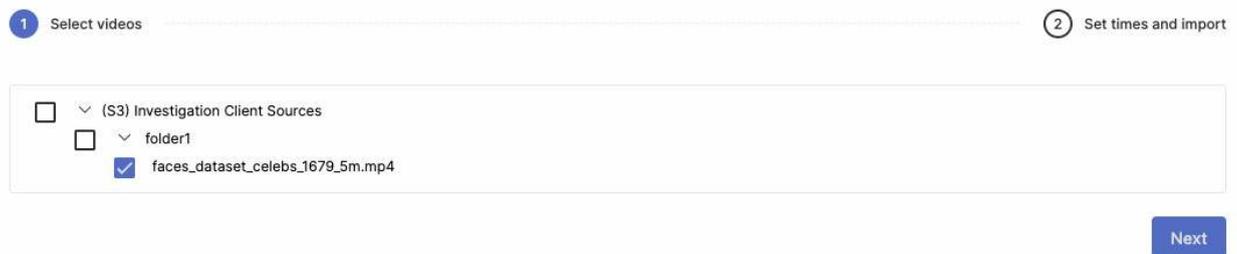


Once you've selected your files, they'll appear on the right hand side of the pane. Add as many as you like, and when you're done, click *Set time and start upload*.



4.3.1.1 Importing

Importing from other sources follows almost the same flow, except you'll get a list of all files and folders that VisualCortex can see in that source. Simply select the relevant assets, and click *Next*.



4.3.2 Setting Video Information

In order for Investigations to give accurate information about Archived Footage, VisualCortex needs to know some details. This includes:

- Where the footage came from (Which Camera)
 - Optionally, the GeoLocation of this camera
- The Date and Time that this footage was recorded at (including the Timezone)
- The [Edge Cluster](#) where you want to process this footage

VisualCortex will automatically try to determine which pieces of footage are from the same camera. It does this by looking at Folder Paths, as well as Filenames. You can change the "grouping" of footage using the on-screen prompts if desired.

Additionally, VisualCortex will also try to automatically determine the Date and Time which the footage was recorded at by looking at the file name. Most Video Management Systems will export footage and include the date and time in the name. VisualCortex colour-codes the time elements that it has detected.

1 Select videos
2 Set times and import

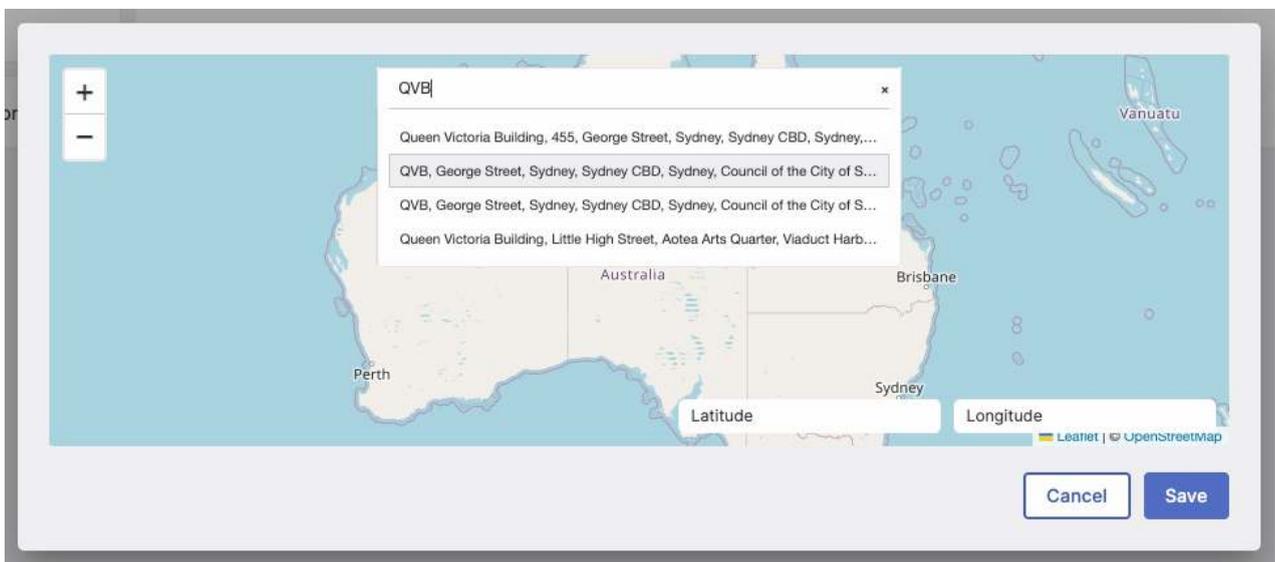
24 Dec, 2025, 7:54:40 AM
/ 2025 12 24 _07 54 40_faces_dataset_celebs_1679.mp4

Type	Name *	Location	Time Zone *	Edge Cluster *	
New Camera v	QVB East Cam 01	-33.87169, -33.87169	Australia/Sydney v	Sydney City Edge Cluster	<input type="button" value="Start uploading files"/>

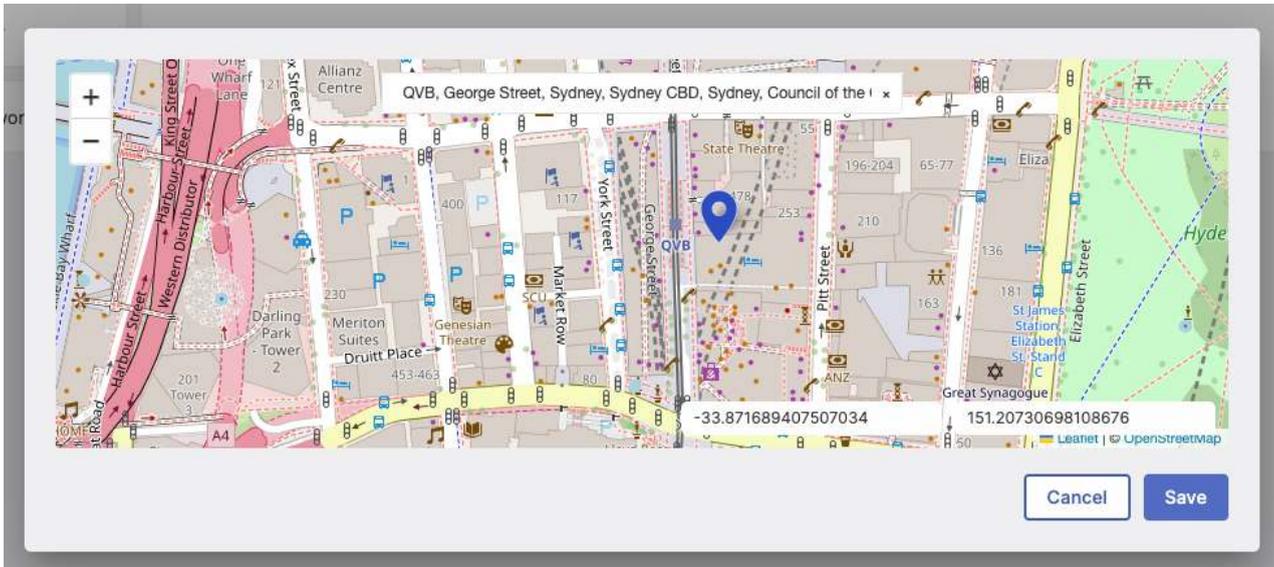
Note: If you're uploading footage to a camera which is already in an Investigation, be sure to choose *Existing Camera* to avoid having to re-enter all the camera information.

4.3.2.1 GeoLocation

Cameras can optionally have Geolocation information. Simply search for a location in the search bar and choose an area.



You'll be presented with a map. Simply long-press on the point of the map in order to place the map pin representing the camera's position.



Note: You can also add in Latitude and Longitude manually if you have that information.

4.3.2.2 Uploading

If you're uploading files (as opposed to importing them), you'll see a button titled *Start Uploading Files*. Click this to start the upload process. You'll see a progress bar, indicating how much of the footage has been uploaded into VisualCortex.

1 Select videos 2 Set times and import

24 Dec, 2025, 7:54:40 AM / 2025 12 24 _07 54 40_faces_dataset_celebs_1679.mp4 ↻

Type	Name *	Location
New Camera ▾	QVB East Cam 01	-33.87169, -33.87169 ▾

Time Zone * Edge Cluster *

Australia/Sydney ▾ Sydney City Edge Cluster

Importing footage 1/1 into investigation... Start uploading files

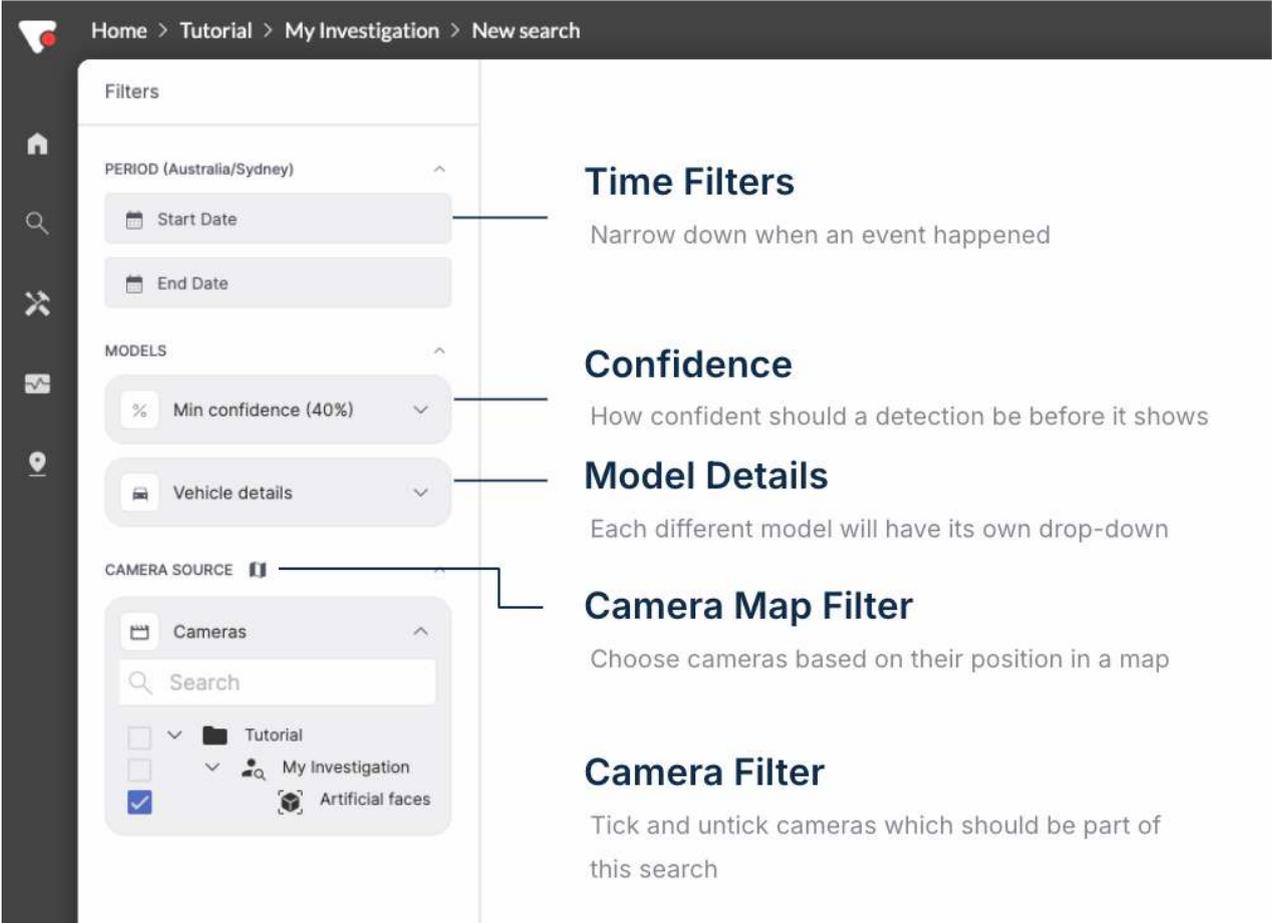
Note: Uploaded footage can take several minutes to get ready, while VisualCortex scans the files for threats and normalises the video.

