Trustworthy Telemetry for AI: Securing Cellular–Connected Sensor Data with Zscaler Cellular



©2026 Zscaler, Inc. All rights reserved.

WHITEPAPER

Executive Summary

AI is only as trustworthy as the data it is allowed to learn from — and that data increasingly originates outside the enterprise perimeter. Across industries, organizations are embedding AI into their most consequential operational workflows: predicting equipment failures before they occur, optimizing inventory in real time, detecting fraud at the point of transaction, and managing public infrastructure at scale. In every one of these use cases, the AI is only as reliable as the telemetry feeding it. Corrupt, manipulated, or unverifiable sensor data doesn't just produce wrong answers — it produces confidently wrong answers, often without any visible signal of failure.

The Integrity Gap: The majority of this telemetry travels over cellular networks — SIM-connected sensors, edge devices, and machines that operate far beyond the reach of enterprise firewalls and endpoint management tools. Yet most organizations have no security policy layer at the SIM or carrier level. There is no firewall sitting between the sensor and the internet. There is no agent on the device enforcing destination rules. There is no per-device traffic log that a security team can produce when an auditor or incident responder asks: "Where did this data go, and did anything touch it in transit?" This is the cellular telemetry integrity gap — and as AI pipelines grow more autonomous, the consequences of leaving it unaddressed grow with them.

What Zscaler Cellular Does: Zscaler Cellular closes this gap by moving security enforcement to the layer where cellular telemetry originates: the SIM and the carrier network itself. Rather than relying on device-resident agents or perimeter controls that these devices never pass through, Zscaler Cellular applies Zero Trust policy at the point of data origin — before traffic reaches the enterprise network, and before it can be intercepted, manipulated, or misdirected in transit. The result is telemetry organizations can verify: per-device traffic logs, enforced destination allowlists, and real-time anomaly detection — without deploying a single endpoint agent. Security policy is enforced at the SIM and carrier level, ensuring that every data point entering an AI pipeline has a defensible, auditable chain of custody from source to system.

This paper is written for security and technology leaders: CIOs, CTOs, and CISOs — responsible for AI initiatives that depend on cellular-connected sensor data. It examines why cellular telemetry integrity is now a strategic AI risk issue, where traditional network security approaches fall short, and how Zscaler Cellular establishes a policy-governed security layer at the origin of the data. By the end, readers will be able to articulate the integrity gap in their own environments and evaluate the controls needed to close it.

©2026 Zscaler, Inc. All rights reserved.

Why Telemetry Integrity Is Now a Strategic AI Requirement

AI is no longer a back-office analytical tool. It is embedded in operational workflows — triggering maintenance orders, managing inventory, authorizing transactions, and automating responses to real-world events. What every one of these systems shares is a common upstream dependency: they act on telemetry from distributed sensors, machines, and edge devices operating far outside the enterprise perimeter.

This dependency reframes telemetry security as an AI risk issue. Compromised or unverifiable sensor data doesn't just create a network problem — it corrupts the inputs that AI systems trust, and by extension, the decisions those systems make.

How Modern AI Patterns Create Structural Exposure: Three AI patterns now dominate enterprise operational deployments — and each creates a specific risk when the underlying telemetry is compromised:

- **LLMs** fed stale or manipulated sensor data generate confidently wrong outputs with no internal signal of failure
- **RAG** pipelines retrieve and surface corrupted data as authoritative context — a falsified vibration reading, for example, can suppress a legitimate equipment fault alert that a human operator would have caught
- **Agentic AI** amplifies the damage further — automatically dispatching crews, issuing orders, or triggering safety responses based on inputs that were never legitimate

In each pattern, the model itself may be perfectly well-managed. The system is still vulnerable if the telemetry feeding it is not.

When Telemetry Is Compromised, AI Fails in Four Specific Ways

- Inaccurate or stale inputs drive wrong outputs — with no visible signal of failure inside the model
- Manipulated telemetry can bias outcomes, trigger unsafe automated actions, or mask a security incident in progress
- Unverifiable provenance makes it impossible to explain why the AI acted as it did — a critical failure in regulated environments or post-incident reviews
- Inconsistent availability creates blind operating states where AI silently degrades rather than escalating to human oversight

©2026 Zscaler, Inc. All rights reserved.

