
Tech4Biz
Excellence Guaranted!

Tech4Biz

CASE STUDY CS14

One operating picture for 186 cameras across three sites

We turned three separate recording estates into a single AI monitored environment, where events find the operator instead of the operator hunting through footage.

Commercial and security operations | Three sites | Eight week delivery | Client not named at their request

WHERE WE STARTED

Three sites, three systems, and nobody watching everything

The cameras were already there and the recordings were already being kept. What was missing was any way to know when something worth looking at had actually happened.

How the estate was run

Each site had its own recorders, its own login and its own way of naming cameras.

Operators watched live walls of video for hours, with attention dropping off quickly.

Any incident review meant scrubbing continuous recordings by hand, often for a whole shift.

Analytics were limited to whatever the recorder brand shipped, and the brands did not agree.

Retention rules drifted from site to site because nobody owned them centrally.

What that cost the operation

Incidents were found after the fact, usually because somebody complained.

Footage requests from management took a day or more to answer.

There was no single number for how many people were on site at any hour.

Nobody could say with confidence that every camera was still recording.

Replacing the estate with an AI capable brand would have meant a full capital cycle.

THE WORK

What we changed, and what it produced

THE ISSUE	WHAT WE IMPLEMENTED	THE RESULT
Three isolated recording estates	Brought every camera across all three sites into one operating environment with a single login and one shared event stream	Three sites run as one, with common camera naming and one retention policy
Footage found by scrubbing	Indexed every detection as a searchable event carrying time, camera, zone, object type and confidence	Around 85 percent less time to locate relevant footage
Continuous human watching	Configured more than twenty detection scenarios and routed each one as a priority ranked alert	Around 70 percent reduction in manual video review effort
Analytics locked to camera brand	Ran the analytics as a software layer against the existing streams rather than in the cameras	128 cameras enabled for AI without replacing a single camera
No view of system health	Added stream health and tamper monitoring so a dead or blocked camera raises its own event	Camera failures surface in minutes instead of at the next audit

SCOPE BOUNDARY

Where we drew the line on scope

A surveillance system only earns trust if it is clear about its limits, so we agreed all of this with the client before the first camera was connected.

What the system does

Detects and classifies objects in every monitored stream, continuously.

Evaluates each detection against zone geometry, direction and dwell time.

Scores confidence and raises an alert only when a configured rule is satisfied.

Writes every event to a searchable index with its clip and thumbnail.

Keeps video and inference on site, inside the operator's own network.

What the system does not do

It does not identify individuals or match faces against any database.

It never takes a security decision on its own, because every alert goes to a person.

It does not replace guards, patrols or physical access control.

It does not send raw video outside the site network.

It never presents an event as certain, because every one of them carries a confidence score and can be dismissed.

AI CAPABILITY

What the cameras were taught to notice

More than twenty detection scenarios were configured across the estate. They fall into five families, and any of them can be combined into a custom rule.

Access and perimeter

Unauthorised entry into a restricted area

Perimeter intrusion

People entering or exiting a defined zone

Wrong way movement through a controlled route

After hours activity in any monitored space

Behaviour and objects

Loitering and unusual dwell time

Object left behind

Object removed from a monitored position

Crowd density above a set threshold

Queue build up at a counter or gate

Vehicles

Vehicle detection and classification

Vehicle movement into a restricted zone

Number plate recognition where the camera angle, lighting and resolution allow it

AI CAPABILITY

Counting, flow, and the health of the estate

Occupancy and movement analytics

Occupancy monitoring by zone and by hour

Person counting across defined lines

Footfall analytics for entrances and internal routes

Zone based dwell time analytics

System integrity and safety

Camera tampering or deliberate obstruction of the view

Stream loss or degraded video quality

Smoke and fire related visual anomalies, where the model configuration supports it

Custom rules

Any detection can be combined with a place, a duration and a time window to make a rule that fits how the site actually runs. Somebody standing in the loading bay for a few minutes in the afternoon is just work getting done, while the same thing at two in the morning is worth waking a supervisor for. That combination of object, location, duration and schedule is where most of the operational value came from.

HOW A DETECTION BECOMES AN EVENT

The pipeline behind a single alert

A frame goes through several stages before anyone is told about anything, and most candidate detections are dropped somewhere along the way.

01 **Ingest and normalise**

Streams from mixed camera and recorder brands are decoded into a common format and frame rate.

02 **Region gating**

Motion and masked regions decide which frames are worth the cost of inference.

03 **Detect and classify**

Objects are located and typed, with a class confidence attached to each one.

