HYPOTENUSE
ANALYTICS

One Platform Three Signals
One Truth

Predict. Protect. Verify.

Technical White Paper - Version 2.0

Contents

Introduction

Bridges, construction sites, camera networks, and digital records all produce huge amounts of data every day. Today that data mostly sits in separate systems that don't talk to each other. An engineer looking at a bridge sensor has no easy way to check what a nearby camera saw at the same time. A security operator watching camera feeds has no way to know if a structure nearby is under stress. A fraud team checking a document has no way to confirm it against real sensor or camera evidence from the same site.

Hypotenuse Analytics unifies **Critical Infrastructure Intelligence** (structural health monitoring), **Surveillance Intelligence** (camera-based situational awareness), and the **Reality Trust Center** (deepfake and forged document detection) within a single AI-native platform. Built on a shared AI and data architecture, these three intelligence domains enable evidence from one source to be correlated, validated, and enriched using insights from the others.

Critical Infrastructure Intelligence is built to work from day one of construction through decades of operation. Across all three products, the platform uses AI models that are efficient enough to run on cheap, on-site hardware, and keeps sensitive data (raw footage, raw sensor readings, private documents) on-site by design rather than by policy. This document explains how the platform is built and how each part works.

> Our mission is to turn scattered sensor readings, camera footage, and documents into one clear, trustworthy answer to three questions that matter to anyone running a physical site:
>
> - Is this structure safe?
> - Is something unusual happening here?
> - Is this piece of evidence real?

Why an Integrated Platform Matters

Hypotenuse Analytics is not three separate products bundled together. It is one platform — Critical Infrastructure Intelligence, Surveillance Intelligence, and the Reality Trust Center — built on shared infrastructure.

Three individually correct systems can still miss the truth together. A structural sensor can be perfectly accurate and still have no idea who is standing near the asset. A camera system can track a person or vehicle correctly and have no way of knowing whether the inspection video submitted afterward is genuine. A deepfake detector can correctly flag a fabricated video and have nothing to check it against. Each product's blind spot is exactly the other two products' core strength — which is why they are built on one platform, for two distinct reasons.

First, the same customer needs all three, under the same rules.

An airport, a metro authority, or an industrial site needs to know if its structures are physically sound, what's happening on its cameras, and whether the records it relies on are genuine. These aren't three different buyers — they're three questions the same buyer already has. All three products fall under the same regulatory ground in India (DPDP, BIS, and emerging AI governance guidance), and for a customer running more than one, they sit inside the same Command Center rather than as separate, disconnected tools.

Second, each product is stronger because the other two exist.

Detecting a deepfake or a forged document on its own is no longer a differentiator — most serious competitors can do that. What none of them can do is check a flagged document or video against real sensor and camera evidence from the same site, because they don't own that physical sensing layer. We do. The same logic runs both ways: a structural anomaly is more trustworthy once camera and inspection records confirm it, and a camera alert is more trustworthy once it's checked against real sensor readings. Only the combination offers that cross-check.

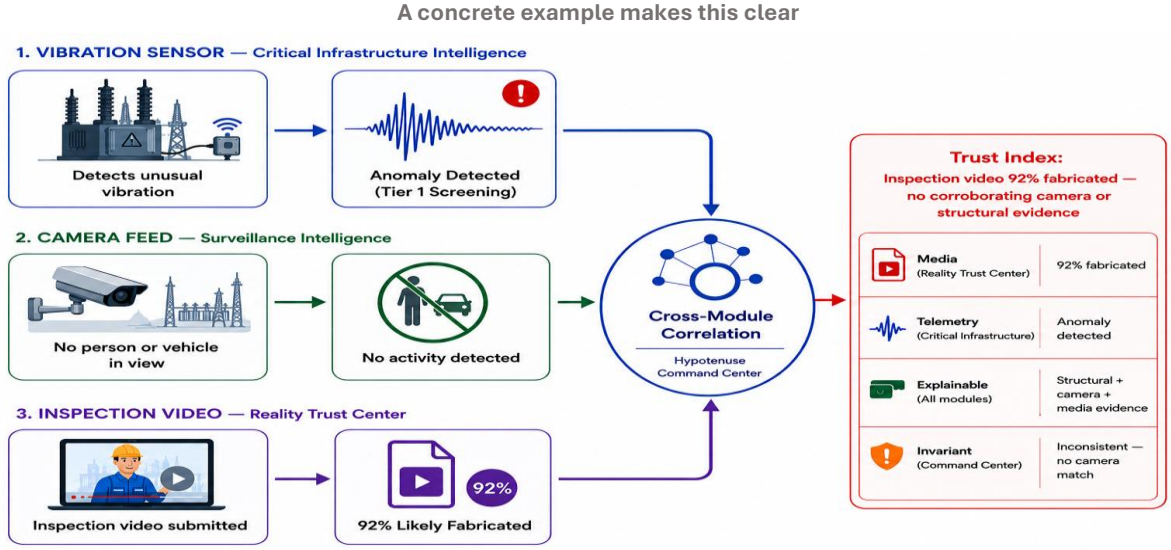

How the Platform Is Built

All three Products share the same four-step pipeline:

Signal

Raw data comes in: sensor readings, camera and drone footage, satellite data, or uploaded documents and media.

Context

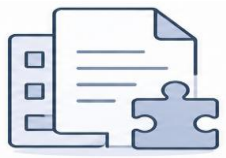

Data is cleaned, synced, and combined with other relevant info (build records, camera layout, etc).

Inference

AI models look for patterns: anomalies, suspicious behaviour, signs of manipulation.

Truth

Result is checked and explained — physics-checked, evidence-backed, not just a bare flag or score.

Hypotenuse Command Center

The Command Center is the single console where all of this comes together. It doesn't receive raw sensor readings, camera footage, or submitted media — only each product's *derived* output: a Factor of Safety from Critical Infrastructure Intelligence, a tracked path or behavioral flag from Surveillance Intelligence, a fabrication probability from the Reality Trust Center. Raw data stays local to each module by design.