The Use Cases: Where Telemetry Becomes Business-Critical

The value of an AI system is inseparable from the reliability of the data it consumes. Across industries, the most consequential AI deployments today are not running on clean, controlled data pipelines — they are running on cellular telemetry from distributed sensors, machines, and edge devices operating in environments that enterprise security was never designed to reach. When that telemetry is compromised, the AI doesn't fail visibly. It continues operating — on corrupted inputs, producing decisions that appear legitimate until the consequences surface downstream.

The following use cases represent environments where cellular-connected sensor data directly informs AI-driven decisions. In each case, a telemetry integrity failure — whether through manipulation, interception, or fabrication — produces consequences that extend well beyond IT security

Use Case	Integrity Risk	Description
Predictive Maintenance	A replay attack — retransmitting a legitimate historical reading to mask a current fault — causes the AI to suppress a valid alert. The missed failure surfaces only when the equipment fails in the field.	Manufacturers, fleet operators, and logistics providers rely on AI to predict equipment failures before they occur, drawing on continuous telemetry from vibration sensors, engine diagnostics, and fault indicators across machinery, vehicles, and refrigeration units.
Intelligent Stocking	Spoofed or manipulated stock signals cause AI-driven replenishment to delay orders that should have been placed — producing measurable stockouts and lost revenue at scale.	Retailers, vending operators, and distributors use AI to manage edge inventory — relying on shelf sensors, POS signals, and machine-level telemetry to determine when to replenish or reorder. This is the last mile of inventory truth: the signal closest to the actual product and the actual customer
Finance and Payments Infrastructure	A masked hardware fault or suppressed tamper alert creates a window of undetected compromise. In a regulated environment, that is not just an operational failure — it is a control failure with direct compliance and liability implications.	Banks, ATM operators, and payment networks use AI to monitor device health, detect anomalies, and flag fraud across distributed terminal estates — drawing on availability telemetry, tamper alerts, and transaction connectivity signals.
Smart Cities and Public Infrastructure	SIM cloning, man-in-the-middle interception, or traffic misrouting to unauthorized endpoints can produce incorrect routing decisions, delayed emergency response, or suppressed safety alerts — at city scale.	Municipal operators deploy cellular sensors to manage traffic, monitor environments, and coordinate emergency response — feeding AI platforms that make real-time routing, safety, and resource allocation decisions. The consequence of compromise here extends beyond operational loss to public safety.

©2026 Zscaler, Inc. All rights reserved.

The Core Problem: Telemetry Often Traverses Links You Don't Control

Most cellular telemetry travels over networks the enterprise does not own, manage, or control. From the moment a sensor transmits a reading to the moment that data enters an AI pipeline, it passes through carrier infrastructure, shared routing paths, and aggregation endpoints that sit entirely outside the enterprise security boundary — with no firewall, no policy enforcement point, and no audit trail in between.

The Chain of Evidence Is Broken: Even when the sensor itself is legitimate, organizations typically cannot prove — end to end — that:

- The transmitting device is the device it claims to be
- The data was not altered in transit
- The data only reached authorized ingestion endpoints
- The communication pattern reflects expected, legitimate behavior

This is the cellular telemetry integrity gap. It is not a configuration failure or a missing patch. It is a structural visibility gap built into how cellular IoT connectivity works today — and most organizations have no control layer in place to close it.

Known Attack Classes That Exploit This Gap

The risks are not theoretical. Security researchers and threat intelligence teams have documented active exploitation of cellular IoT infrastructure across four specific attack patterns:

- **SIM cloning** — an adversary duplicates a device's SIM identity, impersonating a legitimate sensor while injecting false telemetry into the pipeline
- **Command-and-control (C2) over cellular** — a compromised device uses its cellular connection to beacon to an attacker-controlled endpoint, entirely outside perimeter visibility
- **Man-in-the-middle on unencrypted backhaul** — passive or active manipulation of telemetry data while in transit across carrier infrastructure
- **Telemetry replay attacks** — legitimate historical sensor readings retransmitted to mask a current fault condition or anomalous state

Each of these attack classes is executable today against devices that rely on standard cellular connectivity without Zero Trust enforcement at the network layer.

