MINING · OT & PROCESS CONTROL · VIDEO ANALYTICS

The Eyes Your Control System Doesn't Have

Why video analytics is the perception layer for predictive maintenance and the last line of safety — and why it never scales by putting video on the internet

AUDIENCE	FOCUS	REGION
Plant, control/OT, reliability & cyber	Scaling video analytics into the control loop	LATAM

Contents

01

THE MANAGER WHO WANTED TO GIVE THE SYSTEM EYES

A plant manager we worked with runs a mature APC[1] and the plant control system. He was not chasing a new dashboard. He put it to me in a sentence: he wanted to give his control system eyes. His loop, however well tuned, optimizes the variables it can measure — temperatures, pressures, flows — and is blind to everything else: a belt starting to track off, a bearing housing two degrees above its neighbours, a spill forming under a chute, a person stepping into a zone that should be empty. It has no context for the situations that actually take a plant down.

That is the gap. The industry is racing to put AI at the centre of the operation — autonomy, optimization, production — while the layer that decides whether the operation stays safe and whether the assets survive lags behind. Maintenance and safety perception move slower than operational AI, and the asymmetry is the risk.

KEY INSIGHT — A control system, however advanced, is blind to the physical context around it by design. Video analytics is how you give it eyes, without asking the control loop to do anything it was never built for.

02

WHAT THE CONTROL SYSTEM CANNOT SEE

Advanced process control is a numerical instrument. It reasons over sensor streams and acts within a model of the process. What it does not do is see. A recent engineering study of process-integrated computer vision in a steel rolling mill put it plainly: sensor-based systems operate on time-series data and cannot visually assess what is happening in plain sight — a surface defect, an out-of-spec material, the exact place where a failure begins to form.[2] For

[1] APC — Advanced Process Control, the model-based layer that optimizes a process within its measured variables.
[2] The study reports that visual anomalies — surface defects, material problems, the spatial location of a developing failure — often precede any measurable change in a process parameter.

a VP of operations the consequence is simple: the problem is almost always visible before it shows up in a trend, and right now no one in the room is watching for it.

Video analytics fills exactly that hole. It turns cameras the plant already owns into a sensing layer that sees what the instruments miss, and feeds that perception back to the people and systems that act. This is not prettier surveillance: it is the control loop gaining the situational awareness the loop structurally lacks, catching a failure while it is still cheap to fix.

03

THE SAME CAMERA, TWO JOBS

A camera with analytics does two jobs at once, and both move the business result.

Predictive maintenance. Heat and wear show on camera before they show in the trend. Peer-reviewed work on thermal imaging of conveyor idlers detects overheating rollers with high reliability — the classic heat-before-failure signature — and practitioners report that thermal anomalies often surface days or even weeks ahead of mechanical failure.[3] The failure modes a two-hour human round misses — a misaligned belt, an out-of-spec ore size chewing a crusher liner, a leaking gland — are precisely what continuous vision catches. One case documented by an analytics vendor (Razor Labs), on an iron-ore port-conveyor fleet — using vibration sensors rather than cameras — put the number on the table: roughly 800 hours of unscheduled downtime a year, valued at tens of millions, against failures detected more than three months ahead.

People safety. In an operation where presence has been pulled back from the field, the AI camera becomes the first line of defence, not a backstop. Mobile-equipment interaction is the single largest killer in mining: ICMM members reported it as the leading cause of fatalities, and the EMESRT round table attributes 30 to 40 percent of industry deaths to failures of vehicle-interaction controls. In Chile, the Mining Safety Regulation (DS 132 of the Ministry of Mining), enforced by SERNAGEOMIN, exists precisely to protect the life and physical integrity of those who work the operation; with fifteen fatalities reported by the service in 2025 — the highest toll since 2018 — the pressure to strengthen equipment-person interaction controls is not theoretical.[4] A camera that watches the exclusion zone every second, instead of every two-hour round, is what stands between an autonomous haul road and the person who should not be on it.

[3] The peer-reviewed idler thermography reports precision above 0.97 for detecting overheating rollers; the technical figure does not change the executive decision, which is to instrument continuous watching.

[4] SERNAGEOMIN — the National Service of Geology and Mining — is the body that applies and enforces Chile's Mining Safety Regulation, currently under update. DS 132 is not cited here for a specific proximity-detection article; it is invoked as the Chilean mining-safety framework within which equipment-person interaction risk is managed.