4.4 Searches

An Investigation Search is the place where users of VisualCortex can look specific people, faces, license plates and more.

4.4.1 Filters

When you start a new Search, you'll get a blank screen with filters.



The screenshot shows the 'New search' interface with a sidebar on the left and a main content area on the right. The sidebar contains a navigation menu and a 'Filters' section with the following categories:

- PERIOD (Australia/Sydney)**: Includes 'Start Date' and 'End Date' filters.
- MODELS**: Includes 'Min confidence (40%)' and 'Vehicle details' filters.
- CAMERA SOURCE**: Includes a 'Cameras' filter with a search bar and a list of sources: 'Tutorial', 'My Investigation', and 'Artificial faces' (checked).

The main content area on the right provides detailed descriptions for these filters:

- Time Filters**: Narrow down when an event happened
- Confidence**: How confident should a detection be before it shows
- Model Details**: Each different model will have its own drop-down
- Camera Map Filter**: Choose cameras based on their position in a map
- Camera Filter**: Tick and untick cameras which should be part of this search

4.4.1.1 General Filters

All searches can have "general" filters, not related to a [Model](#). These include:

- **Time Filters** - Choose the Time and Date you want to look for detections in
- **Confidence** - All [Models](#) have the concept of *confidence* for what they've detected. A lower confidence generally means more (but less accurate) results, while a higher confidence means fewer, but more high quality results.
- **Camera Filters** - You can change which Sources you want to look for detections in. You can do this by ticking and unticking sources, or by choosing an area on a map (if you've added in Geolocation information)

4.4.1.2 Model-Specific Filters

Every [Model](#) which you add into an Investigation comes with their own options for searching. VisualCortex has a range of built-in models, as well as partner models. VisualCortex even supports custom models. The idea is that anything that a model can detect can be selected here in the model-specific filter panel.

The image shows a mobile-style interface for a 'Vehicle details' model. It features several filter sections:

- Vehicle Attribute:** A dropdown menu with 'Select Option'.
- Numberplate LIKE:** Three input fields with minus or plus signs. The first contains 'AB%', the second 'QNQ640', and the third 'Fuzzy search'.
- Vehicle Attribute / Color:** A dropdown menu with 'Blue, Green'.
- Vehicle Attribute / Type:** A dropdown menu with 'Vehicle'.
- Vehicle Attributes / Country / Australia:** A dropdown menu with 'NSW'.
- Vehicle Attributes / Make:** A dropdown menu with 'Holden'.
- Vehicle Attributes / Power Source:** A dropdown menu with 'Select Option'.

Most model filters are Select boxes, where you can choose one or more options. An example of this might be *Vehicle Colour*. Select boxes usually contain mutually exclusive items, and

work as **OR** filters - meaning if you were to select both *Blue* and *Green*, VisualCortex would show you vehicles which were either blue or green.

Free Text Search

Some models allow free-text search. An example of this is the *Vehicle Details* Numberplate field.

Free text search follows these rules:

- `_` means a single wildcard character
- `%` mean 0 or more characters

Examples include:

- `T%` means all plates starting with `T`
 - eg. `TFM112`, `T00330`, `TSR012`
- `ABC%3` means all plates starting with `ABC` and ending in `3`
 - eg. `ABC123`, `ABCGH3`, `ABC77773`
- `ABC_3` means `ABC` + any one character + `3`
 - eg. `ABC13`, `ABC23`, `ABCD3`
- `ABC123` means match this plate exactly

4.4.2 Results

Once you've set your filters and pressed search, results will appear in a grid. Each detection has metadata and indicators as to whether or not it's been previously viewed.

View Type
Choose between tiles, list, or Map view

Indicators
 means that somebody has already viewed this detection.
 Means that the detection has been added to a report

Time Length
How long the detection was seen for

Click to Open
Click any detection to open up the details

View Controls
Change sorting order, size of tiles, and space around detections

Download Detections
Download detection metadata as a CSV

You can change the results to display in a list view if desired. This shows the detection, as well as the area in the full frame in which it was detected.

All detections Not seen before Seen before

9:17 AM, 24 Apr 2024 - Street view

9:17 AM, 24 Apr 2024 - Street view

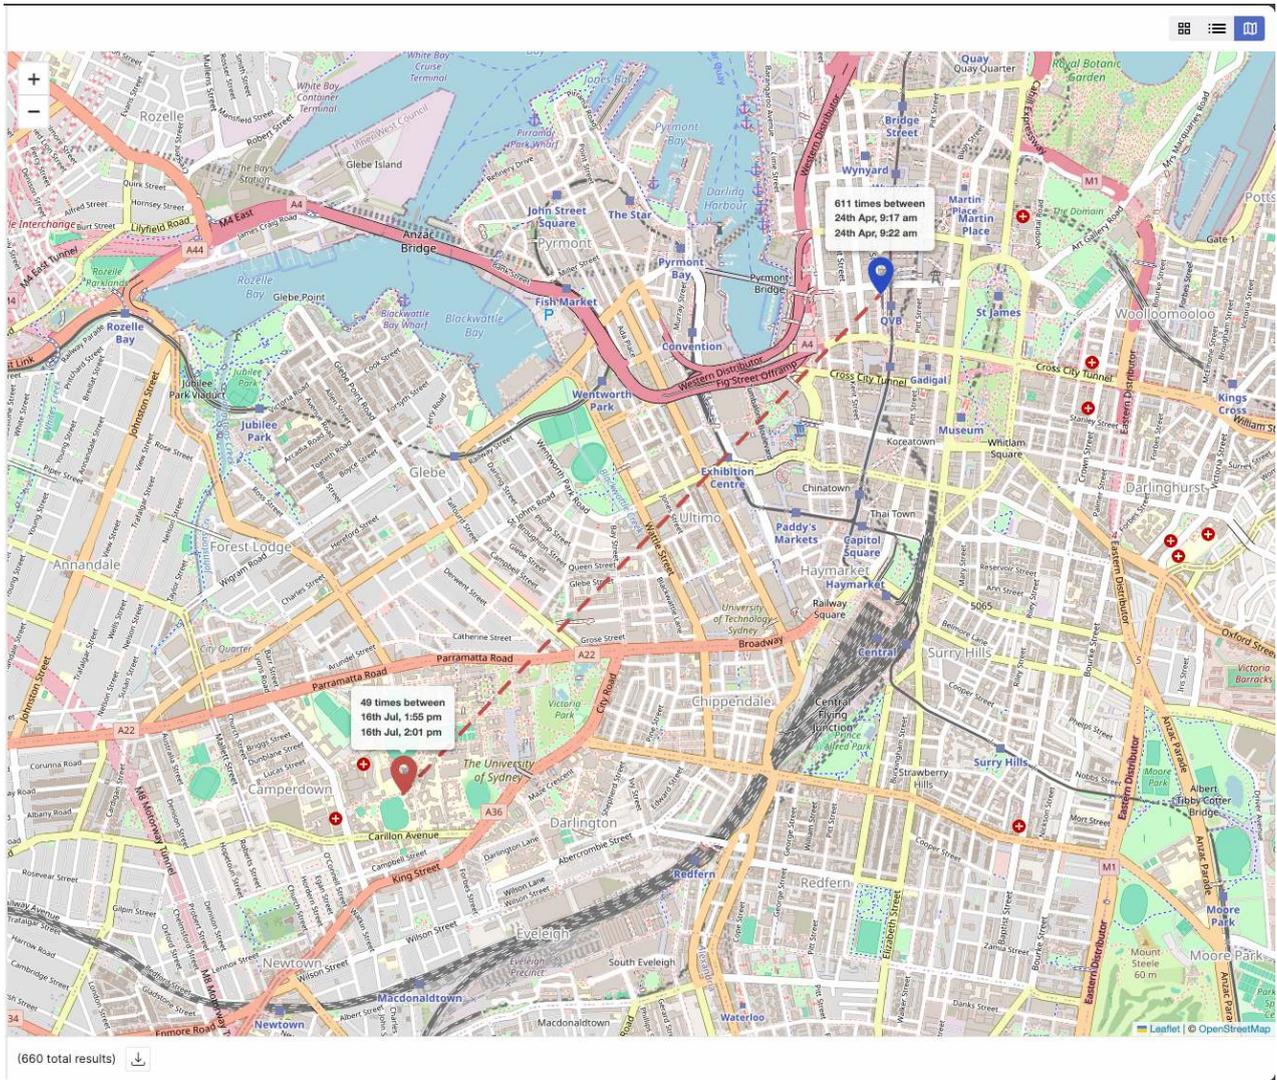
9:17 AM, 24 Apr 2024 - Street view

Add comment

Add comment

Add comment

Lastly, map view allows you to see detections on a map if applicable. Click on any of the map markers to view the detections seen at that marker in a list.



4.4.3 Detections

When you click on a result from the main results page, you'll be presented with a popup window. This window has a variety of information about the detections, and gives you the opportunity to add information to the Investigation report. The image itself is also interactive; you can drag it around and zoom in and out.