04 **Spatial evaluation**

Each object is tested against zone geometry, crossing lines and direction of travel.

05 **Temporal evaluation**

Dwell, persistence and schedule are applied so momentary noise never becomes an alert.

06 **Confidence scoring**

Class, spatial and temporal signals are combined into one score against the rule threshold.

07 **Event and clip**

A qualifying event is written to the index with its snapshot and its clip.

08 **Alert routing**

Priority decides who is told, how quickly, and on which channel.

TECHNICAL ARCHITECTURE

How the system is put together

Everything inside the dashed boundary runs on the client's own hardware at each site.

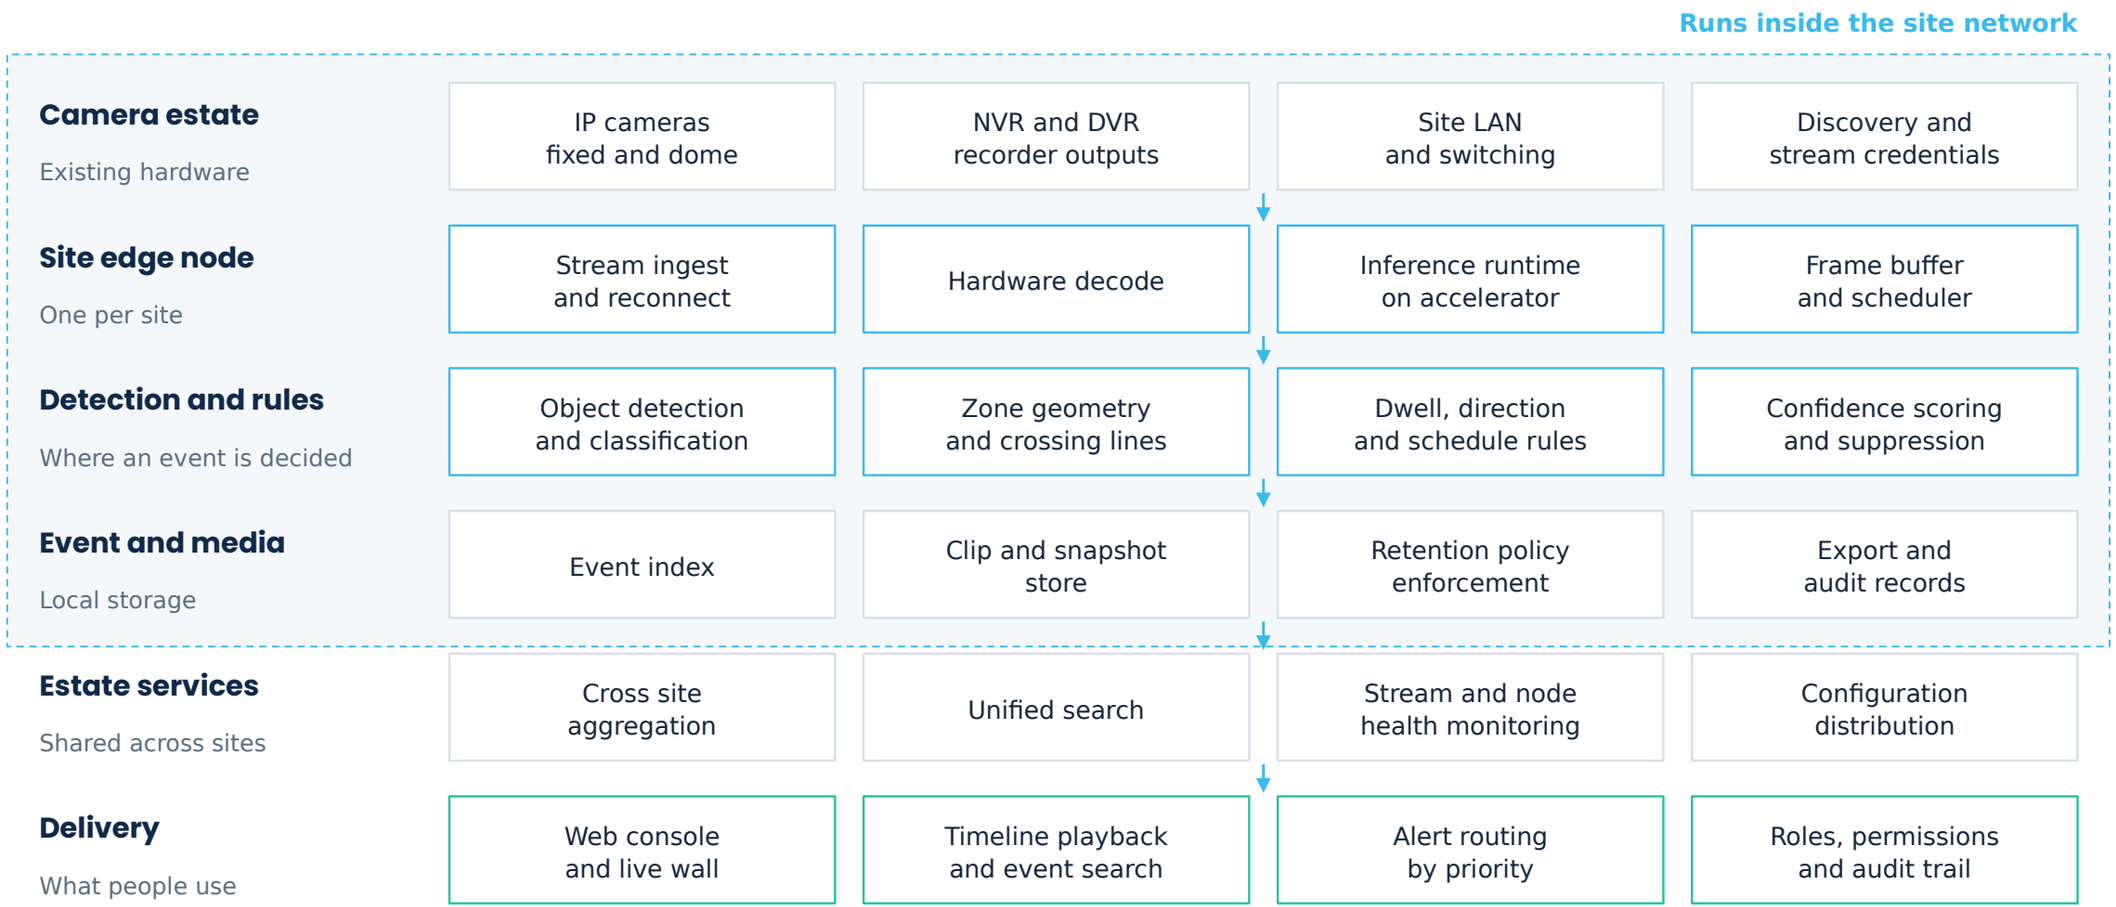

RUNTIME DATA FLOW

How video, events and alerts travel