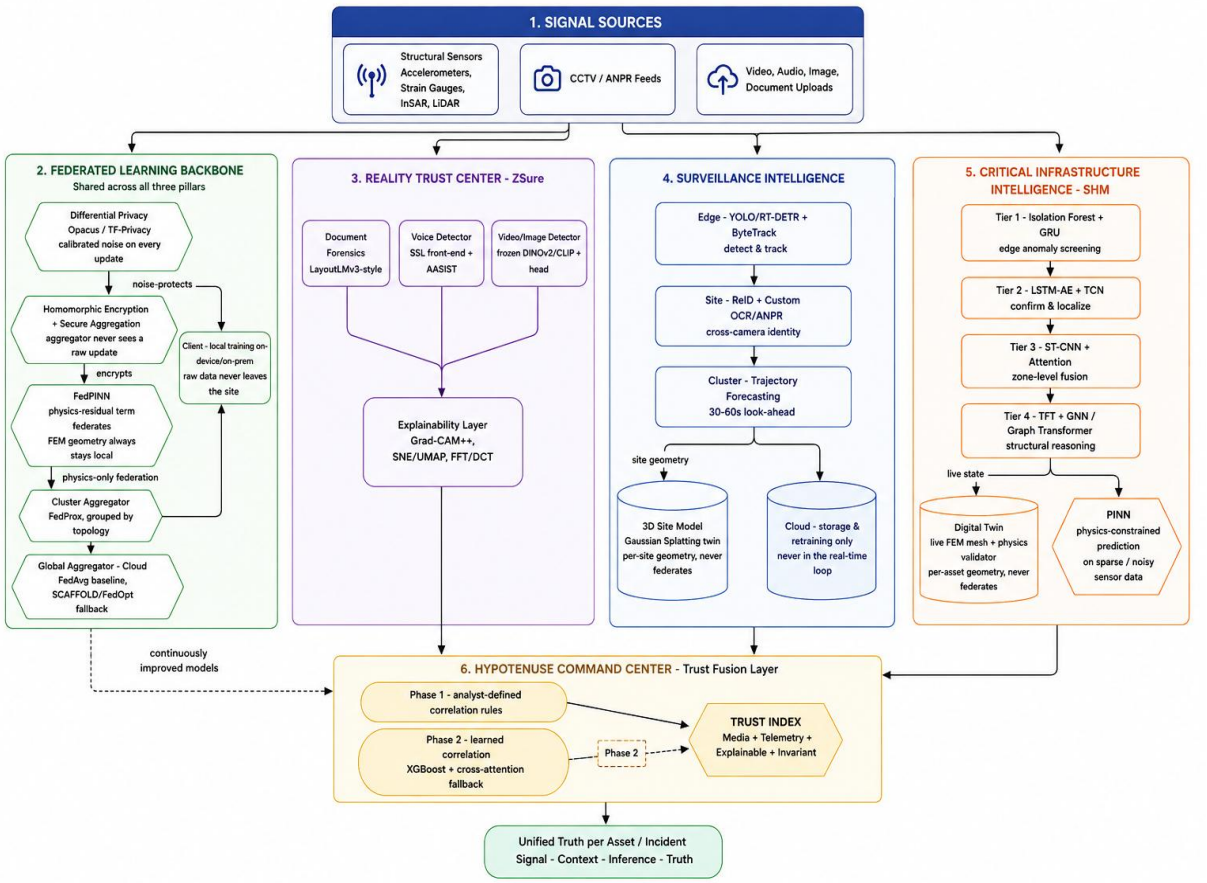

As the diagram shows, structural health monitoring, surveillance intelligence, and deepfake detection are not three separate systems reporting to three separate places — they converge into one Command Center.

Today, correlation runs on analyst-defined rules — for example, flagging a case when a structural anomaly and an absent camera detection occur in the same window. A learned correlation model is planned for a later phase, once enough real multi-module deployments exist to train it on.

The result of this correlation is the **Trust Index** — a live score per asset or incident built from four parts, each kept individually visible rather than blended into one number: **Media** (Reality Trust Center's fabrication probability), **Telemetry** (structural Factor of Safety and sensor confidence), **Explainable** (each module's own evidence layer), and **Invariant** (whether the three modules agree with each other — the newest part of the Command Center, and the one this section exists to explain).

Digital Twin

A Digital Twin is a live, constantly updated model of a physical site, built from real sensor readings, camera and drone data, and scans, not a 3D picture someone updates by hand once a year.

We use Digital Twins in these products:

Critical Infrastructure Intelligence — the twin represents the structure itself (a bridge, dam, or building). It is checked against real engineering physics, so it can answer questions like "what happens if we add more load here" honestly, instead of just showing a picture. It is kept accurate through periodic laser scans (LiDAR) compared against the original design, and it updates as construction changes the structure, for example when material is excavated or supports are added.

Surveillance Intelligence — the twin represents the site layout, built using LiDAR and drone imagery. It gives camera alerts a physical location, so an operator sees exactly where on the site something is happening, instead of just which camera caught it.

Both twins are kept accurate the same way: a scan-rescan cycle (LiDAR, drone imagery) corrects the geometry over time, while live sensor and camera data updates the twin's current state continuously in between scans.

The result is a twin that is useful for real decisions, not just for viewing. It flags impossible readings, meaning a prediction that breaks known physics, supports "what-if" planning, and gives every part of the platform a shared, trustworthy picture of the physical world to check things against.

