
Tech4Biz
Excellence Guaranted!

Tech4Biz

CASE STUDY CS15

A second pair of eyes that never looks away from the room

We built continuous AI safety observation for a childcare operator, so that a fall, a scuffle or a child drifting toward an exit reaches a staff member in seconds rather than at the end of the day.

Early years childcare | Two pilot locations | Six week delivery and calibration | Client not named at their request

WHERE WE STARTED

Cameras were recording everything and telling nobody anything

In a childcare room, the people who need to know something happened are already in the room, already holding a child, and already looking the other way.

How it worked before

Cameras recorded continuously and were watched by nobody during the day.

Incidents were reconstructed afterwards, from memory and from footage, often days later.

A parent question about a bruise meant an educator scrubbing hours of video by hand.

Ratio and coverage were tracked on paper, at set times, not continuously.

Blind spots from a knocked or failed camera were found late, if at all.

What the operator wanted

To be told within seconds when something in a room needed an adult, right now.

To answer a parent the same day, with the actual clip, not a recollection.

To see whether a room had been left unattended, without asking anyone.

To do all of it without turning the setting into a surveillance environment for staff.

To keep every frame of video inside the building.

THE WORK

What we changed, and what it produced

THE ISSUE	WHAT WE IMPLEMENTED	THE RESULT
Nobody watching during the day	Continuous analysis of every monitored feed, with event based alerting to the staff on duty	Forty plus feeds under continuous observation across operating hours
Incidents reconstructed from memory	Every safety event written to a searchable record with its own clip and timestamp	Around 75 percent less retrospective review effort
Safety events described loosely	Fifteen plus defined event categories, each with its own rule, duration and threshold	A shared vocabulary staff and management both use
Coverage and ratio checked on paper	Visual occupancy counting per zone, with unattended zone and threshold alerts	Coverage gaps surface as they happen, not at the next check
Blind spots found late	Camera obstruction, movement and stream loss raise their own events	A blocked or dead camera is known about within minutes

SCOPE BOUNDARY

What we allow this system to claim

This is the most important page in the document, because in a childcare setting a system that overclaims is worse than having no system at all.

What it does

Observes movement, posture and position continuously across monitored zones.

Recognises visual patterns that may indicate a safety incident.

Raises a flag to a human, with the clip attached, so a person can look and judge.

Records occupancy and zone coverage as counts, not as identities.

Keeps every frame of video and every inference inside the building.

What it does not do

It never diagnoses an injury, because all it does is flag a visual event for a person to go and assess.

It does not identify or recognise any child or any adult. There is no facial matching.

It does not judge intent, and it never labels a person as aggressive or negligent.

It does not score, rank or performance manage staff.

It never acts on its own, since every flagged event goes to a person for review.

AI CAPABILITY

Child safety events the system was taught to flag

Each of these is a visual event, described as movement and position over time. None of them is a conclusion about a child.

Falls and posture

- A child falling to the ground
- Sudden collapse or unusual loss of posture
- A child remaining on the floor beyond a set duration
- Unusual rapid movement or distress related activity patterns

Contact and collisions

- Potential collision between children
- Potential collision with furniture or a fixed object
- Rough or forceful physical interaction, including pushing, striking, kicking and forceful grabbing or pulling
- A child being knocked down

Position and supervision

- A child entering a restricted or unsafe zone
- A child approaching an exit or access controlled door
- A child separating from the monitored group
- A child left unattended in a defined area
- Excessive crowding within a zone

AI CAPABILITY

Supervision, access and the room itself

Adult and child interaction, subject to policy configuration

Potentially forceful physical interaction between an adult and a child

A child being moved or handled outside configured activity patterns

Entry of an unauthorised person into a protected childcare zone

Adult to child ratio thresholds, derived from visual occupancy counts

A monitored zone becoming unattended

Environmental and operational events

Smoke or visible fire anomalies, where the model configuration supports it

Blocked emergency access routes and exits

Unexpected occupancy in a restricted zone

Entry and exit events, and occupancy against a threshold

Camera obstruction, tampering or stream loss creating a blind spot

Unusual activity outside operating hours

On the adult interaction category

This group is off by default and only enabled where the operator's own policy, consent position and local regulation allow it. When enabled, the system describes force and motion, never intent. A flagged event means a person should look at a clip. It is never treated as a finding about a member of staff, and it is never used for performance management.