Important: Ensure you click the *Add to report* button to add a detection and any commentary to the Investigation Report

Camera Name Street view

Annotation
Add in commentary about the detection for the Investigation Report

Detection Information
Dates, Times, Detection Type, Length etc.

Manual Tags
Put your own custom tags here to group similar detections

Labels
24 Apr 2024 9:17:51 AM
Timezone: Australia/Sydney
for 6 second/s
...b8-75b6-30a1-7242-da331985a4f2
Person

Manual tags +

[Add to report](#)

Picture Controls
Sharpness, Contrast, Brightness

4.4.3.1 Controls

The detection popup comes with a number of controls to help you gain more insight. These are outlined in the below image.

Reset View

Reset any panning or zooming applied to the image

View Other Detections

Click to view other detections in the frame. Click these other detections to create a new search.

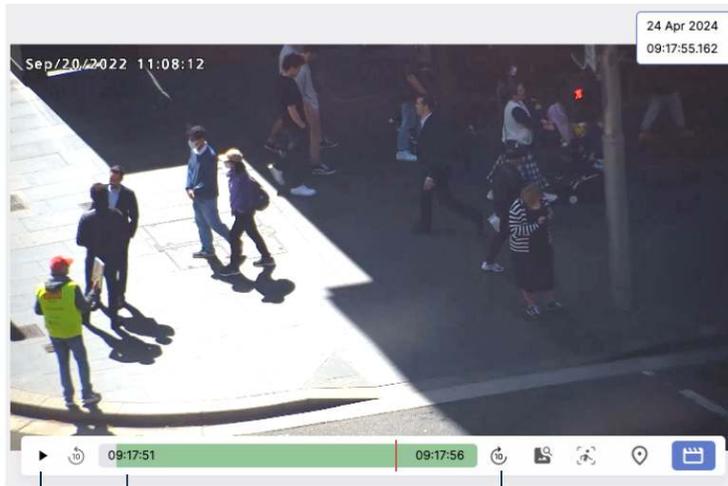
Map View

View this camera's position on a map

Video Snippet

See the associated video snippet for this detection

When viewing video snippets, a different set of controls appears. The functionality is outlined in the below image.



Play / Pause

Timeline

Click to change the video playback playhead. The green highlight represents the time where the current object was detected

Snippet Extension

Click the forward and back buttons to load an additional 10 seconds of footage

Time Indicator

See the exact time in the timezone in which it was recorded

4.5 Report

The Investigation report takes all annotations made in any Investigation [Searches](#), and puts all of that information in one consolidated place. Every detection with an annotation is displayed, along with a thumbnail, the annotation note, and geolocation (if available).

The events are displayed in chronological order. Additionally, custom notes can be added. Notes are a useful way to outline other events which may have happened in an investigation, which weren't necessarily captured on video.

The screenshot shows the 'Report' tab in the VisualCortex interface. On the left, a list of detections is shown, including timestamps and camera types like 'Street view' and 'LPR Rounda...'. The main content area displays a detailed report for a specific detection: 'Apr 24, 2024, 9:17:51 AM - Street view'. This report entry includes a video thumbnail, a note with the text 'Main Suspect E. Carol spotted on Pavement. Suspect appears to be on their phone, making calls. Extremely', and a map showing the geolocation. Below this, another entry for the same time shows a note: 'Suspect leaves main footpath area. Enters shopping centre.' followed by 'Last updated by Ben Evans' and a map. A third entry is a custom note by Ben Evans: 'No video footage, but at this time, prime suspect is reported to have been seen giving a knife to Asian male. Several eyewitness accounts' with 'Last updated by Ben Evans'. A fourth entry shows a note: 'Unknown Suspect - Asian appearance. Has concealed knife - threatening of pedestrian happens at 11:08:17' with 'Last updated by Ben Evans' and a map. The bottom of the screenshot shows the start of another entry: 'Apr 24, 2024, 9:18:42 AM - Street view'.

On the left hand side, all detections are displayed in order. You can click any of the items to scroll the report to that section. Additionally, if you click on the notes, map, or thumbnail, you'll be brought to the relevant [detection](#).

You can also view the detections on a map. In this view, the geolocation points of the relevant cameras are display, along with a line going between them. You can hover over the events in the left hand pane to highlight the relevant map point. Additionally, you can click on a map point and get a list of detections (in the report) which happened at that location.

All

- 2024-04-24, 9:17 AM - Street view
- 2024-04-24, 9:17 AM - Street view
- 2024-04-24, 9:17 AM - Ben Evans
- 2024-04-24, 9:17 AM - Street view
- 2024-04-24, 9:18 AM - Street view
- 2024-04-24, 9:18 AM - Street view
- 2024-06-19, 9:08 AM - With P/IB/T...
- 2024-07-16, 1:55 PM - LPR Round...

24 Apr 2024 9:18 AM
Australia/Sydney

1) 5 times between
24th Apr 9:17 am
24th Apr 9:18 am

Leaflet | © OpenStreetMap

5. Integrations

VisualCortex integrations allow VisualCortex to interact with other systems. This section details the different integrations available.

Authentication Providers

VisualCortex allows a deep integration with a range of Social and Enterprise login providers throu...

Event Integrations

VisualCortex has the ability to send event data to other sources. This data is pre-aggregated, me...

Video Integrations

For Virtual Cameras, VisualCortex allows simple uploads of video files into the platform. While thi...

VixVizion

VixVizion Logo

Video Management System Integrations

VisualCortex has the ability to integrate with several Video Management Systems (VMS's). This al...

5.1 Authentication Providers

VisualCortex allows a deep integration with a range of Social and Enterprise login providers through deep integration with [Auth0](#).

Auth0 allows flexible and robust integration with a variety of providers, and should work with most Enterprise Single-Sign-On (SSO) solutions. Some of the identity providers Auth0 integrates with include:

- Microsoft Azure / Entra ID
- Google Workspace
- SAML
- Okta
- PingFederate
- ADFS
- OpenID Connect

All VisualCortex requires is an Auth0 tenant and application. From there, Authentication can be customised for the customer through Auth0's comprehensive user interface.

5.2 Event Integrations

VisualCortex has the ability to send event data to other sources. This data is **pre-aggregated**, meaning that it's not raw detection data, but rather data which has already been aggregated according to any [Event Definitions](#) you've set up.

This means that rather than sending information about detections (eg. I found a person at coordinates `X,Y`), VisualCortex will instead send information about relevant [Event Definitions](#) (eg. person with ID `abc` crossed the Line `Intrusion Zone A`).

Event integrations are near-real-time, with an average end-to-end time of between 20 and 30 seconds. Events are also [batched](#) so as to not overload the downstream systems.

VisualCortex currently supports the following Event Integration types:

- [Google BigQuery](#)
- [Google Pub/Sub](#)
- [Webhook](#) (Control Plane)
- [Webhook](#) (Edge)
- [Apache Kafka](#)
- [Postgresql](#)

5.3 Video Integrations

For [Virtual Cameras](#), VisualCortex allows simple uploads of video files into the platform. While this is convenient, many organisations already have data in a variety of locations and don't wish to duplicate this in VisualCortex.

To help with this, VisualCortex offer a set of *Video Integrations*. These are sources which you can *import* footage from.

Note: When importing videos from a defined Video Integration, VisualCortex will only ever store a *reference* to the video. No video data will be copied, but a thumbnail will be stored in VisualCortex's internal object store.

Integrations are defined by [Organisation Administrators](#) or [Super Administrators](#). Once defined, they are available to *all* users of the platform.

To add an integration, as an Organisation or Super Administrator, navigate to *Configuration* and select *Integrations*.

5.3.1 Integration types

5.3.1.1 Google Cloud Storage

[Google Cloud Storage](#) (sometimes shortened to GCS) is a managed cloud storage service for storing unstructured data. It is ideal for storing large files, such as video.

When adding a GCS source into VisualCortex, you'll need the *Project ID*, the *Bucket Name*, the *Bucket Location* and the JSON key for a Google Cloud [Service Account](#). Any defined service account needs to have permissions to *Read* objects and *Sign URLs*.

5.3.1.2 S3-Compatible Object Stores

Amazon's [Simple Storage Service \(S3\)](#) is an industry-leading scalable, managed cloud storage service for storing unstructured data. Like Google Cloud Storage, it's ideal for storing large files, such as video.

S3 also pioneered an API, which has led to a suite of "S3-Compatible" services, some of which are completely on-premises. Examples include [Minio](#), [Dell ECS](#), [Cloudian Hyperstore](#) and [Ceph](#).

VisualCortex should be compatible with any S3-compatible storage, as long as the following conditions are met:

- The S3-compatible source is accessible (via network) from
 - The installed VisualCortex Control Plane
 - Edge Clusters which are to process data from the source
 - End Users using the platform
- The Access key and Secret key provisioned have permission to both *Read* and to create *Signed URLs*.
- The bucket supports [CORS configuration](#) and allows videos to have a web origin from the installed VisualCortex Control Plane.

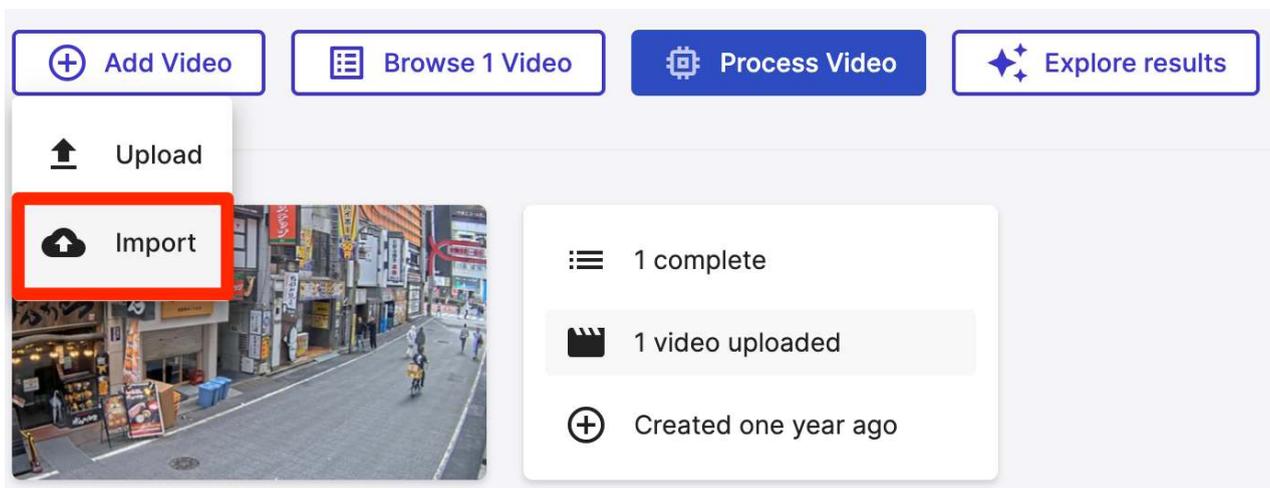
5.3.1.3 Edge Cluster Attached Storage

VisualCortex allows importing from storage attached to any Edge Cluster. This is useful for ingesting footage on internal network shares or USB disks. While not enabled by default, [VisualCortex Support](#) can enable this feature for your instance.