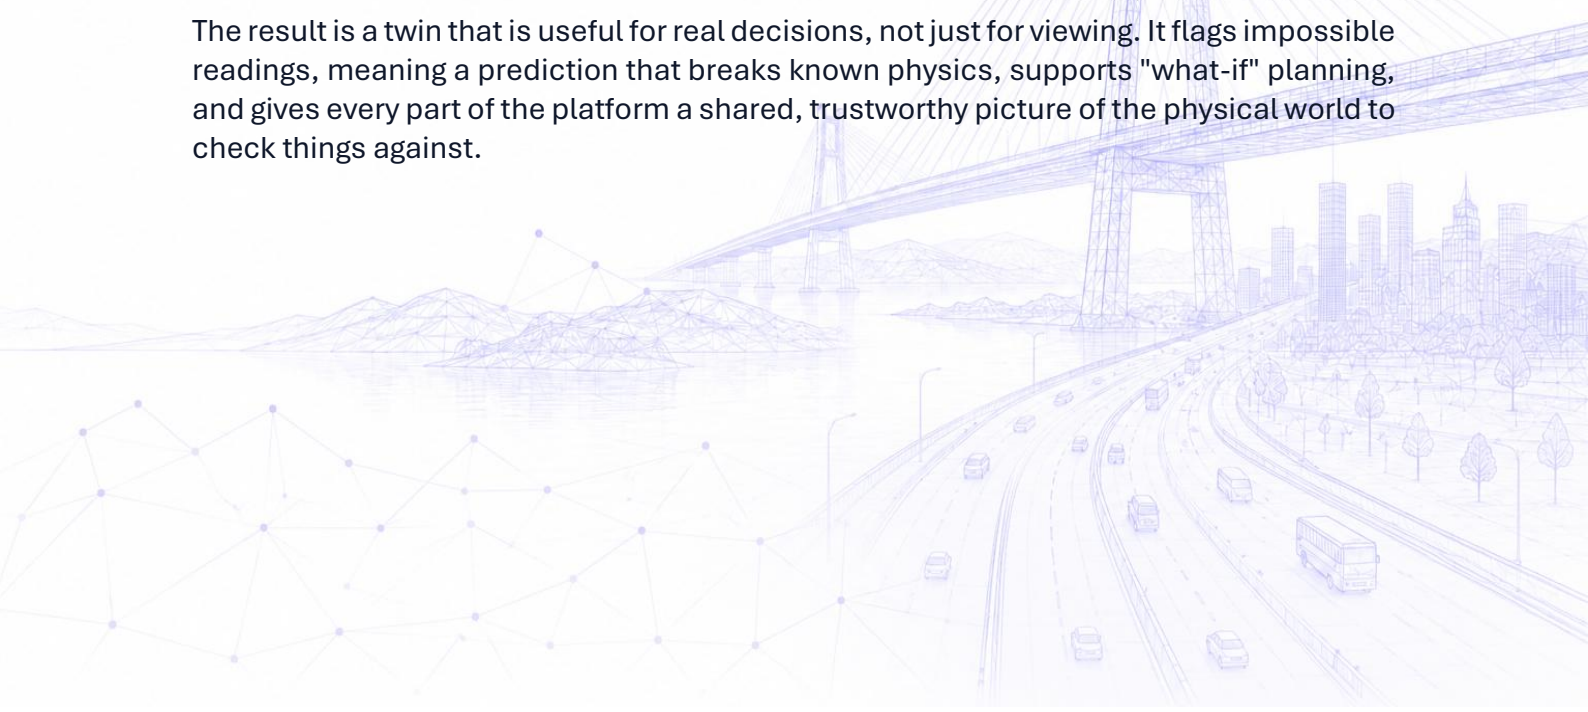

Architecture

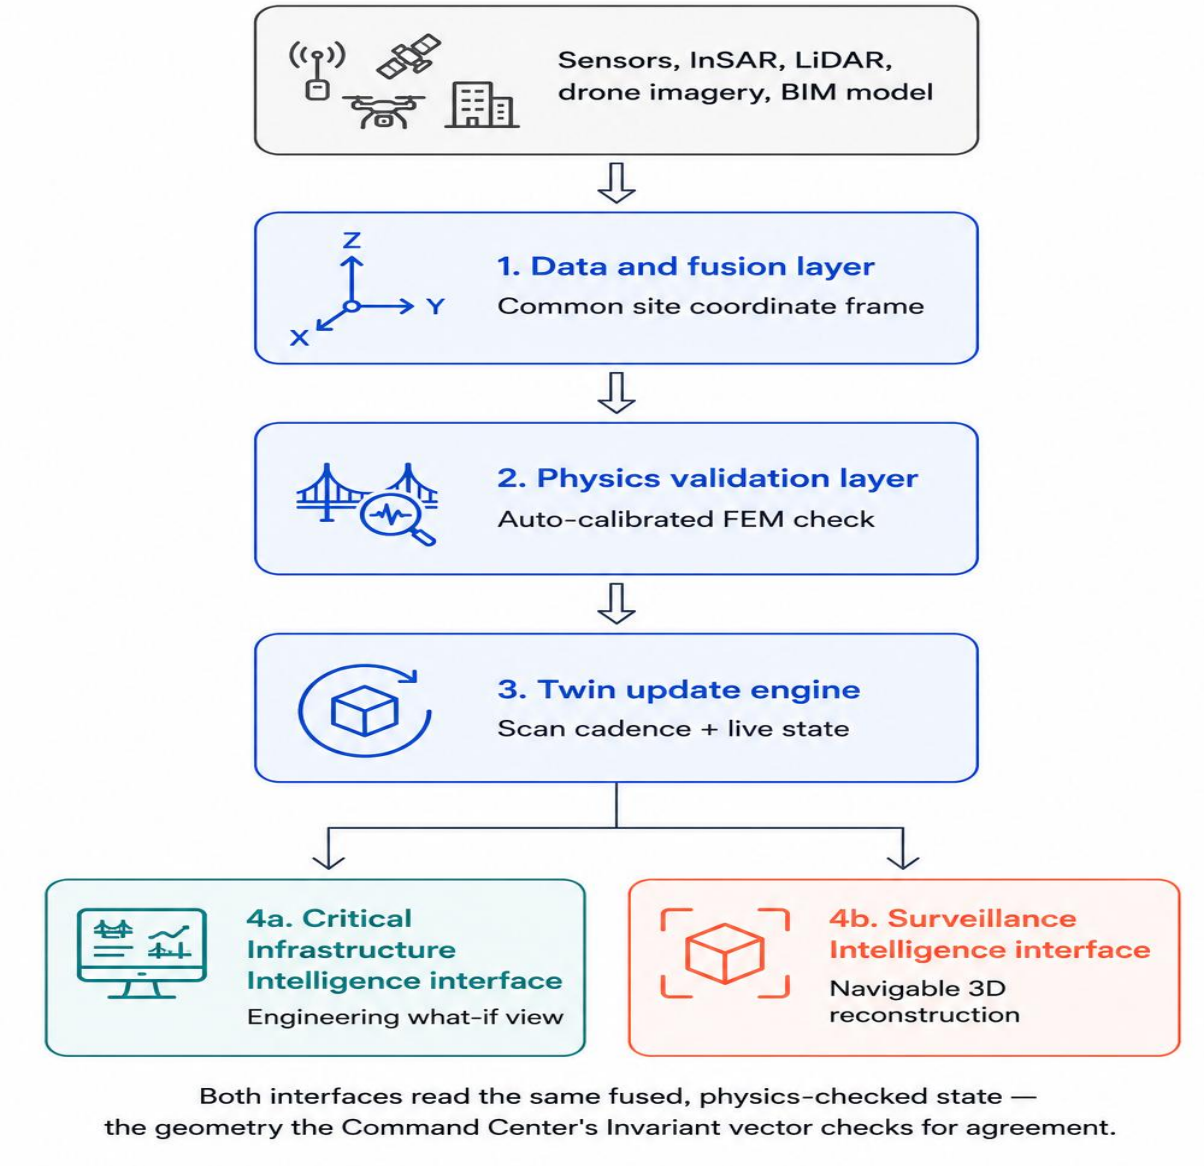

Underneath the visual model, the twin runs as four layers, not a single render job that happens once:

Data and fusion layer –

Every input that touches the twin — structural sensor streams, InSAR displacement, periodic LiDAR scans, drone imagery, and the BIM design model — is registered into one shared site coordinate frame before it is used. This is what lets a LiDAR scan, a drone photo, and a sensor reading all be compared against the same geometry instead of three separate references, and it's what makes an objective scan-to-BIM "built vs. planned" percentage possible rather than an eyeballed overlay.