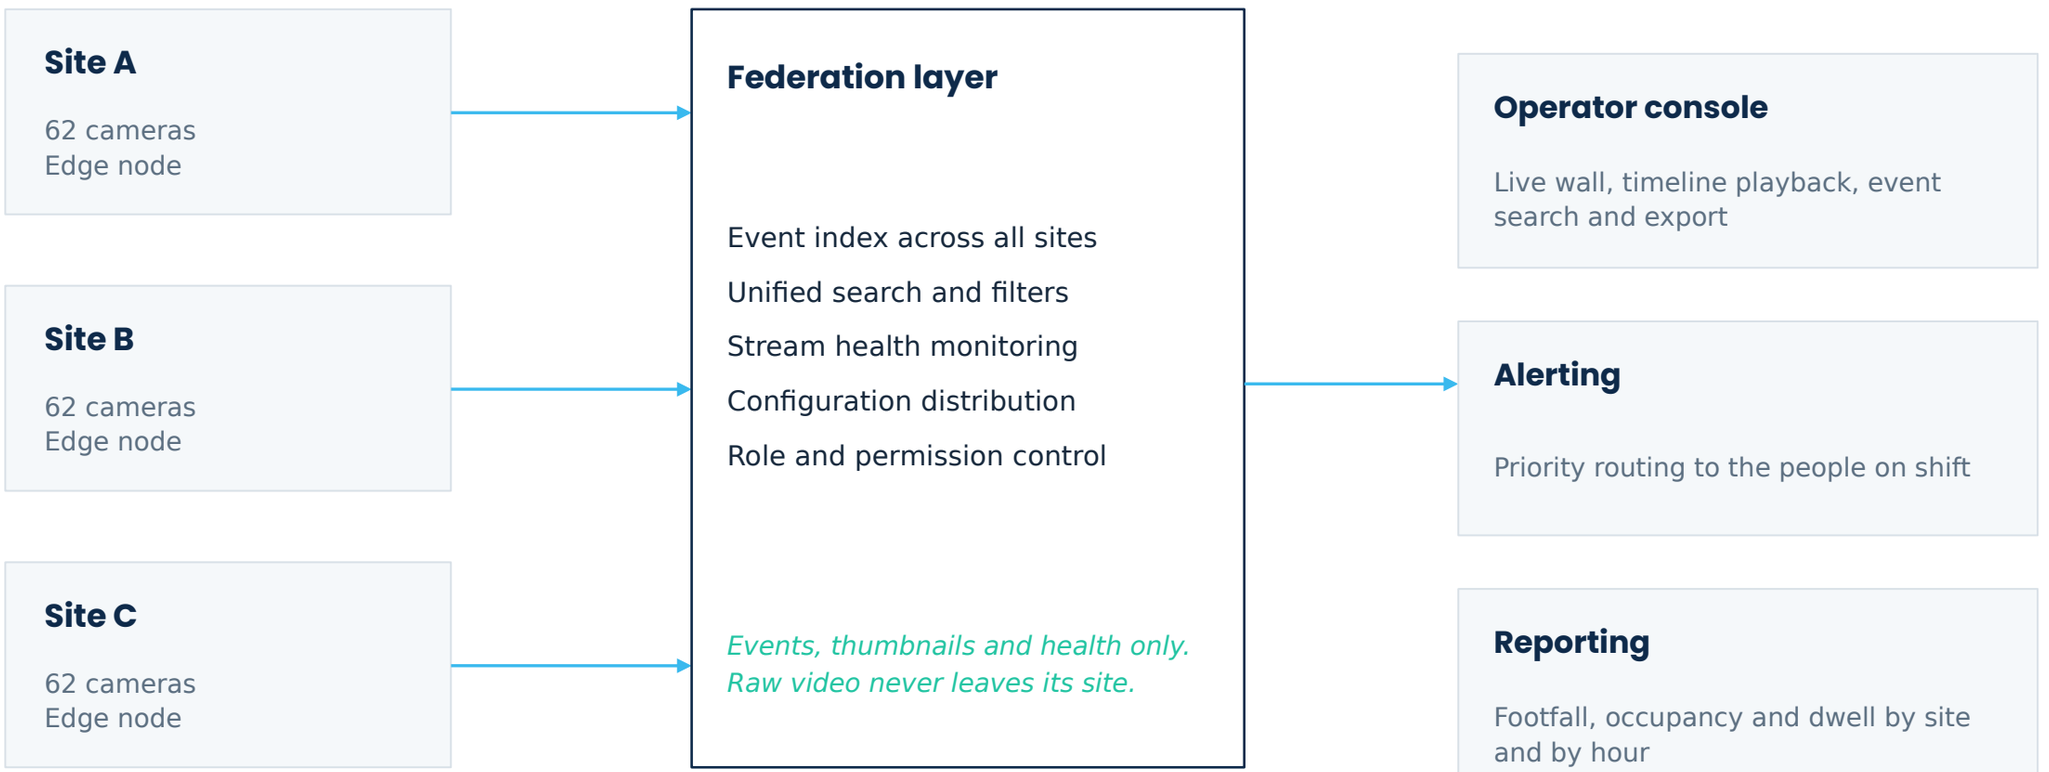

WORKED EXAMPLE

A pallet left by a fire exit at 21:40

This is the sequence as it happened in testing, timed end to end.

21:40:02	A pallet is set down inside the fire exit zone on the loading corridor camera.
21:40:04	The object is detected and classified. It sits inside a zone marked as a keep clear area.
21:41:04	The dwell rule for that zone is sixty seconds. The object has not moved, so the rule is satisfied.
21:41:05	Confidence is scored across class, position and persistence. It clears the threshold set for that zone.
21:41:06	An event is written with a snapshot and a fifteen second clip either side of the detection.
21:41:09	A priority alert reaches the duty supervisor's console and phone. Total elapsed time under ten seconds from rule satisfaction.
21:47	The supervisor reviews the clip, confirms it, and the pallet is moved. The event is closed with a reason code.

SCALE

What was actually deployed

3

Sites brought into a single environment

186

Cameras integrated across the estate

62

Average cameras per site

180+

Concurrent video streams processed

128

Cameras configured with one or more analytics

20+

Detection scenarios configured

8

Weeks from survey to handover

0

Cameras replaced to enable the analytics

SECURITY AND GOVERNANCE

Where the video sits, and who can see it

Where data sits

All video and all inference stay inside the site network.

Only events, thumbnails and health signals cross between sites.

Retention is set per site and enforced automatically, with no manual purging.