The storage must be attached or [mounted](#) to *Every node* in the desired [Edge Cluster](#), and must be available in the same location under `/media` on the Edge Cluster Nodes.

5.3.2 Importing into VisualCortex

Importing video from a Video Integration is simple. When adding footage to a [Virtual Camera](#), click **Import**.



From there, you'll be presented with a File Browser. You can then select the videos you wish to import. You'll be asked to provide timestamps for each of the videos. This helps in providing accurate data for [Reporting](#).

- 1 Choose external storage
- 2 Choose files to import**
- 3 Import footage

Choose files to import

2 files selected

Sort by

Name: A - Z



Street Camera - 10_03_2023.mp4



Issue Register



Video Exports



move-convert



1_ISSUE_REGISTER_FRONT_855_Morning_20260816131124.mp4



The import process can take a while as thumbnails for each video must be generated. Please be patient.

5.4 VixVizion



VisualCortex has partnered with leading Facial Recognition company, VixVizion, to facilitate Facial Recognition in [Investigations](#).

To integrate, you need to have an existing licensed VixVizion instance. From there, you'll need to configure VisualCortex to reference this server. Contact [VisualCortex Support](#) to do this.

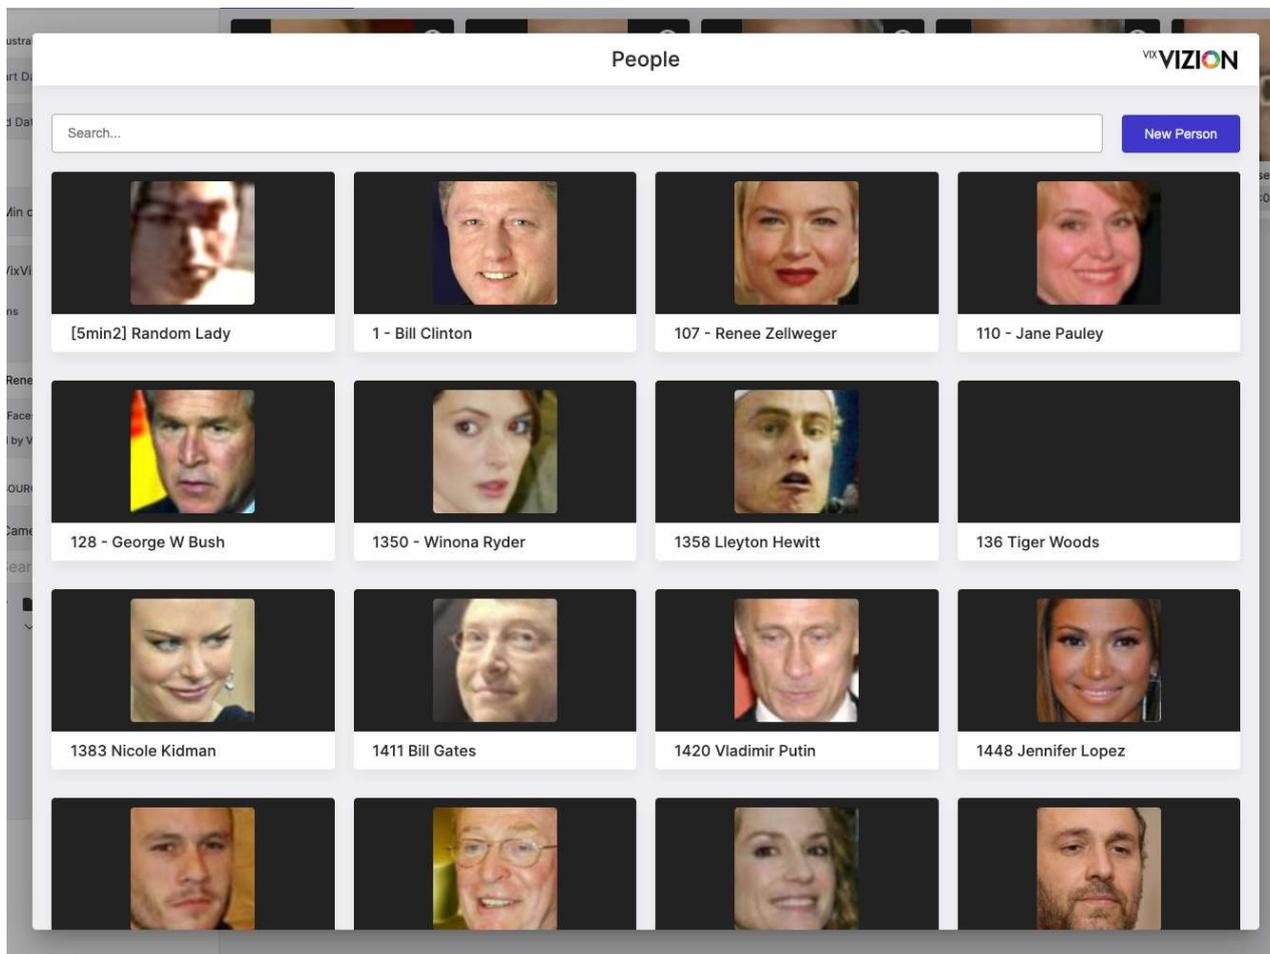
5.4.1 In Investigations

In [Investigations](#), VixVizion works just like any other model. You get a list of enrolled faces as a drop-down menu, which you can choose and search with. Results appear like any other model.

The only difference is that VixVizion also has a *Manage Faces* button in the Filters section.

A screenshot of the VixVizion interface. On the left is a sidebar with filter options: "PERIOD (Australia/Sydney)" with "Start Date" and "End Date" fields; "MODELS" with a "Min confidence (68%)" dropdown; and "VixVision" with a "Detections" section, a "Person" dropdown menu showing "107 - Renee Zellweger, 128 - ...", and a red-bordered "Manage Faces" button. Below the dropdown is the text "Provided by VixVision". On the right is a grid of three search results, each showing a face image, the label "Celebrity Dataset", a timestamp "4 Dec 2025 1:03 PM", and a duration "3s", "6.5s", and "6.5s" respectively.

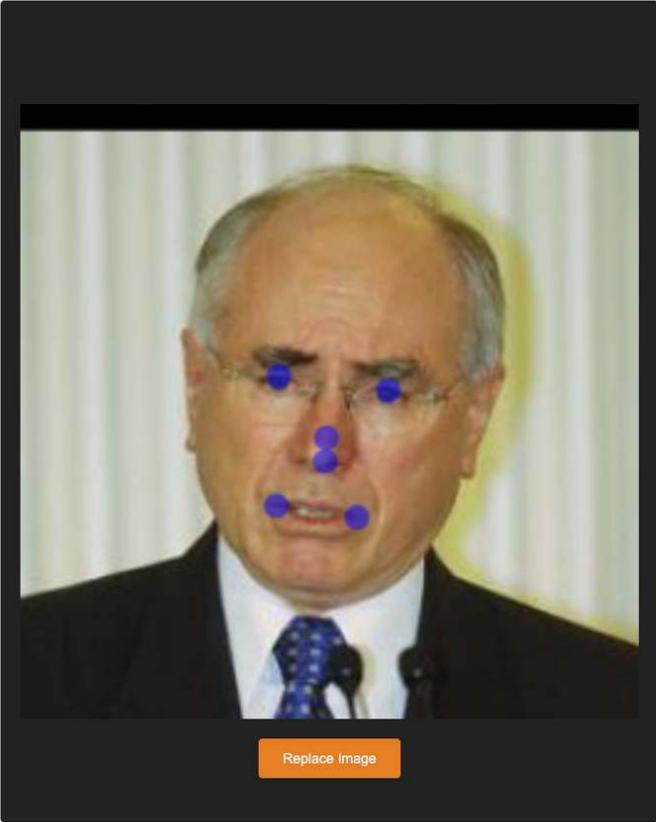
Here, you'll get to see all the enrolled people in the VixVizion database. You can click into people to add additional photos to them, or to remove an enrolled face.



Adding a new person enrolment is as simple as uploading, dragging-and-dropping, or pasting (from your computer's clipboard) into the application.

When adding a face to the application, VixVizion will perform a series of checks before it uploads. Once these checks have been verified, you can add enrol the face with a name.