Physics validation layer –

For Critical Infrastructure Intelligence, the twin is an auto-calibrated Finite Element Model (FEM) — a structural model continuously re-tuned against real sensor and scan data instead of fixed at design time. Every reading is checked against this model before it's surfaced, which is what lets the twin flag a prediction as physically impossible rather than just displaying it.

Update engine — two clocks.

Geometry and current state are kept accurate on different cadences. The twin's shape only needs correcting on the scan-rescan cycle (LiDAR weekly or monthly, drone imagery periodically), since a site's physical geometry doesn't change quickly. Its current state is updated continuously from live sensor and camera data in between scans. Splitting these two rates is what keeps the twin accurate without re-scanning a site every time a single reading comes in.

Interface layer.

The same fused, physics-checked state is exposed two ways depending on the product. Critical Infrastructure Intelligence exposes it as an engineering surface that runs what-if load simulation against the FEM. Surveillance Intelligence exposes it as a navigable 3D reconstruction of the site — built from LiDAR mapping and drone photogrammetry, rendered with Gaussian Splatting for real-time performance rather than the slower NeRF approach — so an operator can place a camera alert inside real spatial context. Because both interfaces read from the same underlying state, a structural anomaly and a camera detection for the same location are always describing the same physical point — which is also the geometry the Command Center's Invariant vector checks for agreement.

This four-layer design keeps the twin from ever running ahead of what the data can support: geometry only moves when a scan confirms it, structural state only moves when sensors and physics agree, and neither the engineering surface nor the 3D reconstruction can display a state that has not cleared validation. That is what separates it from a conventional 3D model — a live, physically-constrained record of the asset, not a picture that merely looks current.

Critical Infrastructure Intelligence

Critical Infrastructure Intelligence monitors the physical health of structures such as bridges, dams, buildings, tunnels, railways, and industrial sites, from the day construction starts through decades of use. It combines vibration, strain, displacement, and temperature sensor readings with two data sources that most infrastructure-monitoring setups leave largely untapped: satellite radar (InSAR) and laser scanning (LiDAR).

InSAR detects millimetre-level ground and structure movement over a wide area, catching slow-building problems, such as a pier settling gradually, that ground sensors alone would miss. **LiDAR** is a periodic laser scan (weekly or monthly) that produces a precise 3D point cloud, compared against the previous scan or the original design to show exactly what has moved and how much has actually been built. Together with live sensor readings, this same scan-rescan cycle is what keeps the Digital Twin's shape and current state accurate over time

Running one large, highly accurate model on every sensor everywhere does not work in practice. It draws too much power and runs too slow once hundreds of sensors are involved on ordinary site connectivity. Instead, the work is split by how urgent and how heavy it is. Small, efficient models run constantly on the sensor itself, just to screen for anything unusual. Slightly heavier models are only triggered once something needs a closer look, confirming and pinpointing a flagged anomaly.

A mid-level model reads patterns across a cluster of sensors together, without needing a full round trip to the cloud. And the most demanding, cross-structure reasoning happens centrally, where full computing power is available, to understand how stress in one part of a structure affects the parts connected to it. Every model here is judged not only on accuracy, but on whether it can actually run continuously in the field, including how fast it responds, how much power it draws, and how it behaves when a sensor briefly drops out.

The result for a customer is a live **Factor of Safety** (how much capacity a structure has left before its load limit), a **Remaining Useful Life** estimate, and a maintenance schedule built around the structure's actual condition instead of a fixed calendar. The platform works with sensors a customer already has installed, so there is no need to rip out and replace existing equipment to start getting value.

Live Construction Intelligence

Most monitoring only starts once a structure is finished and handed over, which means the years where many structural problems actually begin, such as excavation, temporary loads, early settlement, and poor concrete curing, are never recorded at all. Live Construction Intelligence closes that gap by monitoring from day one of construction, not after.

A sensor reading a change in settlement, on its own, does not say why. It could be excavation moving into weaker ground, recent blasting, a delayed support installation, heavy rainfall, or a genuine structural issue. The Construction Context Engine automatically connects sensor readings with what was actually happening on site at that moment, including excavation progress, blasting records, support installation, concrete pours, weather, workforce and equipment activity, and the BIM design schedule. This means an engineer is not left manually digging through logs to explain an alert, since the system already links cause to effect.

As construction physically changes the structure, the physics model updates to match, so a reading is always checked against what the structure should be doing right now, given its current state, not against a fixed baseline from before construction began. The output includes a Site Health Score, a Structural Risk Score, live construction progress tracking, and a plain-language explanation alongside the engineering one, so a site manager without an engineering background can still understand and act on it.

Physics-Informed Neural Networks (PINN)

The next model being brought into Critical Infrastructure Intelligence is a Physics-Informed Neural Network, or PINN. A PINN is a machine learning model that is trained to obey known physical laws, such as how materials bend or transmit heat, so it stays accurate even when sensor data is sparse or noisy, instead of just pattern-matching from data alone.

Most AI models only learn from examples, so they can only guess at situations similar to what they've already seen. If sensor data is missing, delayed, or noisy, a normal model can produce a confident but wrong answer. A PINN is different because it already knows the basic rules a structure has to follow, so even with incomplete data it stays within what is physically possible, rather than inventing a result that a real structure could never actually produce. This matters most in exactly the situations where accuracy matters most: during bad weather, sensor faults, or early construction, when data is at its most unreliable.

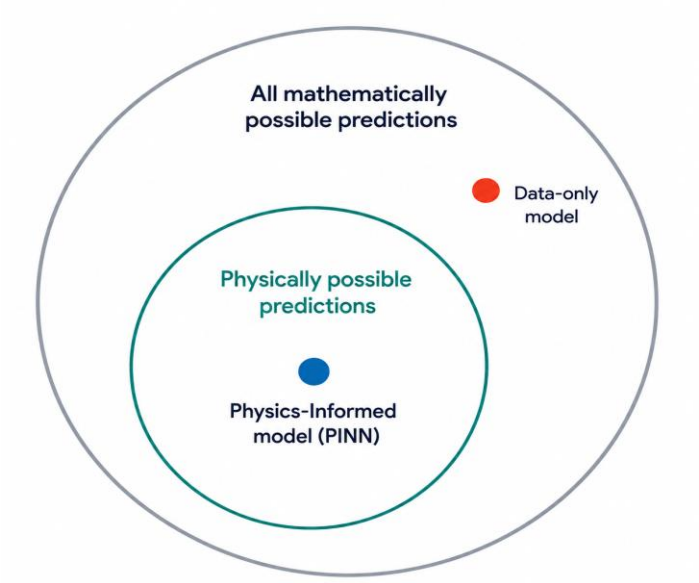

Conceptual illustration — a PINN's predictions are constrained to stay within what is physically possible

> Predict before failure, not after. Monitor from the first day of construction, not decades later. Trust every reading, because it's checked against physics, not just data.