Exports are watermarked and logged against the user who made them.

Who can see what

Access is role based, down to individual cameras and zones.

Every view, search, playback and export is written to an audit trail.

Private areas are masked at ingest, so they are never recorded or analysed.

Administrative changes to rules and retention are versioned and attributable.

On identification

The system detects people. It does not recognise them. There is no facial matching, no biometric template and no watchlist in this deployment. Every alert describes a situation, not a person, and a human decides what to do about it.

DELIVERY

Eight weeks, in five stages

STAGE	WHAT WE DID	WHAT CAME OUT OF IT
Weeks 1 to 2	Site survey, stream inventory, camera by camera assessment of angle, lighting and resolution against each intended analytic	A per camera capability map, and an honest list of what each camera could and could not support
Weeks 3 to 4	Edge nodes built and commissioned at each site, all streams ingested, storage and retention configured	All 186 cameras live in one environment
Weeks 5 to 6	Zones, lines, schedules and rules drawn per camera, analytics enabled on the 128 suitable cameras	The full detection scenario set running
Week 7	Tuning against real site traffic, false positive reduction, threshold and dwell adjustment per zone	Alert volume brought down to a level operators could act on
Week 8	Operator training, runbooks, escalation matrix, handover to the client security team	The client team running the system without us

OUTCOMES

Measured after handover

MEASURE	BEFORE	AFTER
Cameras in one operating environment	0 of 186	186 of 186
Concurrent streams under analysis	None	More than 180
Manual video review effort	Baseline	Around 70 percent lower
Time to locate relevant footage	Baseline	Around 85 percent lower
Cameras with AI analytics enabled	0	128
Detection scenarios in production	0	More than 20
Priority alert delivery	Not applicable	Under 10 seconds in normal network conditions

The result in three numbers

70%

less manual video review effort

85%

less time to find the footage that matters

186

cameras under one operating picture

All of it on cameras the client already owned.

APPLICABILITY

Where this same approach works

The pattern holds anywhere there is an existing camera estate, more video than people can watch, and a need to know what happened without watching all of it.

Retail and malls

Footfall, queue build up, after hours access, object removal, loading bay control

Manufacturing

Restricted area entry, walkway compliance, occupancy by shift, tamper and camera health

Multi branch operations

One picture across many small sites, with the same rules applied everywhere

Warehousing and logistics

Yard movement, dock occupancy, wrong way traffic, keep clear zones, vehicle classification

Corporate campuses

Perimeter, tailgating patterns, visitor flow, after hours presence, occupancy reporting

Critical infrastructure

Perimeter intrusion, loitering, vehicle control, and continuous proof that cameras are alive

BEFORE YOU START

What decides whether this works for you

What we would ask you

How many cameras, at how many sites, and what brands of recorder are in place?

What resolution and frame rate does each stream actually deliver, not what it is rated for?

Which areas must be masked for privacy, and who signs that off?

How long must recordings be kept, and does that differ by site?

Who is on shift to receive alerts, and on what channel?

Where these projects go wrong

Camera angles that were fine for recording but are wrong for analytics.

Low light performance that looks acceptable to a human and is not to a model.

Bandwidth between site and control room that was never sized for this.

Alert thresholds set once and never tuned, which trains operators to ignore alerts.

Nobody owning the rules after handover, so the configuration slowly stops matching the site.

We survey every camera before committing to what each one can do. A camera that cannot support an analytic is listed as such rather than enabled and left to underperform.

Talk to us about this work

We are happy to walk through the deployment in detail, including the parts that are not in this document.

COMPANY

Tech4Biz Solutions Pvt Ltd

WEB

www.tech4bizsolutions.com

EMAIL

contact@tech4biz.io

OFFICES

Bengaluru, India
Atlanta, United States

NOTICE

About this document

This case study describes work delivered by Tech4Biz Solutions. The client has asked not to be named, so the deployment is described by sector, scale and geography only.

All figures in this document are drawn from the delivered deployment. Percentage improvements are measured against the client's own pre deployment baseline for the same activities, over comparable periods.

Detection capability depends on camera placement, resolution, frame rate and lighting. Capabilities listed here were configured and validated on cameras assessed as suitable during the site survey. Number plate recognition in particular is conditional on camera angle, shutter behaviour and illumination.

Alert timing figures assume normal network conditions between the site edge node and the operator console. No detection system is complete or infallible. This system supports human security operations. It does not replace them.

Product names, internal component names and implementation detail are omitted deliberately. We are happy to discuss them under a mutual non disclosure agreement.