The Question Every CISO and Auditor Will Eventually Ask: Can you prove that the telemetry feeding your AI systems is legitimate, unaltered, and delivered only to authorized endpoints? If the answer is no — or even "probably" — then the decisions your AI is automating are built on a foundation you cannot defend.

©2026 Zscaler, Inc. All rights reserved.

Why Traditional Approaches Struggle (Especially at Scale)

Most organizations don't ignore cellular telemetry security — they address it with tools they already have. The problem isn't intent. It's that the tools were designed for a different environment.

The Standard Toolkit — and Where It Breaks

Approach	Where It Breaks
VPNs / private network extensions	Extend the attack surface by making sensors routable into the enterprise; lateral movement risk grows with every device added
Backhauling traffic for inspection	Introduces latency, adds cost, and still relies on the device behaving legitimately before traffic reaches the inspection point
Local firewalls / per-site stacks	Create per-site implementation variation that compounds into ungoverned drift across a distributed estate
Endpoint agents	Simply not viable on most constrained IoT and OT devices — the hardware cannot run them

Each of these approaches can work in a bounded, single-site deployment. The structural failure point is scale. At 500+ devices across 10 or more sites or carriers, the per-site management model breaks down. Every new site or carrier relationship introduces implementation variation. That variation accumulates as policy drift — and policy drift creates gaps that an auditor or adversary will find before the security team does.

At that scale, patching the perimeter is not a solution. The only sustainable model is one that enforces trust consistently at the connectivity layer — not reconstructed site by site, device by device.



©2026 Zscaler, Inc. All rights reserved.