Architecture

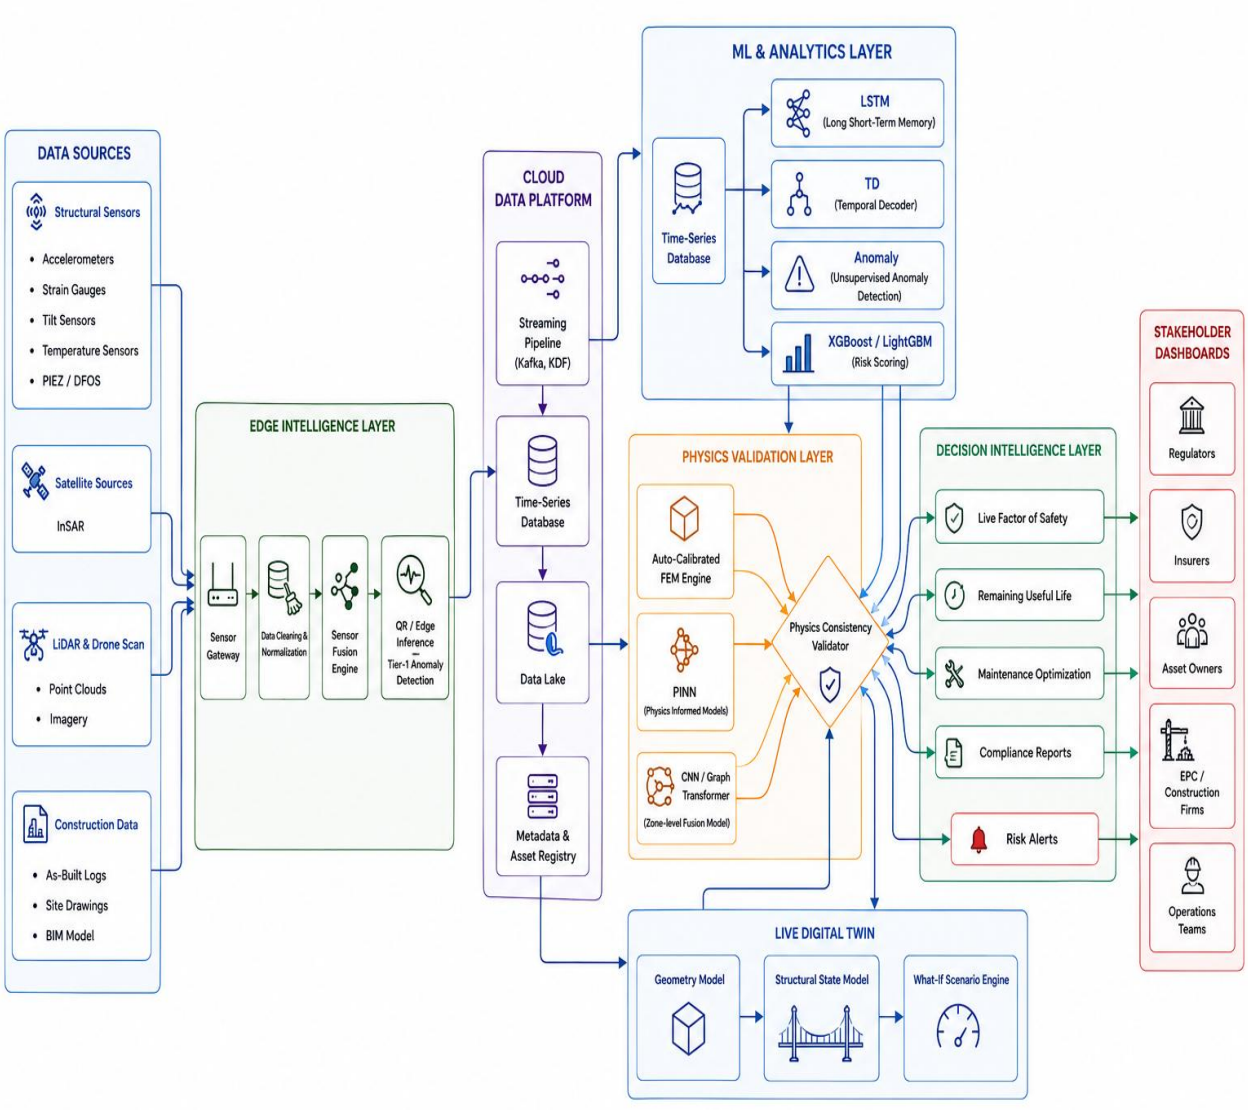

This architecture shows the full path from raw data to decision: sensors, InSAR, LiDAR, and construction records are fused and screened at the edge, then processed by the ML layer before every prediction is checked against the Physics Validation Layer — ensuring no output is surfaced unless it's physically possible. Validated results feed the Live Digital Twin and generate Factor of Safety, RUL, and maintenance outputs, which flow into role-based dashboards for regulators, insurers, owners, and operations teams.

Surveillance Intelligence

This product turns existing CCTV cameras into an active monitoring system, instead of a wall of feeds nobody can fully watch. It works with cameras a site already has installed.

Instead of only reacting to motion, which triggers constantly on wind, shadows, or stray animals, the system tracks where a person or vehicle is actually heading and predicts likely movement 30 to 60 seconds ahead. It follows someone across multiple cameras as one continuous path, and it only raises an alert based on how long and in what context something is happening, for example loitering for ten minutes, rather than a single odd frame.

Every clip is indexed the moment it is captured, allowing operators to instantly search by event, person, vehicle, or time instead of manually reviewing hours of footage. The system automatically generates a plain-language incident timeline by combining footage from all relevant cameras. Around **80–90% of processing** happens on-site at the edge, enabling faster real-time alerts while keeping footage on the premises. Every alert is assigned to a responsible person and closed with a logged timestamp, ensuring complete accountability and a reliable audit trail.

When something needs action, the system can recommend a response, such as raising an alert, requesting drone deployment, or triggering a lockdown workflow, but it never executes a physical action on its own. A human operator always has to approve and carry out anything that affects the physical site.

High-speed ANPR and vehicle tracking

The platform reads vehicle number plates at highway speed, even with motion blur, low light, an angled camera, rain, or partial obstruction, conditions where most standard ANPR systems fail. Once read, a vehicle's plate, type, colour, make, and movement history can be searched across every connected camera, which matters for logistics hubs and city streets as much as for highways.