KEY INSIGHT — Where the people have been pulled back from the field, the AI camera is the first line of defence, not the last: it watches the exclusion zone every second, not every round, on the same optical assets that already do predictive maintenance.

04

WHY IT STAYS A PROOF OF CONCEPT

Almost every operation has run a video-analytics pilot. Almost none have scaled it. The reason is not the vision model: detection models are now the most mature part of the solution. The real walls are the same three every time:

- **Scalability:** a demo on one camera, in one corner of the plant, proves nothing about a hundred cameras streaming inference across the site.
- **Interoperability:** a detection that does not close the loop with the control or maintenance system creates no value — it creates noise.
- **Cybersecurity:** scaling without the right architecture exposes the OT[5] network to attack vectors that operations teams are not willing to accept, and rightly so.

A one-camera pilot proves none of those three. That is why it does not scale.

KEY INSIGHT — The vision model is not the bottleneck. Scalability, interoperability and cybersecurity are the three factors that decide whether a pilot becomes a program. There, and not in the model, is where the proof of concept dies.

[5] OT — Operational Technology, the control and instrumentation systems that run physical operations.

05

CYBERSECURITY: THE RISK THAT STALLS THE ROLLOUT

This is the factor operations and cybersecurity teams flag most consistently as the real barrier to scaling. The seemingly fastest path to "scale" a camera fleet is to push the streams to a cloud over the public internet. It is also the most direct way to hand an attacker the keys to the equipment you were trying to protect.

The evidence is not theoretical. The Mirai botnet hijacked more than 600,000 internet-exposed devices — most of them IP cameras and routers reachable through a short list of default passwords — and weaponized them to take down a major DNS provider, knocking a significant fraction of the internet offline. Claroty named the pattern directly in its 2025 state-of-exposure analysis: the single riskiest behaviour in industrial environments is connecting an OT device straight to the internet, and a large share of organizations still do it.

Here it is worth dismantling the objection we hear at every site: "I already have CCTV; I am not touching my OT network." The point is exactly the opposite. The video analytics we propose runs *on top of* the cameras and feeds the plant already has; it does not replace the network or the existing video infrastructure — it adds the inference layer and integrates it in a segmented way. What is non-negotiable is not the camera: it is the data path. A poorly segmented camera — new or legacy — is a pivot point onto the same network that runs the control system, and that is the scenario cybersecurity owners will not accept. It is also the argument that stops projects that technically work.

KEY INSIGHT — Video analytics runs on top of the CCTV you already have; it does not replace the OT network, it adds segmented inference. What is non-negotiable is the data path: an exposed camera is a lateral access vector into the control system, not a privacy problem.

06

CONNECTIVITY: THE ENABLER THAT GETS UNDERESTIMATED

The second structural wall is interoperability, and it is where value is won or lost. A detection that lights up a standalone dashboard and goes nowhere is noise. The value appears only when the detection closes the loop — into the control system, the APC, or the maintenance management system (CMMS) — so that a hot bearing becomes a work order and a person in an exclusion zone becomes an interlock, automatically. Detection without integration is a science project.

The inference architecture decides whether that integration is possible. Safety-critical analytics — collision, intrusion, proximity — must run at the **edge**, because the response is measured in tens of milliseconds and the data has to stay local. Maintenance analytics, which watch trends over hours and days, can **aggregate to a central on-premises server**. Both demand a network that is high-bandwidth, low-latency, and — the non-negotiable — private.

This is why operators choose private cellular over public links and Wi-Fi. Newmont moved to private 5G after finding it could not reliably connect more than two machines beyond a hundred metres on Wi-Fi, and ran its own cores and SIMs so only its own equipment could join the network. In LATAM the pattern is already on the ground: Nokia's private LTE at Antofagasta Minerals' Centinela operation in Chile was built to carry exactly this kind of low-latency, controlled traffic. These are precisely the operators this brief addresses — Codelco, BHP, Antofagasta Minerals — the ones already solving the private-connectivity layer on which video analytics can finally scale.

Where a photonic backbone such as NTT's IOWN[6] reaches, aggregated video and inference can behave as if local across long distances: its 2030 targets of 100x power efficiency, 125x capacity and a 200× improvement in latency point at a transport that scales without ever touching the public internet.

KEY INSIGHT — Not everything scales, and it is worth saying so: safety-critical analytics belong at the edge and maintenance can aggregate to a central on-premises server, but neither scales over the public internet. The enabler is a private network sized for the whole fleet, not for the demo.