< Back Review face VIX VIZION



Replace Image

Face Detection

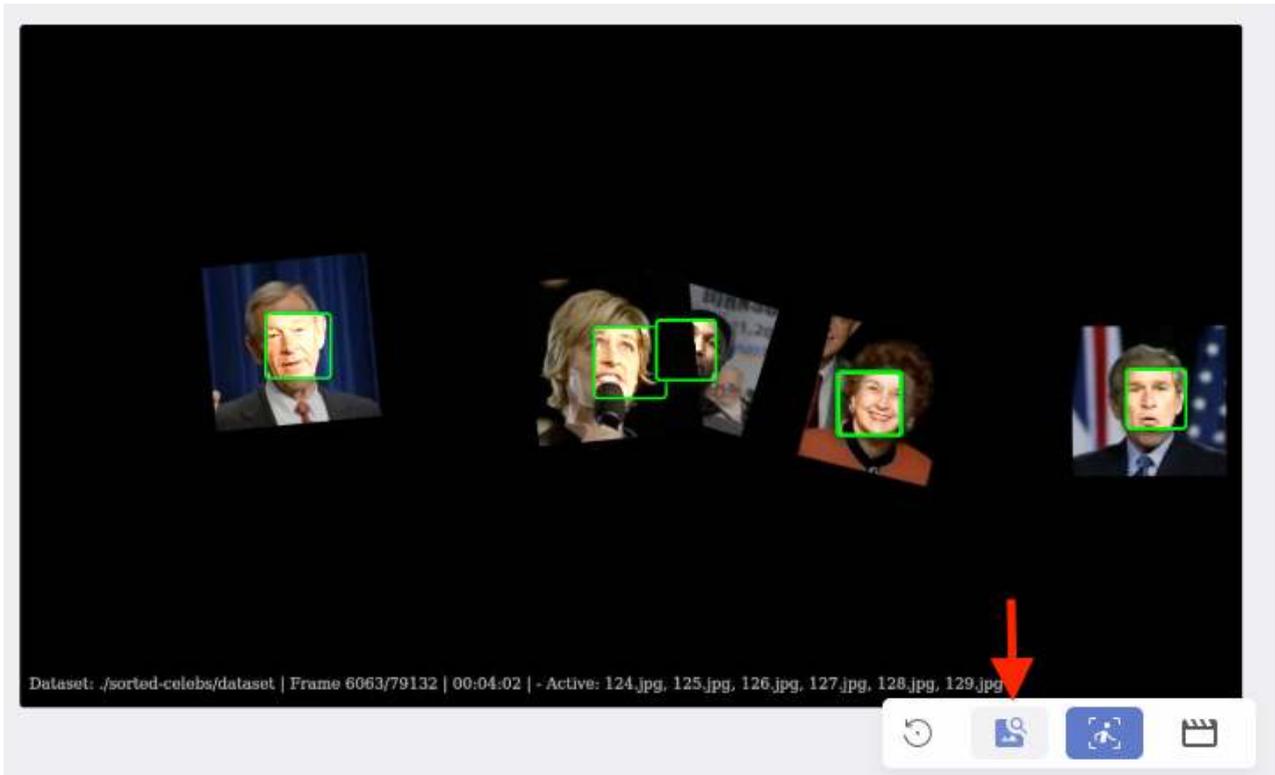
- ✓ Can see a face?
- ✓ Contains 1 face?
- ✓ Is looking forward?
- ✓ Is clear?
- ✓ Clear contrast and lighting?
- ✓ Not too similar to another face?

John Howard

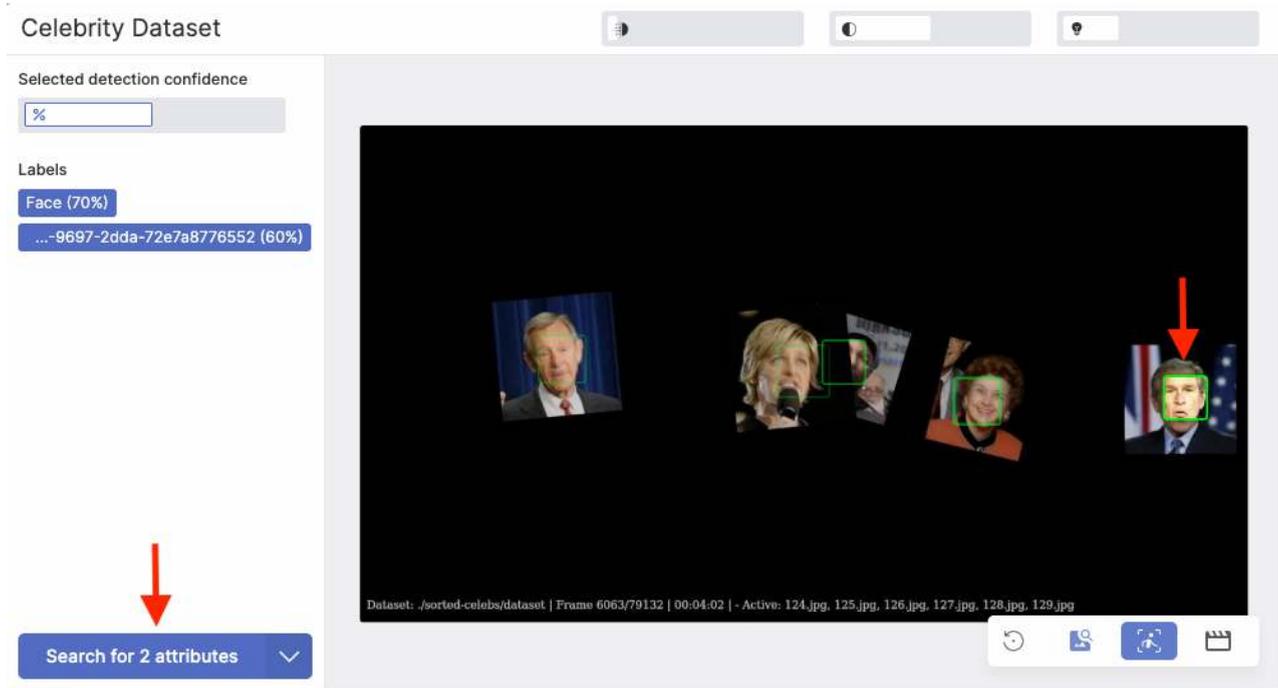
Enrol as new person

5.4.2 Searching un-enrolled faces

When you're performing an Investigation, you may not have all your known faces defined before you get started. You can search for these other detections by clicking on the "Other detections in frame" button in the [Investigation Search Detection](#).



From there, click on the face of interest, and click on the *Search for Attributes* button.



You'll be prompted to name your new search. After this, your search will appear as normal, but the unique Face ID for the face you've selected will be selected.

Filters

PERIOD (Australia/Sydney)

Start Date

End Date

MODELS

Min confidence (83%)

VixVision

Detections

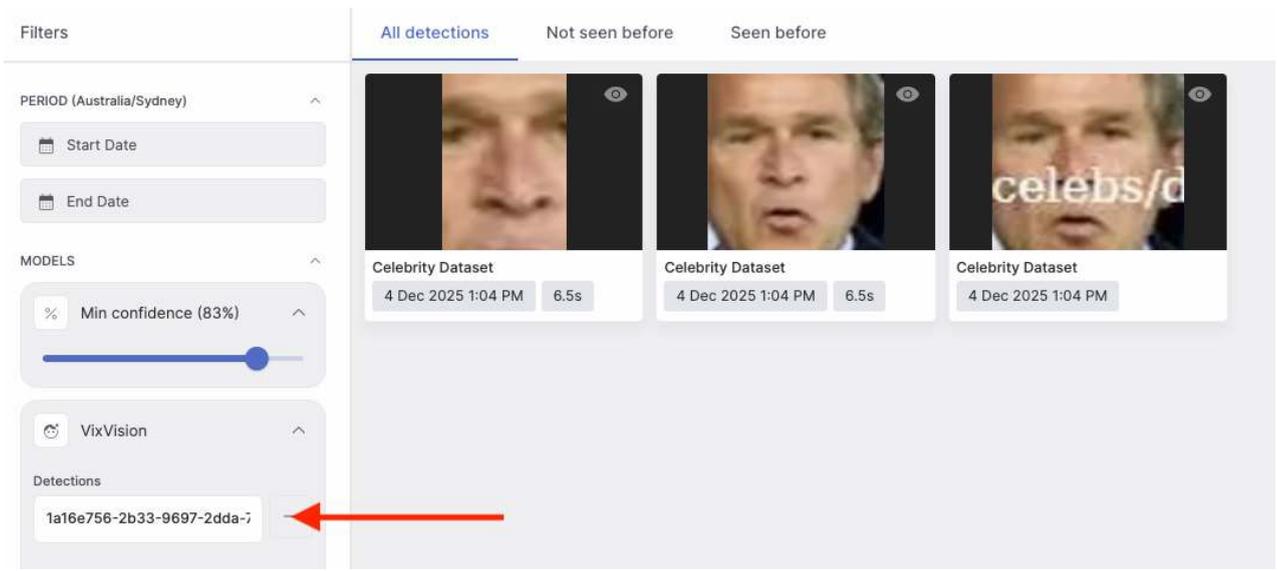
1a16e756-2b33-9697-2dda-7

All detections Not seen before Seen before

Celebrity Dataset
4 Dec 2025 1:04 PM 6.5s

Celebrity Dataset
4 Dec 2025 1:04 PM 6.5s

Celebrity Dataset
4 Dec 2025 1:04 PM



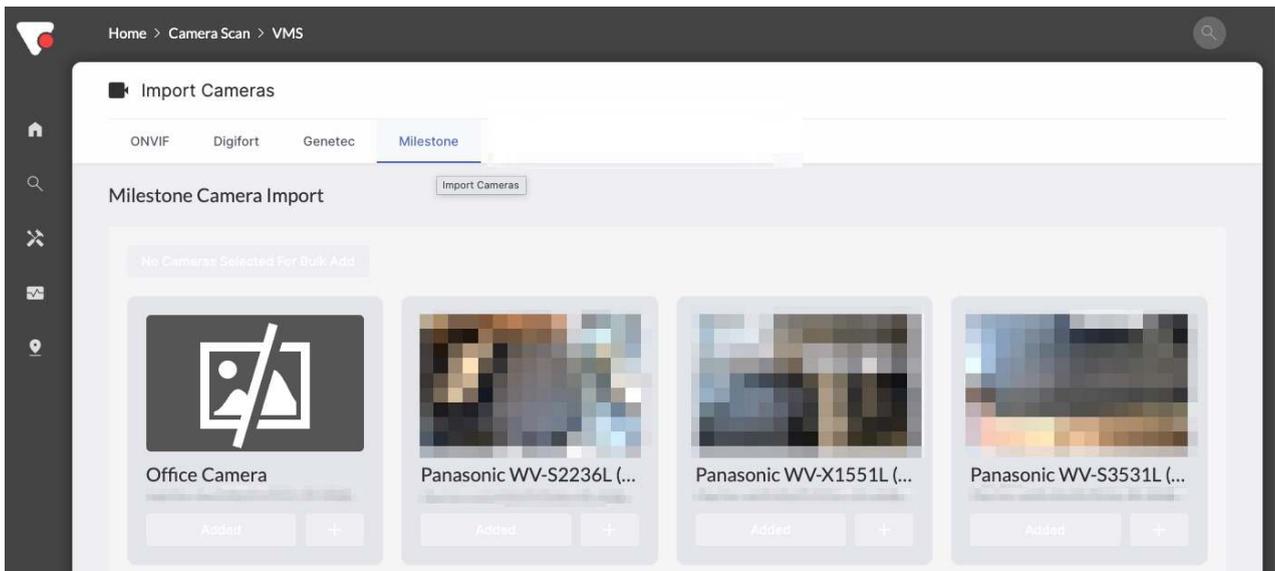
5.5 Video Management System Integrations

VisualCortex has the ability to integrate with several Video Management Systems (VMS's). This allows deeper integration with existing systems, and facilitates a range of extra functionality. Not every Video Management System supports every feature. See the compatibility list below for support for your VMS.