High-speed ANPR reading a plate under motion blur and low light — conditions where standard ANPR typically fails

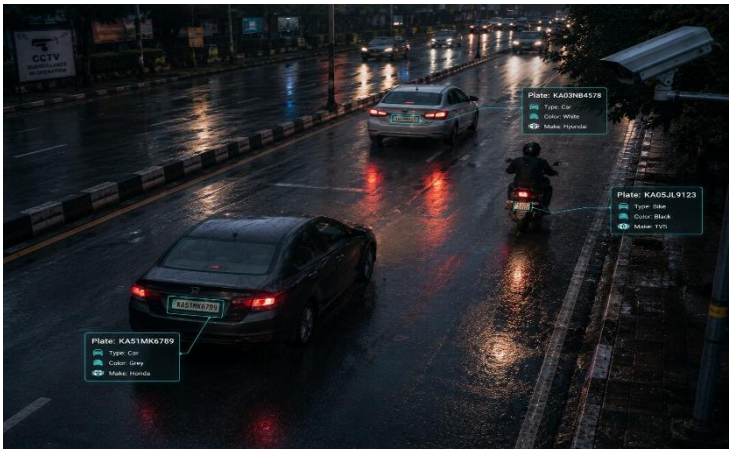

On-site hardware

The AI models run on standard edge-compute hardware already common in this kind of deployment, such as NVIDIA Jetson, Intel OpenVINO-based units, or Edge TPUs, compressed and optimised so they run efficiently on that hardware rather than needing a server-grade GPU on-site. A proprietary Hypotenuse edge device is planned for sites that want a single vendor for both software and camera-side compute.

A 3D model of the site

For the next phase of the product, sites that have been mapped using LiDAR and drone photography, reconstructed with a fast rendering technique called Gaussian Splatting, will show detections inside a navigable 3D Digital Twin of the site instead of a flat grid of separate camera tiles, giving an operator a real sense of where an event is happening relative to everything around it.

Alerts and notifications go through tools a site already uses, such as WhatsApp, existing ticketing systems, and regional-language messages, instead of requiring operators to learn a new dashboard.

Architecture

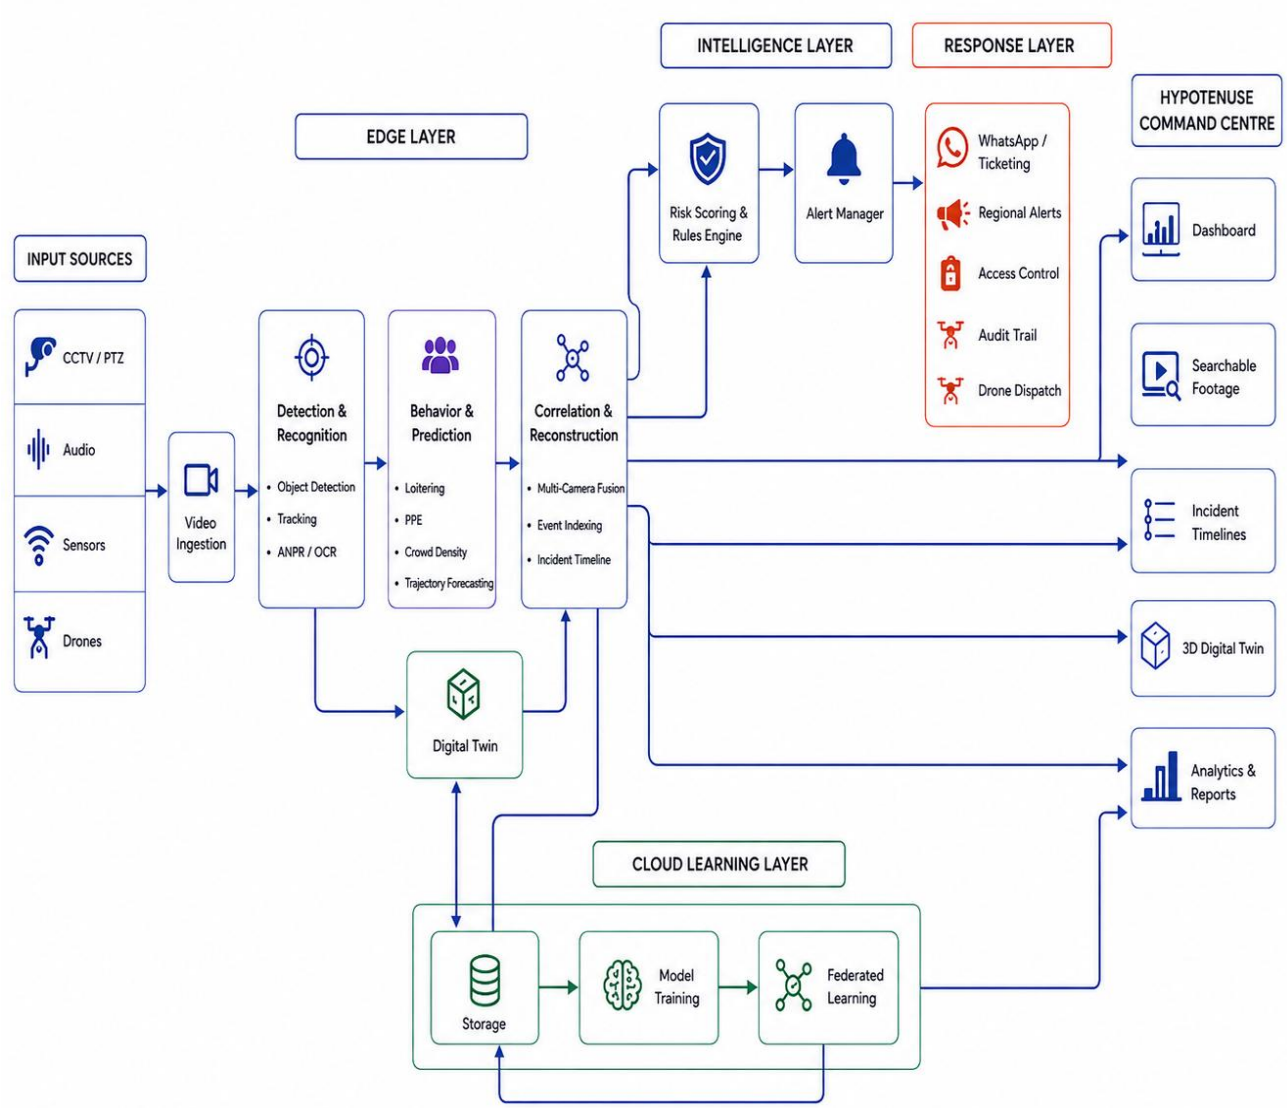

This architecture demonstrates how raw camera, audio, sensor, and drone feeds move from ingestion to action: the Edge Layer handles detection, tracking, ANPR/OCR, and behavior prediction (loitering, PPE, crowd density, trajectory forecasting) on-site, while correlating multiple feeds into one event timeline. From there, the Intelligence Layer scores risk and raises alerts, and the Response Layer routes them through WhatsApp, ticketing, regional-language messages, access control, or drone dispatch, all captured in an audit trail. In parallel, the Cloud Learning Layer retrains models and runs federated learning across sites without moving raw footage off-site. The result surfaces in the Hypotenuse Command Center as a live dashboard, searchable footage, incident timelines, a 3D Digital Twin, and analytics reports.