HOW A FLAG IS MADE

How the system builds up to a single flag

A child sitting on the floor is just a child sitting on the floor, so the system only treats it as an event once it has seen the rapid posture change that got them there, how close they are to a fixed object, and how long they have stayed down.

01 **Presence**

A child is detected and tracked within the monitored zone.

02 **Activity**

Movement over consecutive frames is analysed rather than a single still image.

03 **Posture**

A change in posture or an unusual loss of upright position is identified.

04 **Spatial interaction**

Proximity to other children, to furniture, to exits and to zone boundaries is evaluated.

05 **Duration**

How long the state persists is measured against the rule for that zone and that event type.

06 **Confidence**

The signals are combined into one score and tested against the threshold for that category.

07 **Flag and clip**

A qualifying event is recorded with its clip and a short window either side of it.

08 **Human review**

The alert reaches the designated staff member. The judgement is theirs, always.

TECHNICAL ARCHITECTURE

How the system is put together

Everything inside the dashed boundary runs on hardware in the building itself.

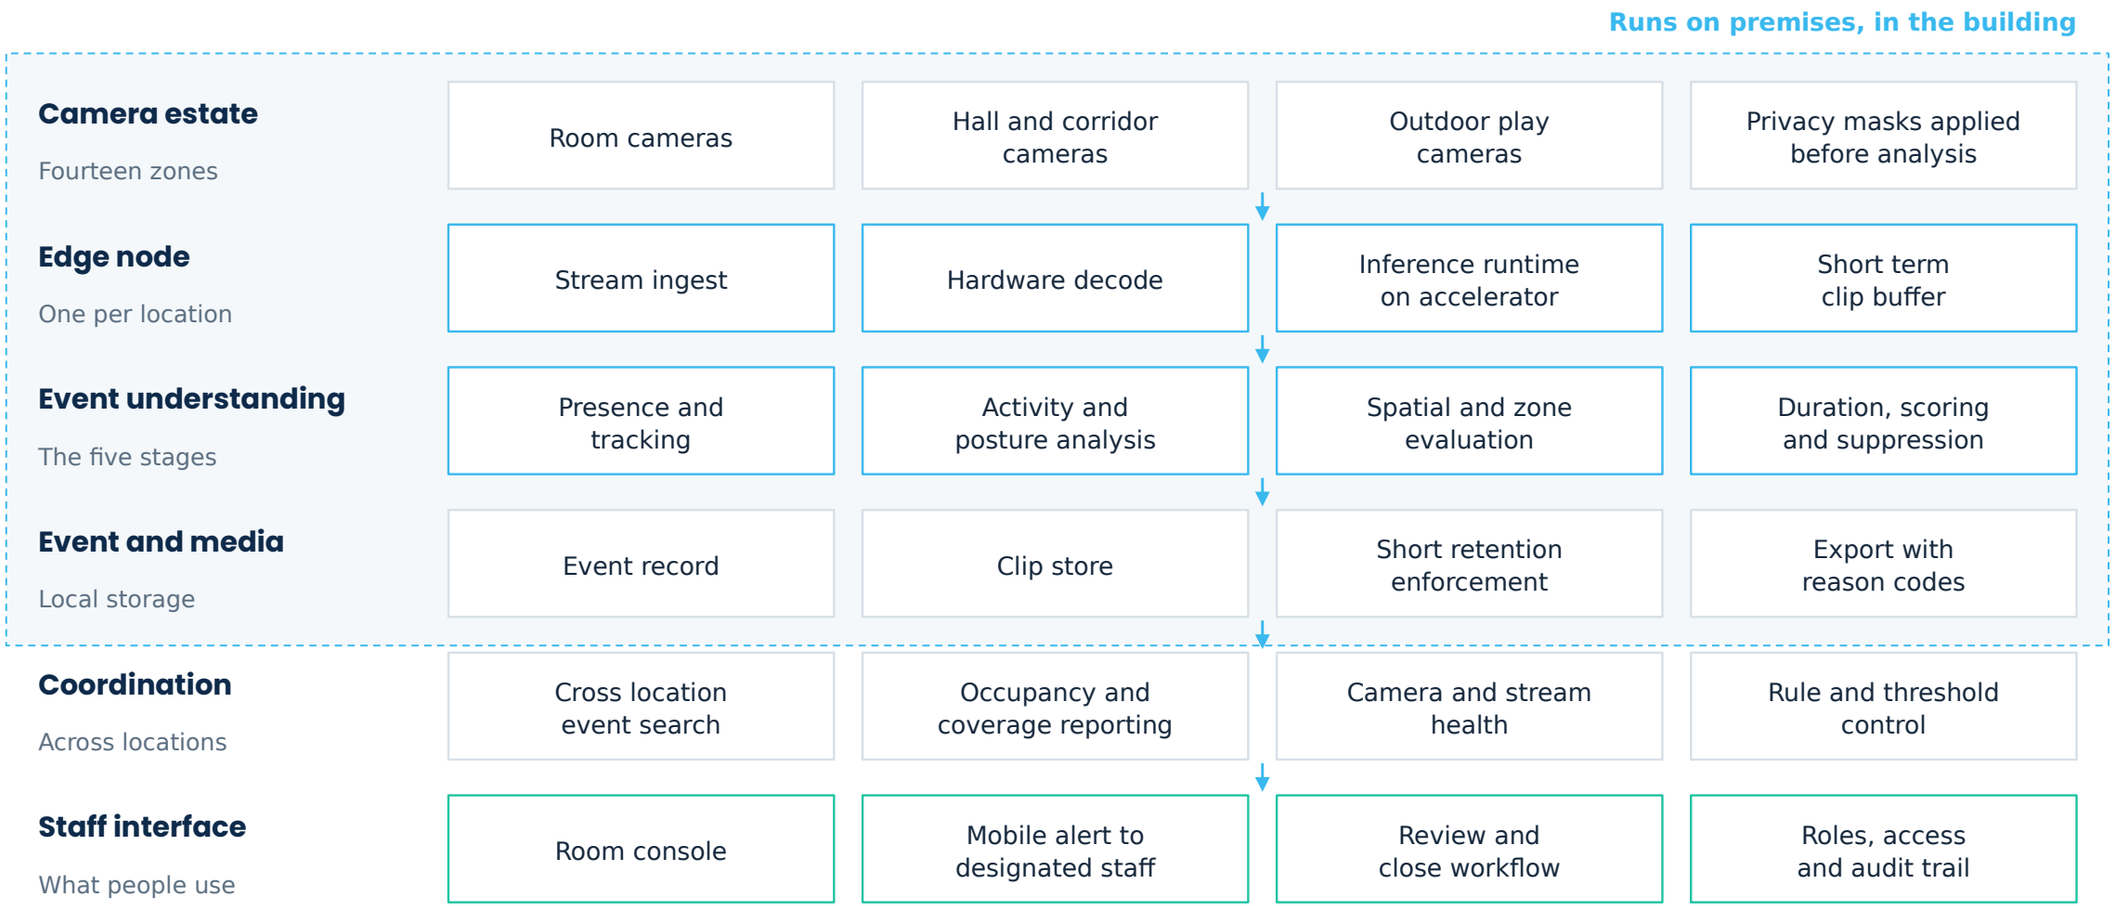

RUNTIME DATA FLOW

How an alert reaches a member of staff

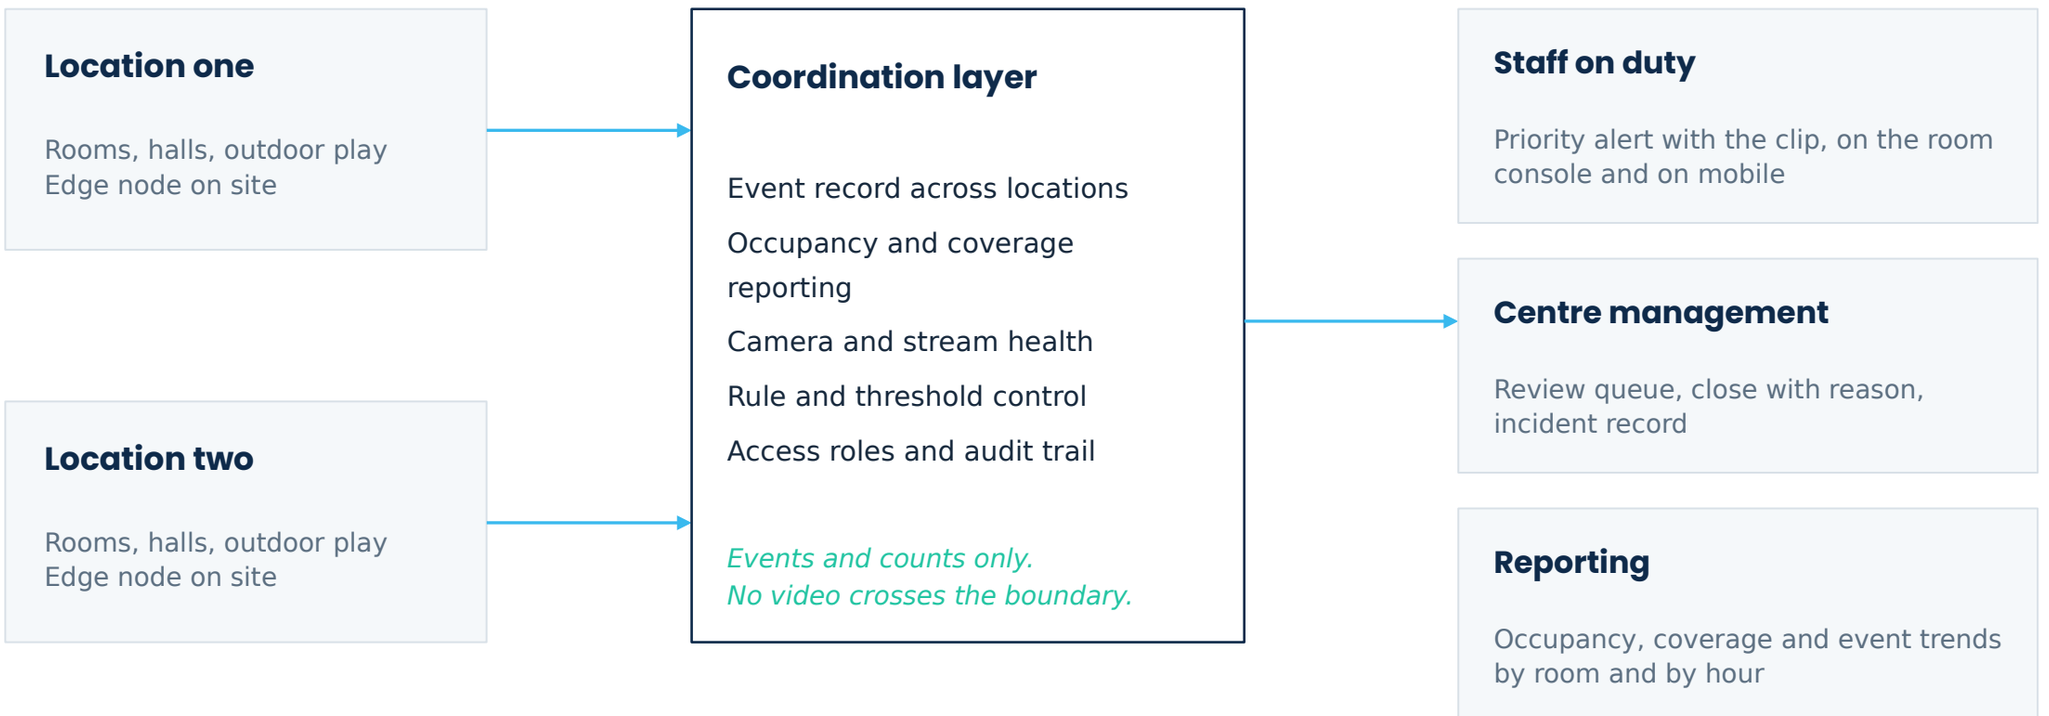

WORKED EXAMPLE

A child goes down beside a play frame at 10:22

This is the sequence as configured for the toddler room, timed end to end.

10:22:31	A tracked child undergoes a rapid change of posture close to a fixed play structure.
10:22:32	The activity stage sees a fall pattern rather than sitting down, based on movement across frames.
10:22:33	The spatial stage notes proximity to a fixed object and to one other child.
10:22:38	The duration rule for this room is five seconds on the floor without recovery. The child has not got up.
10:22:39	Signals are combined. The score clears the threshold set for the fall category in that zone.
10:22:41	A priority alert reaches the room console and the duty lead's phone, with a clip from six seconds before the event. Under fifteen seconds from the first frame.
10:22:55	An educator is with the child. The event is reviewed, closed with a reason code, and the record is available if a parent asks.

SCALE

What was actually deployed

2

Pilot locations under a single coordination layer

42

Cameras deployed and calibrated

14

Monitored zones across rooms and outdoor areas

40+

Concurrent AI video feeds analysed

15+

Safety and operational event categories

80 to 120

Children monitored across operating periods

6

Weeks of deployment and calibration

0

Frames of video sent outside the premises

PRIVACY AND SAFEGUARDING

A higher bar, designed in from the start

How the data is handled

All video and all inference stay on the premises. Nothing is processed off site.