5.5.1 Functionality offered

5.5.1.1 Native Streaming

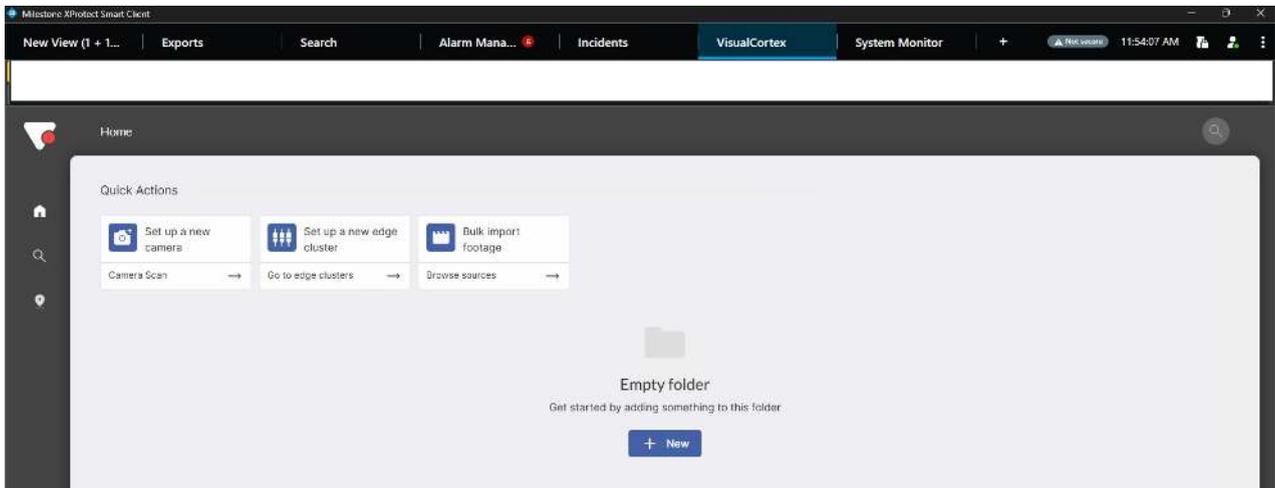
Native Streaming means that VisualCortex can natively add a [live camera](#) stream from the Video Management System. Where available, VisualCortex will always try and use a *native* connection type (if offered). When enabled for a VMS integration, users will be able to perform camera scans by going to the "Camera Scan" section, and clicking on the integration.



5.5.1.2 Embedded Access

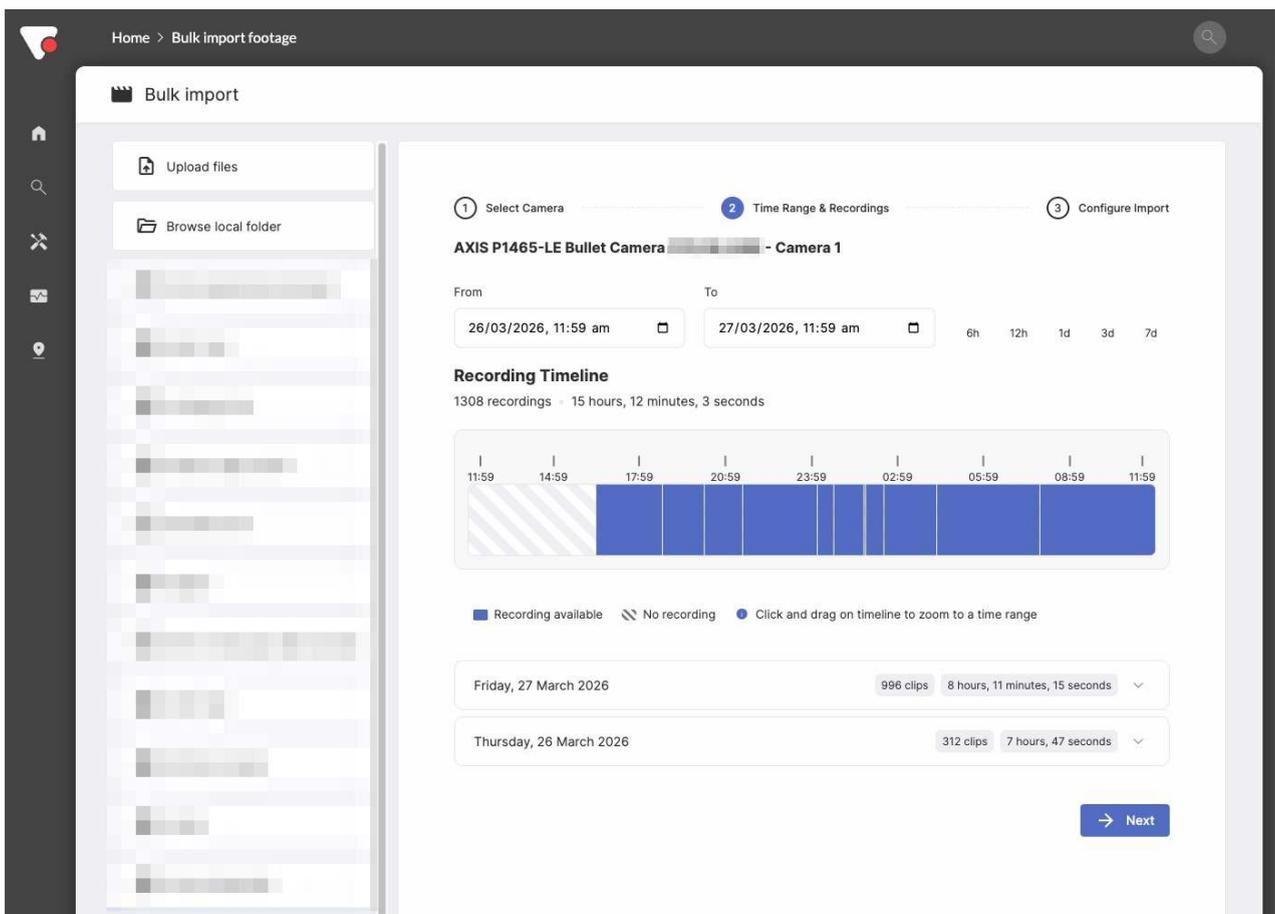
For VMS's which support embedded access, the VisualCortex interface can be made available as part of that VMS. This means that operators who are generally familiar with a particular VMS don't need to open a separate web browser in order to interact with VisualCortex.

Embedded access will pass through the authentication information (where supported) from the VMS client into VisualCortex. There are also options to automatically provision users and groups from these integrations.



5.5.1.3 Import Footage

When a VMS supports importing footage, VisualCortex can reach out to the VMS and import footage directly from it, without manual exports and imports having to be performed. It's important to note that this creates a *copy* of the footage on the Edge Cluster allocated to the VMS; it is *not* a reference. The motivation for this is so that footage for offline [investigations](#) can persist, even if the VMS clears footage after a certain period of time.

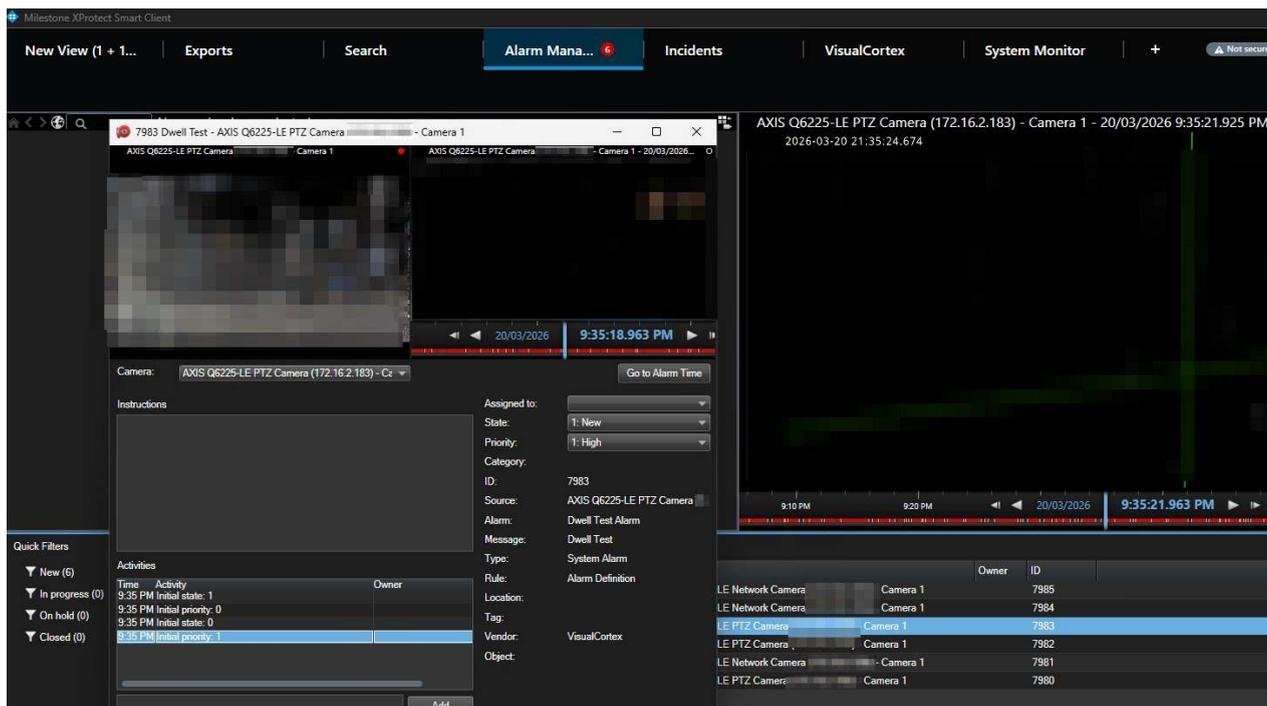


5.5.1.4 Thumbnail Generation

Thumbnail generation allows [Live Investigations](#) to use the VMS for thumbnail and video generation. In practice, this means that all image and video data are 100% stored on the VMS, but still accessible through the [Main Investigations](#) screen.

5.5.1.5 Integrated Alerts

When Integrated alerts are enabled, VisualCortex will show the VMS as an [Alert Target](#). When VisualCortex sends an alert to this alert target, this alert shows up in the Video Management System's alerting or alarm interface.