Reality Trust Center

Inside the Reality Trust Center is ZSure, Hypotenuse Analytics' deepfake and forged-document detection product. It runs standalone, including as a browser extension for checking something in the moment, and it also plugs into the Command Center, where a flagged document or video can be checked against real camera footage or sensor data for the same site or incident, something a standalone deepfake-detection company cannot do.

ZSure checks whether a video, voice recording, image, or document is real or AI-generated, and gives a probability score with visual evidence, not just a plain yes-or-no answer. That evidence matters, because for anything used in a fraud case or a compliance report, a bare label isn't enough on its own; a reviewer needs to see why the system reached its conclusion.

It is tested against detection methods the underlying AI models have never seen before, not just methods it was trained on, since a detector that only performs well on familiar fakes will fail the moment a new generation tool appears. It checks video, cloned voices, images, and documents (spliced text, altered records, fully AI-written statements) as distinct problems, since each requires different checks.

How results are tested

ZSure is tested against several public benchmark sets covering video deepfakes (DFDC, Celeb-DF, and a wide multi-generator set), and separate benchmark sets for document forgery and fully AI-written documents. The important test is not how well it does on fakes it was trained on, but on fakes generated by a method it has never seen before, since that's the realistic test of whether it will catch next month's new fake, not just last year's.

Why cross-generator testing matters

To show why this is hard, the same identity was generated using one prompt across four different AI platforms — Stable Diffusion XL, Adobe Firefly, Leonardo Diffusion, and Realistic Vision v6.

Sample from ZSure's internal test set. Even a single synthetic image preserves fine hair detail, realistic eye reflections, natural skin texture, and accurate facial proportions — making it visually indistinguishable from a real photo.

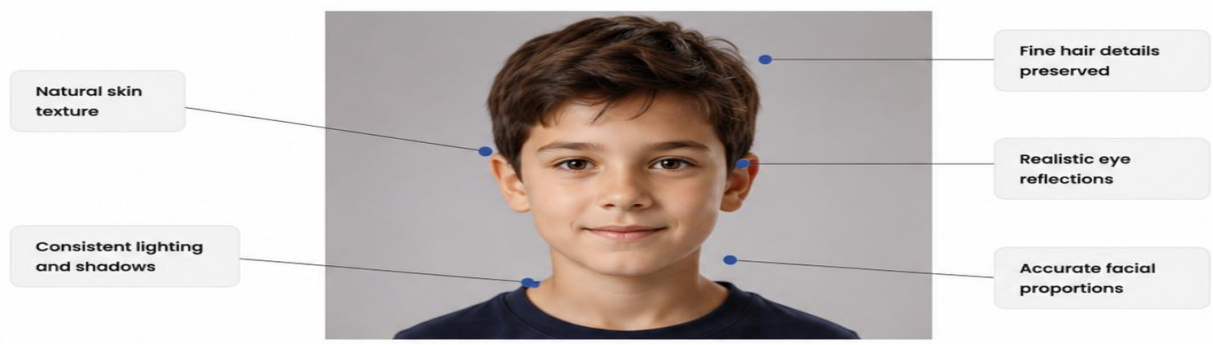

One real photo vs. four AI-generated versions, all from the same prompt

Showing the evidence, not just a score

Alongside every score, ZSure highlights the exact part of an image or video frame that drove the result, using a visual heat-map technique, and checks for content-authenticity metadata, an open industry standard called C2PA, where it's present. This is what turns a flagged result into something a fraud investigator or compliance officer can actually use as evidence, rather than a number they have to take on faith.

This evidence trail is what separates ZSure from a simple detector: a fraud team, an auditor, or a court doesn't just get told something is fake, they get shown exactly why, in a form they can act on, challenge, or present as proof. That is the difference between an alert and a defensible finding.

A bare probability score, on its own, is easy to dismiss. Anyone challenged on a flagged result can simply ask "how do you know," and a number alone has no good answer. Evidence changes that. When a specific pixel region, a metadata mismatch, or a document splice boundary is pointed to directly, the finding stands on something concrete, not just a model's confidence.

It also protects against a different failure mode: a detector that's simply wrong. If a result includes reasoning, an error is easier to catch and correct before it causes harm, since the evidence itself can be checked, questioned, and, if needed, overturned.