The Solution: Zscaler Cellular for Trusted, Policy-Governed Telemetry

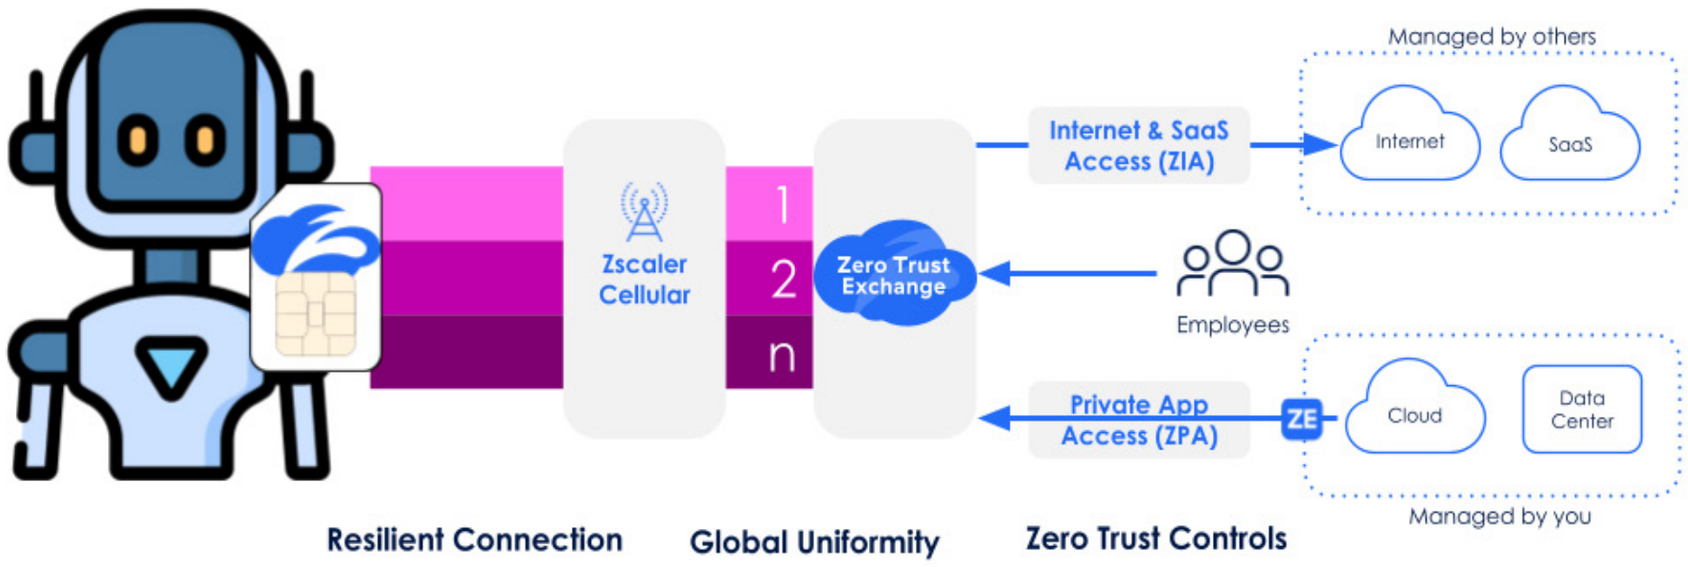

Figure 1: Architectural view of East-West traffic flow, illustrating how Zscaler Cellular bridges cellular-connected sensors directly to the Zero Trust Exchange.

Closing the cellular telemetry integrity gap requires enforcement at the layer where the gap exists — not at the enterprise perimeter, and not on the device. **Zscaler Cellular** extends Zero Trust policy enforcement to the cellular network itself, operating at the SIM and carrier layer to govern every device's traffic before it traverses the public internet or reaches an enterprise ingestion point. This is a fundamentally different architectural model from perimeter-based inspection or endpoint-resident agents — and that difference is what makes it applicable to the constrained, distributed IoT and OT environments where traditional controls have consistently fallen short. This architectural distinction has three practical implications:

- **No device agent required** — enforcement is in the network, not on the device. Constrained IoT and OT hardware that cannot run software agents is fully covered
- **No per-site infrastructure** — policy is enforced in the Zscaler Zero Trust Exchange, consistently, regardless of which carrier the device is connected to or where it is physically located
- **No reliance on perimeter inspection** — traffic is governed before it arrives at the enterprise boundary, not after

How It Works: Two Core Components

The Zscaler Cellular architecture is designed to bridge the visibility gap in distributed IoT environments by shifting security enforcement to the carrier layer. This approach creates a unified traffic model that handles both East-West traffic—the direct connection between sensors and the Zero Trust Exchange—and North-South traffic, which routes telemetry securely to its final destination. By centralizing this flow, we eliminate the need for per-site infrastructure and ensure that every telemetry packet is governed before entering the enterprise.

©2026 Zscaler, Inc. All rights reserved.

Zscaler Cellular operates through two integrated components:

- **Zscaler SIM (physical SIM or eSIM)** — functions as the device's secure gateway, automatically routing all outbound cellular traffic through the Zscaler Zero Trust Exchange for policy enforcement and inspection. No configuration change is required on the device itself.
- **Zscaler Cellular Edge** — an intelligent carrier-layer gateway that bridges mobile network traffic with the Zscaler security platform. It enforces granular, identity-based policy per device, isolates devices from each other to prevent lateral movement, and provides a per-device traffic log for every session.

Together, these components ensure that every device in the cellular estate — regardless of hardware capability, carrier, or location — is subject to the same enforceable Zero Trust policy, establishing the foundational controls required to maintain a defensible chain of custody for all telemetry data.

What This Means for Telemetry Integrity

For AI pipelines that depend on cellular sensor data, Zscaler Cellular directly addresses the three integrity requirements that traditional approaches cannot meet:requirements that traditional approaches cannot meet:

- Verified identity — each device's SIM identity is bound to a policy, making impersonation and SIM cloning detectable and blockable at the network layer
- Governed data paths — destination allowlisting ensures telemetry can only reach authorized ingestion endpoints; any deviation is blocked and logged
- Auditable chain of custody — per-device session logs provide the evidence trail that security teams, compliance teams, and incident responders need to answer: "Where did this data go, and did anything touch it in transit?"



©2026 Zscaler, Inc. All rights reserved.