[6] APN — All-Photonics Network, the end-to-end photonics architecture at the core of NTT's IOWN initiative.

07

WHAT TO VERIFY BEFORE THE NEXT PILOT

It is not the technical questions that stall these programs; what we have seen stall them is organizational. Before committing resources to another proof of concept, it is worth having an honest answer to these four:

1. **Who owns the budget, and who will fight for it?** Video analytics falls between OT, reliability and cybersecurity, and each assumes it belongs to another. We have seen pilots die not from the technology but because no clear owner defended the budget line when the cut came. Settle this before you start.
2. **Which group holds a veto, and why would it use one?** Cybersecurity can veto the architecture; operations can veto anything that touches the OT network. If you did not put cyber at the table on day one, you will hit the wall on day ninety. Pull their objection forward, not back.
3. **What happens to the legacy CCTV?** There is almost always an installed base of cameras and a network nobody wants to rebuild. The program that scales is the one that runs *on top of* that existing estate and segments it, not the one that asks to replace it. If the proposal opens with "let's swap the cameras," it has already lost operations.
4. **Does the detection close a loop, and is the network sized for the fleet?** If an alert does not become a work order or an interlock, you bought a dashboard, not a defence; and a hundred cameras is a qualitatively different bandwidth and latency problem than a single-camera pilot. Size for the operation you are building, not the pilot you are showing.

KEY INSIGHT — What kills these programs is rarely the vision model: it is failing to settle who pays, failing to neutralize the cyber veto early, and ignoring the CCTV that is already installed.

08

KEY INSIGHTS

The claims of this brief, each able to stand on its own.

1 You are asking the control loop to see something it never had sensors for. Video analytics gives it that perception using the cameras you already own, not by replacing them.

2 The same optical asset does two jobs that rarely share a budget: it anticipates mechanical failure and it guards people safety. Split them and you pay twice.

3 The safety frame — ICMM, EMESRT, and in Chile the DS 132 enforced by SERNAGEOMIN — puts equipment-person interaction at the centre as the risk to control. Continuous watching, not spaced rounds.

4 The vision model is the mature part. The pilot dies on scalability, interoperability and cybersecurity — and before any of those, on who owns the budget.

5 Connecting a camera to the public internet turns your video into a lateral access vector onto the control network. The data path, private and segmented, is non-negotiable.

6 The real interlocutor is the operator already building its private network — Codelco, BHP, Antofagasta Minerals — and it is on that layer, and only that layer, that video analytics scales.

09

SOURCES

The public sources behind this brief's load-bearing claims.

- Process-Integrated Computer Vision for Real-Time Failure Prediction in a Steel Rolling Mill — arXiv, 2025. *Sensor-based systems operate on time-series data and cannot visually assess surface defects or misalignments.*
- Automated identification of overheated belt-conveyor idlers in thermal images using CNN — Sensors (MDPI), 2022. *Thermal vision detects overheating idlers with high reliability.*
- Vehicle Interaction — 9 Layers of Defence — EMESRT. *30–40% of mining fatalities are attributed to vehicle-interaction control failures.*
- Mining Safety Regulation, Supreme Decree No. 132 of 2004, Ministry of Mining — Chile, enforced by SERNAGEOMIN. *The framework that protects the life and physical integrity of those who work the mining operation.*
- Inside the Infamous Mirai IoT Botnet: A Retrospective Analysis — Cloudflare, 2017. *600,000+ devices, mostly cameras, weaponised to take down a major DNS provider.*
- State of CPS Security: OT Exposures 2025 — Claroty (Team82), 2025. *Connecting an OT device directly to the internet is the single riskiest behaviour.*
- Newmont on how private 5G changes mining — RCR Wireless News, 2025. *Why a private network, not Wi-Fi or the public internet, for latency-critical equipment.*
- Nokia & Antofagasta Minerals private network, Minera Centinela, Chile — Computer Weekly, 2022. *LATAM private network carrying autonomous-fleet traffic.*
- Leading iron-ore company reduces unscheduled downtime of its port-conveyor fleet — Razor Labs case study. *~800 hours/year of unscheduled downtime, >US$70M/year, failures detected 3+ months ahead — via vibration sensing, not cameras.*

For more on this topic, contact: ***Isabel Torres****, Intelligent Industry & OT, NTT DATA* · Isabel.TorresManquelafquen@emeal.nttdata.com.