Architecture

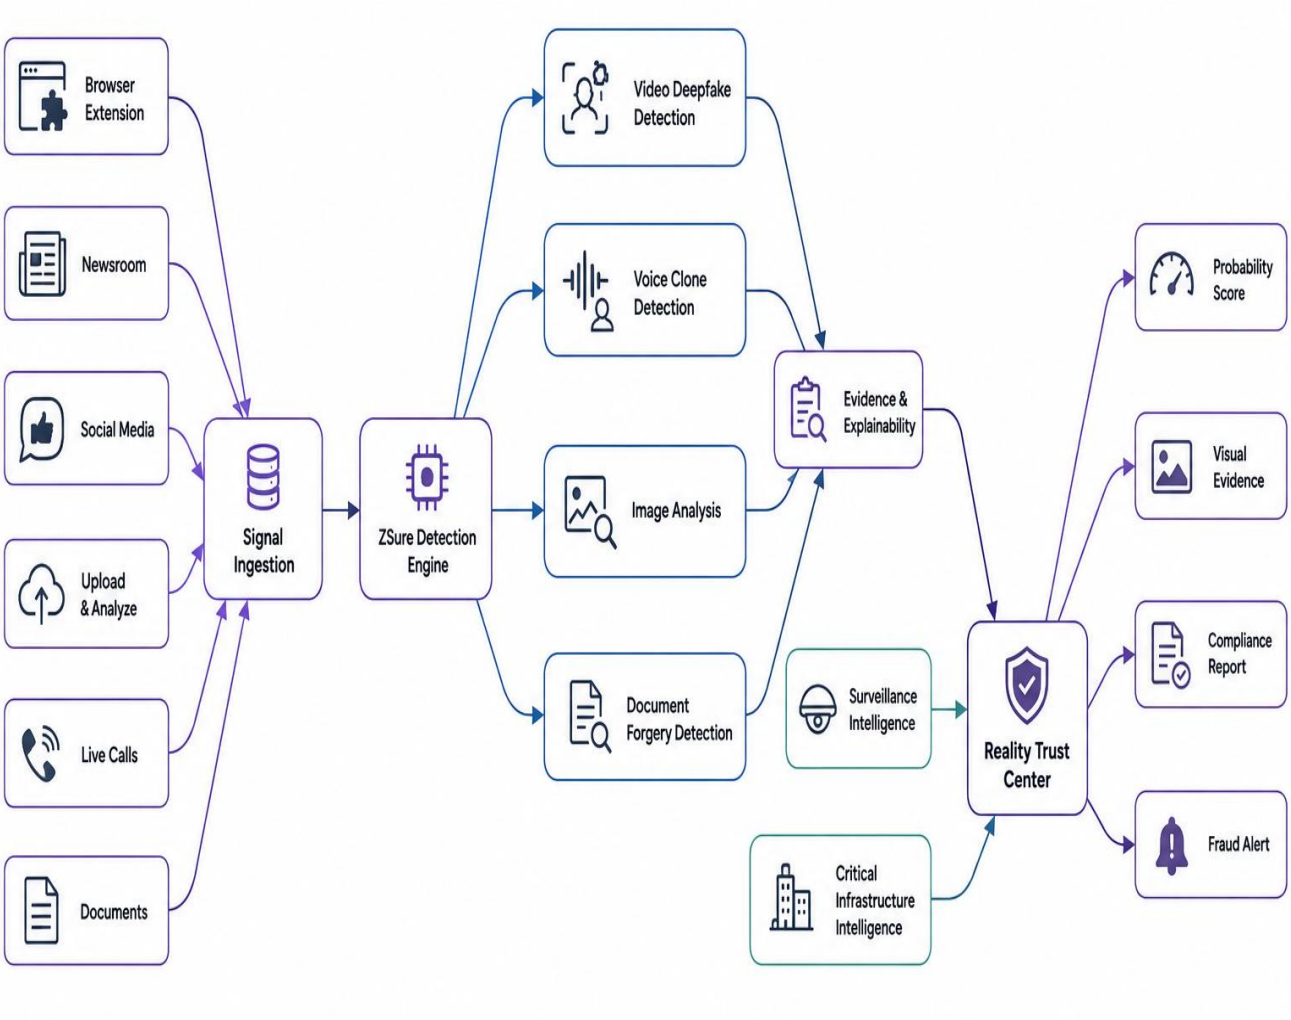

This architecture shows how the Reality Trust Center processes inputs from browser extensions, newsrooms, social media, uploads, live calls, and documents through a unified Signal Ingestion layer into the ZSure Detection Engine. Based on the content type, it performs video deepfake detection, voice clone detection, image analysis, or document forgery detection, with every result backed by Evidence & Explainability. The platform also correlates findings with Surveillance Intelligence and Critical Infrastructure Intelligence before generating a single evidence-backed outcome, including a probability score, visual evidence, a compliance report, and a fraud alert.

Federated Learning Across the Platform

Both products use the same basic idea to keep learning without moving private data around: instead of sending raw camera footage or private documents to a central server, only what the AI model has learned from that data is sent. The raw data stays where it was collected.

This matters for two reasons. First, some customers are contractually or legally required to keep their raw data on-site, particularly in government, defence, and critical-infrastructure deployments. For them, this is the only way their site can benefit from learning happening elsewhere. Second, it keeps footage and documents private by default, which is a significant part of how the platform meets Indian data protection requirements.

Here is what this looks like in each product

Surveillance Intelligence — each camera site learns from its own footage. Only the learned detection patterns, such as what loitering looks like or what a normal versus unusual vehicle path looks like, are shared. Raw footage never leaves the site.

Reality Trust Center — each deployment learns from the media and documents it checks. Only the learned signs of manipulation are shared, never the actual private images, videos, or documents that were checked.

Every site trains locally, and updates combine step by step, first across a few sites, then regionally, then occasionally into one global model, always sharing only learned patterns, never raw data.

Every update is checked before acceptance, so a bad camera feed or a deliberate attempt to bias the model cannot corrupt it. Camera and document updates are checked for consistency, and any site behaving unusually is flagged for review.

Rollout is staged: reliable federated learning across one customer's sites first, stronger privacy and multi-customer learning next, and fully encrypted learning for the most sensitive government and defence customers later.

Deployment and Security

The platform is designed around an edge-first architecture, with the cloud playing a deliberately narrow role: most processing happens on-site, and the cloud is used only for long-term storage and periodic model retraining, not for real-time decisions. This keeps the system fast, keeps ongoing bandwidth costs down, and works even where connectivity is limited, including designed support for low-connectivity sites using Starlink internationally and BSNL, Jio, and Airtel in India.

On-site processing

- Fast, real-time decisions with no round-trip delay to the cloud
- Works reliably even with limited or unreliable internet
- Lower ongoing bandwidth cost, since raw footage and sensor data don't need to stream continuously off-site

Cloud's role

- Long-term storage of records and footage
- Periodic retraining and improvement of models
- Company-wide reporting and analytics across sites

Security and accountability

- Every alert must be closed out by a named person, with that person's name and closing timestamp logged, so nothing is ever silently ignored
- Practices are built around India's data protection requirements (the DPDP Act and DPDP Rules), including purpose-limited processing, data minimisation, and retention controls

Conclusion

Bridges, cameras, and documents are usually managed as three separate problems. Hypotenuse Analytics treats them as one, because a structure's real condition, what's happening around it, and whether a piece of evidence is genuine all check against each other. Critical Infrastructure Intelligence, Surveillance Intelligence, and the Reality Trust Center run on one shared foundation, on-site, learning without moving private data around, giving a customer one dependable answer instead of three disconnected tools.

Predict structural failure before it happens. **Protect** every site, watched without losing the thread. **Verify** what's real, so no flagged deepfake or forged document ever has to stand alone. This is what one platform, built this way, makes possible.

CONTACT & CONNECT

INDIA OFFICE

Innov8 Graphix 2, sector 62, Ground Floor, Graphix Tower 2, A-13, Sector 62, Noida , Uttar Pradesh, India – 201301

info@hypotenuseanalytics.com

+91 63884 33626

FOLLOW HYPOTENUSE ANALYTICS

FOLLOW ZSURE®