

# Why generic analytics fall short

**Paul DiBerardino, Regional Sales Manager for Salient Systems' Southeast Region, looks at purpose-built AI for correctional facilities**

**T**he call comes in at 2:14am. A fight has broken out in a housing unit. By the time the responding officer reaches the scene, it's over. Two individuals are being separated. A third is on the ground.

Now the shift supervisor must piece together what happened, who threw the first punch, how long it lasted, whether anyone else was involved. That means pulling footage, scrubbing through the timeline, writing a narrative account from camera angles that weren't designed to answer those questions. It's 4am before the incident report is done. This is the routine that purpose-built AI video analytics is starting to change but only when the system is actually designed for corrections in the first place.

In corrections, operators are watching for altercations in progress, prohibited items, individuals who've gone down and aren't moving, fire and behaviors that indicate escalation before something happens.

Generic video analytics platforms, even good ones, are designed around the use cases their largest customer segments need. Corrections are specialized verticals. The behavioral

patterns that matter, the event types that require immediate response, the reporting formats that work inside a facility; these aren't default configurations. They're purpose-built requirements.

**“A corrections-specific platform needs to reliably detect altercations between individuals, not just motion anomalies.”**

That gap is why most facilities running AI analytics today are either running the wrong tool for the environment or spending significant time customizing a generic platform to approximate what a purpose-built solution would do natively.

**What purpose-built actually means in this context**

When we talk about purpose-built AI for corrections, we mean systems that are trained and designed around the specific behavioral signatures that matter in custodial environments,

**“A solution that requires hardware replacement isn’t a solution; it’s a capital project that may not see approval for two or three years.”**

not repurposed from retail loss prevention or logistics.

That distinction plays out in detection capability. A corrections-specific platform needs to reliably detect altercations between individuals, not just motion anomalies. It needs to distinguish a fall or man-down event from someone sitting on the floor. It needs to flag prohibited behaviors and unauthorized activity with enough precision that operators aren’t spending half their shift clearing false positives.

It also means the alert experience must match how operators actually

work. An alert that fires with a timestamp and a camera number creates work, the operator has to pull the clip, review it and write up what they saw. An alert that fires with a natural-language description; the number of individuals involved, what they were wearing, a summary of how the incident unfolded, changes the operational dynamic entirely. End-of-shift reporting goes from a burdensome reconstruction exercise to a documentation task. That matters in facilities where staff-to-coverage ratios are already stretched.

Consider what that shift means in practice. An officer who previously ▶



**Paul DiBerardino,**  
Regional Sales Manager  
for Salient



spent 45 minutes reconstructing a single incident from raw footage can instead review a structured summary, verify it against the clip and close the report. Multiply that across a full shift in a busy facility and the time savings are real, not as an abstract efficiency metric, but as hours returned to active monitoring and response. That's the difference between a system that records what happened and one that actually helps staff manage what's happening.

None of that happens automatically. Even a purpose-built platform delivers its best results when it's configured for the specific facility; camera placement, zone definitions, event thresholds calibrated to the population and layout, alert routing mapped to how staff actually respond. That configuration work is where a lot of deployments fall short, not because the technology isn't capable, but because it was stood up without enough attention to how the facility operates. A good integration partner brings that operational context to the setup, not just the technical one and the right integration makes the infrastructure question far simpler than most facilities expect.

### The integration question: hardware and software together

One of the practical barriers to deploying any new analytics capability in corrections is infrastructure. These facilities often run a mix of IP and analog cameras across aging facilities and budget cycles don't move fast. A solution that requires hardware replacement isn't a solution; it's a capital project that may not see approval for two or three years.

The integration model has to account for that reality. For Salient, the goal has always been to make CompleteView a platform that extends to meet what operators need without forcing a rip-and-replace solution and that the partners we bring into the ecosystem are held to the same standard.

VisionCoreAI's WardenEye One is an example of what that looks like in practice. It works with any existing IP or analog camera, so there's no hardware prerequisite. Connecting a Recording Server to begin camera analysis and event forwarding takes seconds. When an event occurs, CompleteView operators receive a structured alert – number of individuals involved, clothing descriptions, a plain-language summary of how the incident unfolded directly within the interface they're already using. No new screen to monitor, no parallel workflow to manage.

What makes that alert useful rather than just fast is the technology behind it. VisionCoreAI's platform uses end-to-end deep learning neural networks trained specifically for correctional environments going well beyond motion detection to deliver what the company describes as near-human comprehension of on-screen events. The difference is clear in the quality of the alert: specific, contextualized and actionable rather than a flag that still requires an operator to figure out what actually happened.

### What this means for corrections deployments now

The conversation at the facility level is changing. Administrators who've been through a serious incident with inadequate documentation or who've

faced oversight scrutiny over response times aren't asking whether they need better video technology. They're asking why they don't have it yet.

That shift affects how integrators approach this vertical. The facilities asking the sharpest questions now are the ones that have already learned something the hard way. They want specifics: what events does the system detect, what does the alert actually tell the operator, how does it fit into the existing infrastructure. A general answer about analytics capability doesn't satisfy that conversation. A tested, purpose-built integration does.

What's different today is that the barrier to that conversation is genuinely lower. A solution that works with existing cameras, connects in minutes and surfaces alerts inside the VMS operators already know removes most of the friction that used to slow corrections deployments down. The infrastructure argument, the training argument, the budget argument are all smaller than they were.

**“The facilities asking the sharpest questions now are the ones that have already learned something the hard way.”**

The facilities that act on this now build something that compounds over time: a documentation record, an incident baseline, an operational posture that's hard to replicate quickly. In corrections, that advantage isn't measured in efficiency gains. It's measured in what doesn't happen. ■