5.5.2 Support Matrix

System	Native Streaming	Embedded Access	Import Footage	Thumbnail Generation	Integrated Alerts
Milestone	✓ *1	✓ Installer	✓	✓	✓
Milestone AI Bridge	✓		✓	✓ *2	✓
Genetec	✓ *3				✓
NX Witness			✓		

System	Native Streaming	Embedded Access	Import Footage	Thumbnail Generation	Integrated Alerts
Cloud					
Digifort	✓ *4		✓	✓	✓

1. Milestone Native Streaming prior to release 5.19 uses the Open Network Bridge on port 554. Post 5.19, it uses the ImageServer protocol on port 7563
2. Due to the limitations of the AI Bridge, thumbnails for investigations are quite slow. Recommend using Native Milestone where possible
3. You will need to enable the [Genetec Media Gateway](#) for this to work. This is sometimes licensed separately.
4. When streaming, Digifort re-encodes RTSP streams. Be aware that this can *significantly* increase your server load.

6. Administration

This section outlines different administrative functions in VisualCortex.

Code Transformations

While a lot of Video Intelligence platforms give you the ability to upload custom models, many of t...

Edge Clusters

Technical Details

Groups

VisualCortex's permission model is Group-based, which means if you want to Share an item in th...

IP Whitelisting

IP Whitelisting is a security measure that regulates access to a specific location, server, or networ...

Organisations

Every installation of VisualCortex supports multiple organisations.

Roles

VisualCortex comes with a set of predefined roles already defined, but has the ability to customis...

Users

VisualCortex has the concept of Users, much like other web applications. For VisualCortex, a Use...

6.1 Code Transformations

While a lot of Video Intelligence platforms give you the ability to upload custom models, many of them do not allow you to create metadata based on custom logic. VisualCortex has *Code Transformations* which allow you to do just that.

Code Transformations are written in [Typescript](#) and allow you to add, remove and augment metadata detections in your application. Code transformations can be applied to both [Live Cameras](#) and [Virtual Cameras](#).

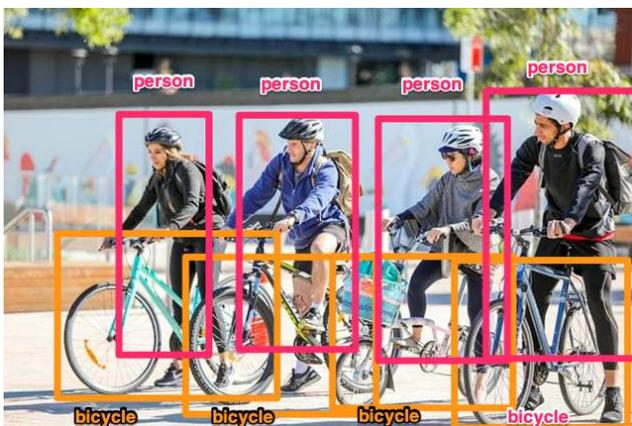
Reach out to [VisualCortex Support](#) to get technical documentation on syntax, interfaces and samples for code transformations.

6.1.1 Sample Use Cases

6.1.1.1 Grouping Classes

You might decide you want to use VisualCortex's built-in *Pedestrian* model, which outputs a variety of classes. These include both `person` and `bicycle`.

A visualisation of what VisualCortex might detect is below:



If you were to define a [Directional Count Event Definition](#), you'd see `4x person` events and `4x bicycle` events. While technically accurate, you might want to count people without bicycles separately and not include people riding bicycles in that figure.

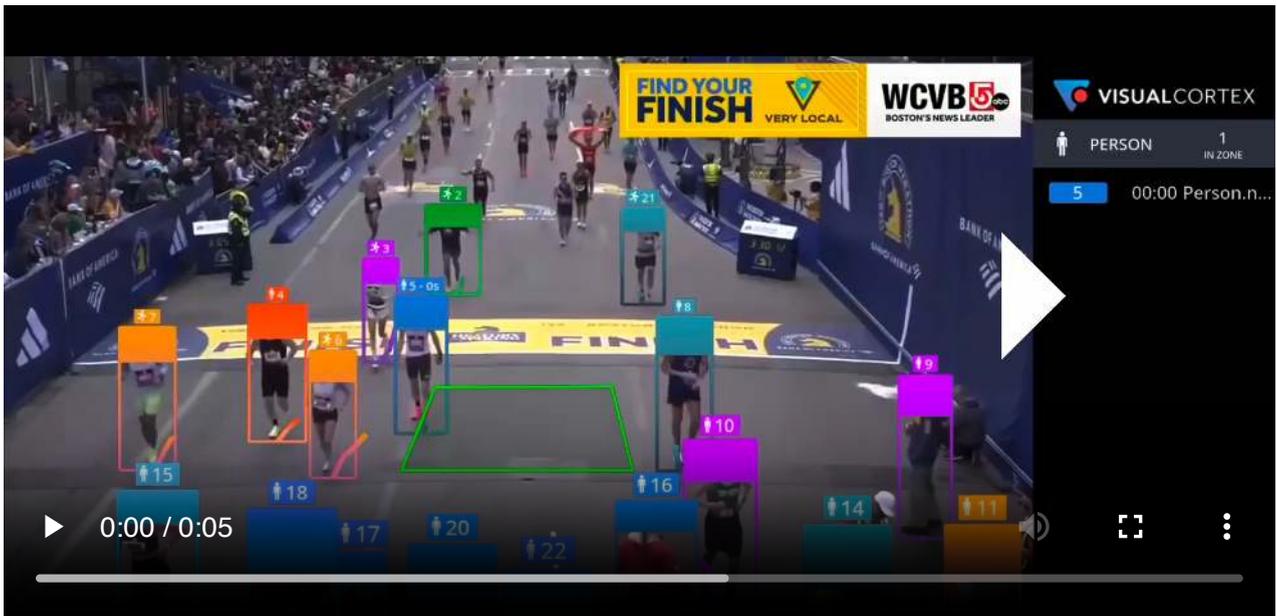
A code transformation could look at the overlapping coordinates of people and bicycles and instead create a **New** classification - `cyclist`.



This way, `person` and `cyclist` become separate objects, while `bicycle` gets removed, allowing you to get a more representative count of traffic in an area.

6.1.1.2 Speed

Code Transformations have access to any [Event Definitions](#) defined for a camera. You may decide to define a `2x2m` zone and utilise that to understand how quick an object is moving. In this example below, icons above the people change depending on whether or not the `running` attribute has been added to a `person`.



6.2 Edge Clusters

6.2.1 Technical Details

Edge clusters are Kubernetes clusters made up of one or more GPU-enabled Kubernetes [Nodes](#).

6.2.1.1 Hardware

VisualCortex supports most current generation Intel or AMD x64 processors. VisualCortex required at least **16GB** of RAM for each node in the Edge Cluster. At least one of the Edge Cluster nodes needs to have an NVIDIA GPU with compute capability of at least **7.5**. See the official [NVIDIA Website](#) for more information. Each node should have storage space of at least **250GB** - preferably SSD. Significantly more storage is required if you intend to utilise the VisualCortex Investigations module.

VisualCortex also supports the [NVIDIA Jetson Orin AGX](#) for very small deployments - [Contact Us](#) if you intend to use one of these devices.

6.2.1.2 Software

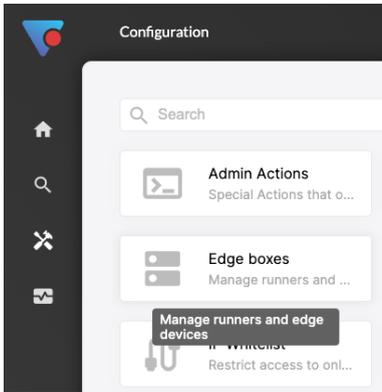
VisualCortex supports [Ubuntu Server](#) as its operating system. Both **22.04 LTS** and **24.04 LTS** are currently supported.

6.2.1.3 Networking

In terms of networking, Edge Clusters need *outbound* network access to the [Control Plane](#). This outbound access needs to support [Web Sockets](#). Most customers elect to have outbound internet access too, but this is not required. **No inbound access to the Edge Cluster's network is required** - this means no port forwarding, no DMZs etc. VisualCortex Edge Clusters will also work behind [NAT](#).

6.2.2 Management

Edge Cluster management is an administrative feature. You'll need to be either an organisation administrator or super administrator to add, remove and modify edge clusters. See [Permissions](#) for more information. You can see a list of Edge Devices by choosing the Edge management tile in the **Configuration** section.



6.2.3 Settings

Edge Clusters have a number of settings. These are outlined below:

- **Status** - This can either be *available* or *disabled*. Disable an Edge Cluster if you don't want VisualCortex to use it
- **Name** - This is the name which appears throughout VisualCortex
- **Description** - An optional description
- **Physical Location** - An optional text field representing where the Edge cluster is physically located
- **Physical GPUs** - This is the total number of GPU cores present across the cluster
- **Usable GPUs** - This is the total number of GPU cores you want used at any one time. This is typically set to a *lower* number than Physical GPUs when operating in a Highly Available cluster
- **Cluster Architecture** - This is the architecture of the cluster itself.

6.2.4 Provisioning

VisualCortex has an automated installer which simplifies the setup process for provisioning Edge Clusters. Upon creating a device, a button titled **Generate Setup Command** appears. Clicking that will give you a short-lived URL which you can paste or type into your physical Edge Cluster node.

Basic Settings

Status*

Disable the edge device to stop streaming and processing of all footage.