Changing areas, toilets and staff rest areas are masked at ingest and never analysed.

Clip retention is deliberately short and enforced automatically.

Alerts carry a clip to the designated staff member only, never to a general distribution.

How access is controlled

Role based access, scoped to specific rooms and zones.

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

Exports require a reason code and are attributable to a named user.

Rule and threshold changes are versioned, so the configuration has a history.

No identity, no profiles, no history of a child

The system counts and tracks within a scene only for the length of an event, and it has no way to recognise a face, build a profile or carry a child's history from one day to the next. What gets recorded is a moment in a room rather than anything about a particular child.

DELIVERY

Six weeks, in five stages

STAGE	WHAT WE DID	WHAT CAME OUT OF IT
Week 1	Walkthrough of both locations, zone definition, privacy exclusion mapping agreed with the operator and signed off before any camera went live	Fourteen monitored zones defined, exclusions fixed in writing
Week 2	Camera placement and calibration for the angles that safety events actually need, which are not the angles security cameras usually take	42 cameras placed, all feeds ingested
Weeks 3 to 4	Event categories configured per zone, with durations and thresholds set for the age group using that room	The full category set running in observation mode
Week 5	Calibration against real room activity. Toddler rooms generate very different motion from a baby room, and the thresholds had to reflect that	Alert volume brought to a level staff could act on without fatigue
Week 6	Staff briefing, review and close workflow, escalation routing, handover to centre management	Staff running the system, with management owning the rules

OUTCOMES

Measured after calibration

MEASURE	BEFORE	AFTER
Feeds under continuous analysis	None	More than 40
Monitored zones	0	14
Defined safety event categories	0	More than 15
Retrospective manual review effort	Baseline	Around 75 percent lower
Time to alert on a priority event	Not applicable	Under 15 seconds on high confidence events
Coverage of operating hours	Spot checks	Continuous
Video leaving the premises	Not applicable	None

The result in three numbers

15s

to reach a staff member
on a priority event

75%

less retrospective
video review effort

15+

safety event categories
under continuous watch

And not one frame of video left the building.

APPLICABILITY

Where the same approach carries across

The pattern suits any supervised environment where the people responsible are occupied with the people they are responsible for.

Early years and daycare

Falls, rough contact, unattended zones, exit approach, ratio by visual occupancy

Elder care and assisted living

Falls, prolonged floor presence, room left unattended, wandering toward exits

Paediatric clinical waiting areas

Falls, crowding, unattended children, blocked emergency access

Schools and after school care

Playground incidents, corridor crowding, restricted area entry, after hours presence

Special needs settings

Distress movement patterns, self endangerment, separation from the group

Indoor play and activity centres

Collisions, crowding on equipment, entry to closed areas, capacity thresholds

BEFORE YOU START

What decides whether this works for you

What we would ask you

Which rooms and outdoor areas are in scope, and which must be excluded outright?

What is your consent position with parents and with staff, and who signs it off?

What does your regulator require you to record, report and retain?

Who receives an alert during the day, and on what device?

Who owns the thresholds after handover, and reviews them each term?

Where these projects go wrong

Camera angles chosen for security, which are wrong for reading posture and movement.

Thresholds set once, so staff learn to dismiss alerts and stop looking.

Treating a flag as a finding, which damages trust with staff faster than anything else.

Skipping the consent conversation with parents and staff until after go live.

Expecting the system to catch everything, when what it really does is reduce how much gets missed.

We calibrate against the actual room before enabling a category. A rule tuned on a baby room will misfire in a toddler room, and a system that misfires is a system staff will learn to ignore.

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 operator's own pre deployment baseline for the same activities, over comparable periods.

The system detects visual events and patterns that may indicate a safety incident. It does not diagnose injury, assess a child's condition, or determine intent. Every flagged event is reviewed by a qualified person, and the clinical or safeguarding judgement is always theirs.

Detection capability depends on camera placement, resolution, frame rate, lighting and the calibration carried out for each room. Alert timing figures apply to high confidence events under normal network conditions. No observation system is complete or infallible. This system supports supervision by trained staff. It does not replace it.

Deployment in any childcare setting is subject to the operator's own safeguarding policy, consent arrangements with parents and staff, and the requirements of the applicable regulator. Product names, internal component names and implementation detail are omitted deliberately, and can be discussed under a mutual non disclosure agreement.