Available ▼

Name of Edge Device

Cluster Architecture*

Choose the architecture of the edge cluster

x86_64 ▼

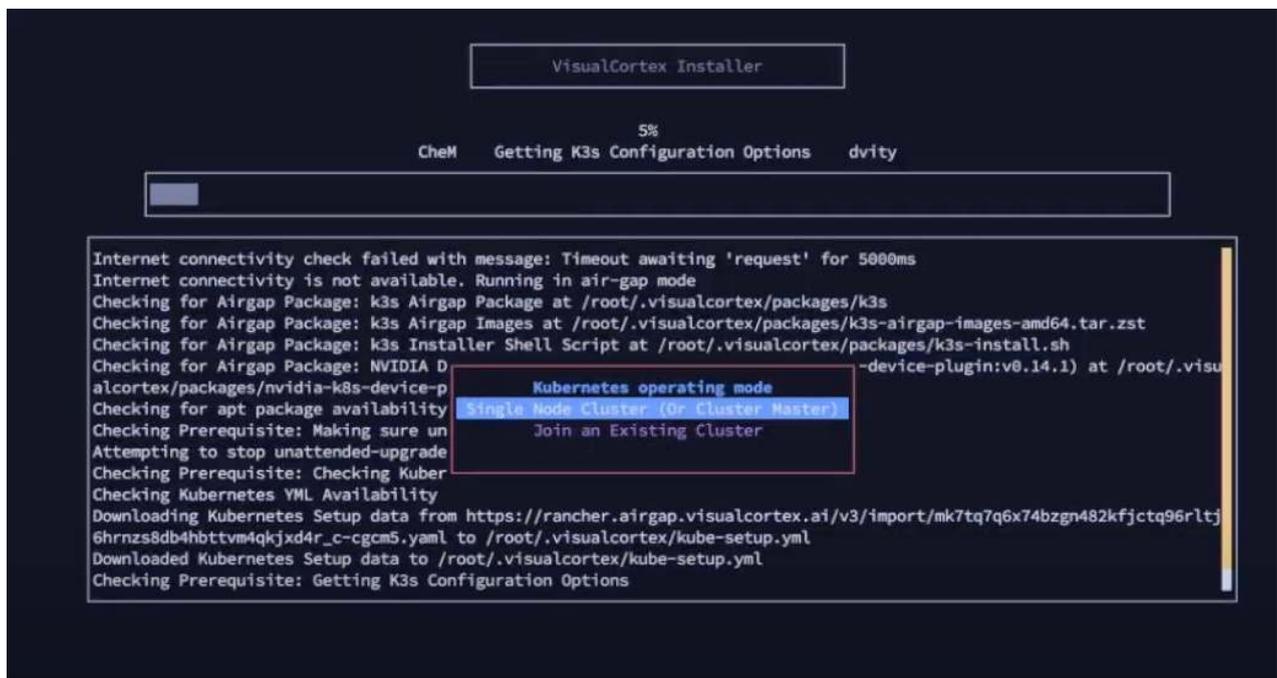
Generate One-Time Setup Command

To utilise VisualCortex on your edge device, you will need a fresh Ubuntu Server 22.04 LTS installation. Once your machine is prepared, click on "Generate Setup Command" to generate a command which you can copy and paste (or type in manually) to your edge device's terminal. Make sure you run the command as root.

Upon running the generated command, the VisualCortex Edge installer will run. You'll be asked several questions, then the installer will start. Once finished, the edge device should appear as active, and you can start adding cameras to it.



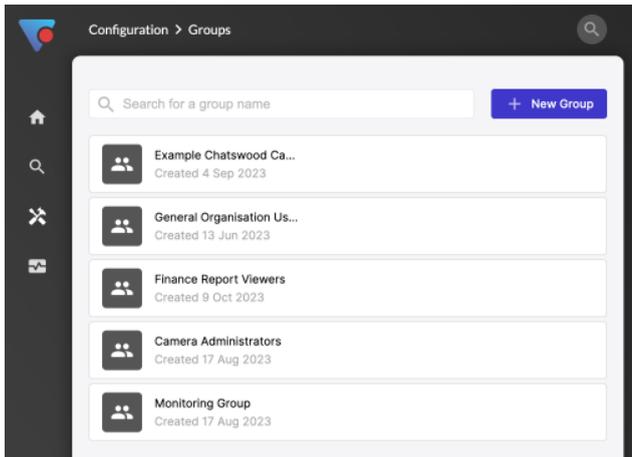
The command generated will spawn a Terminal-based GUI which will guide you through the setup process.



6.3 Groups

VisualCortex's permission model is Group-based, which means if you want to [Share](#) an item in the platform, it has to be shared with a Group.

Groups can be administered by both [Organisation and Super Administrators](#), and is available in the Configuration section of the application.

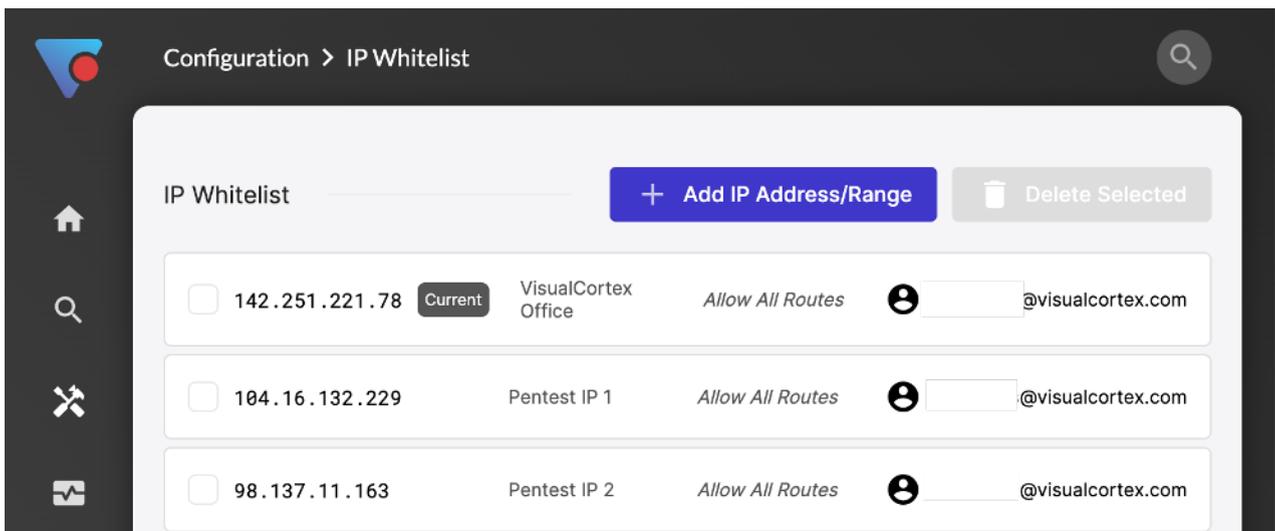


Clicking a group will bring up its member list. From there, members can be added or removed. Permissions are applied instantly.

6.4 IP Whitelisting

IP Whitelisting is a security measure that regulates access to a specific location, server, or network. While this doesn't completely limit connectivity to public-facing VisualCortex installations (which should be done with network firewall rules), it can be useful in scenarios where Edge Clusters or Users are not connected to a [Corporate WAN](#).

IP Whitelisting is only configurable by [Super Administrators](#), and is available in the *Configuration* section of the platform.



There are 2 types of whitelisting available:

- Allow All Routes
 - This is suitable for general users of the application. This option would be suitable if you were whitelisting an office's public IP address
- Allow Edge Routes Only
 - This is suitable if you want [Edge Clusters](#) to be able to push data to VisualCortex, but don't want the wider application to be accessible from that IP.

6.5 Organisations

Every installation of VisualCortex supports multiple organisations.

Organisations allow you to completely segment VisualCortex, so that it can be used by completely different teams or business units. No cameras, dashboards, users, groups, roles, alerts or credits are shareable between organisations.

Each organisation can have a different [Login Mechanism](#) with different authentication providers.

For all intents and purposes, having a separate organisation is the same as having multiple installations of VisualCortex with the following exceptions:

- Control plane resources are shared amongst all organisations
- The same URL is used for all organisations in the installation
- [Edge Clusters](#) can be shared across multiple organisations (if desired)
- Super Administrators can administer everything within all organisations in the VisualCortex installation.

6.6 Roles

VisualCortex comes with a set of predefined roles already defined, but has the ability to customise roles if required.

A role defines what a user can do when granted the role. It covers all core concepts like video sources, dashboards, folders etc. Roles which are defined in the *Roles* area in the *Configuration* section apply to the [Organisation](#) and are administered by [Organisation Administrators](#).

You define a role per *Role Source*. These include:

- [Alert Definition](#)
- [Dashboard](#)
- [Folder](#)
- [Virtual Camera](#)
- [Live Camera](#)
- [Investigation](#)

Once you give the role a name, you can set the permissions for what that role can do. This differs depending on the *Role Source*. Simply tick what the role should be able to perform and press Save.

Editor for Folder ▼

Role

Permissions for Folder

Folder

- Edit all details of this folder
- Can share a folder
- View all details of this folder

Folders Above This

- View the full file path of this folder

Folders Below This

- Edit all child folders
- Share all child folders
- View all child folders

Alert

- Can edit Alert Definitions in folders beneath this one
- Can share Alert Definitions in folders beneath this one
- Can view Alert Definitions in folders beneath this one

Camera

- Edit all cameras in folders beneath this one
- Share all cameras in folders beneath this one
- View all cameras in folders beneath this one

Event

- Edit all event definitions that are attached to cameras folders beneath this one
- View all event definitions that are attached to cameras folders beneath this one

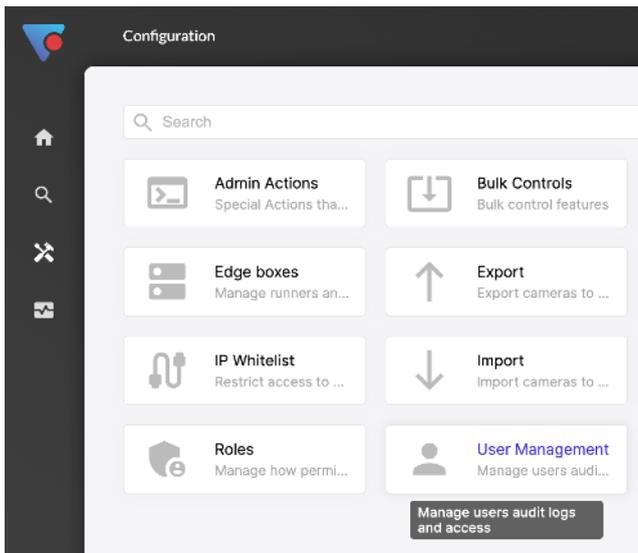
Once saved, users can [Share](#) things with [Groups](#) using the Roles you've specified.

6.7 Users

VisualCortex has the concept of Users, much like other web applications. For VisualCortex, a User's primary identifier is their email address. People can use the same email across multiple [Organisations](#).

[Organisation](#) or [Super Administrators](#) have the ability to administer users in VisualCortex.

To administer users, navigate to the Configuration page and select User Management.



Users have the following attributes:

- Email address
- First Name
- Last Name
- Roles (Optional)
 - Organisation Adminsitrator
 - Super Administrator

Users with no optional roles can be disabled by Organisation or Super Administrators. Users with Organisation or Super Administrator roles can only be disabled by Super Administrators.

Data, Cameras and Reports created by Users which are disabled still persist in the application.

6.7.1 Provisioning

Provisioning a user is as simple as getting an [Administrator](#) to add a user's email address to the platform. If your VisualCortex organisation uses a third party [Authentication Provider](#)

(Highly recommended), you won't need to specify a password. If your organisation does not use a third party Authentication Provider, you'll be asked to supply a password.

Note: Even if your third party [Authentication Provider](#) automatically provisions users based on group access, they still need to be provisioned in VisualCortex.

6.7.2 Groups

While not strictly necessary, we recommend adding Users to [Groups](#) to allow easy [Sharing](#) within the platform.