Security Capabilities That Directly Improve Telemetry Integrity for AI

The preceding sections established a clear and specific problem: cellular telemetry traverses networks outside enterprise control, traditional security tools were not designed to govern it, and the AI systems that depend on it have no native mechanism to detect or correct for compromised inputs. The capabilities documented in this section address each dimension of that problem directly — not as compensating controls layered over an inadequate architecture, but as the functional outputs of Zero Trust enforcement applied at the cellular network layer.

1. Every Device Communicates Only With What It Is Permitted to Reach

Zscaler Cellular enforces least-privilege communication at the SIM level — each device is bound to a policy that defines exactly which endpoints, APIs, and brokers it is permitted to contact. Any traffic outside that policy is blocked before it traverses the network. This eliminates the most common telemetry integrity vectors: unauthorized data exfiltration, C2 beaconing, and lateral movement between devices on the same cellular estate.

2. Telemetry Flows Are Policy-Governed, Not Best-Effort

Traffic between a SIM-connected device and its destination application is routed through the Zscaler Zero Trust Exchange — regardless of whether that application sits in the cloud, a data center, or a partner environment. Telemetry does not travel over uncontrolled paths. It travels over a governed, inspected channel with consistent policy applied at every hop.

3. Anomalous and Unauthorized Traffic Is Detected and Acted On

Zscaler Cellular continuously monitors per-device traffic for anomalous destinations, unexpected communication patterns, and behaviors consistent with device compromise. When a device begins beaconing to an unauthorized endpoint or transmitting at an anomalous volume, the platform detects it and enforces a response — block, isolate, or alert — in real time. This is the visibility layer that organizations currently lack for their cellular IoT estates.

4. Full Coverage — Including Devices That Cannot Run Agents

This is the capability that distinguishes Zscaler Cellular from every endpoint-dependent security model: enforcement is in the network, not on the device. Constrained IoT and OT hardware — sensors, embedded controllers, legacy field equipment — receives the same policy enforcement as any other device in the estate, with no software deployment, no configuration change, and no dependency on device capability. There are no coverage gaps based on hardware limitations.

5. Unified Visibility Across the Entire Cellular Device Estate

Security and operations teams gain a single pane of glass across all SIM-connected devices — SIM activity, data usage, traffic destinations, and policy events — regardless of carrier, region, or device type. This is the audit capability that makes telemetry defensible: when a regulator or incident responder asks “where did this data go?”, the answer exists in a per-device log that is available, searchable, and complete.

6. Consistent Policy Enforcement at Global Scale

With coverage across 520+ carriers in 185+ countries, Zscaler Cellular enforces the same policy regardless of where a device is deployed or which carrier it is roaming on. Organizations with international device estates do not rebuild security per region or per carrier — the policy travels with the SIM.

©2026 Zscaler, Inc. All rights reserved.

Conclusion: Making AI Outcomes Defensible by Making Telemetry Trustworthy

The argument this paper has built is straightforward: AI systems are making consequential operational decisions, those decisions are only as reliable as the telemetry feeding them, and for the majority of distributed IoT and OT deployments, that telemetry currently travels over cellular networks without a verifiable chain of custody.

This is not a speculative risk. It is a structural gap — one that exists today, in production environments, across every industry deploying operational AI at scale. And it is a gap that traditional security architectures were not designed to close.

Zscaler Cellular addresses it at the layer where it exists: the cellular network itself. By enforcing Zero Trust policy at the SIM and carrier layer — governing every device's traffic before it reaches the enterprise boundary or an AI pipeline — it gives organizations something they have not previously had: the ability to verify where telemetry came from, confirm it was not altered in transit, and prove it only reached authorized destinations.

The outcome is not simply a stronger security posture. It is AI that organizations can defend — to auditors, to regulators, to boards, and to the operational teams whose decisions depend on it every day.

The question is not whether telemetry integrity matters. The question is how long your organization can afford to operate without it.

Next Steps

To learn more about how Zscaler Cellular can secure deliver integrity for your telemetry, request a personalized demonstration, contact a Zscaler expert today. Visit the Zscaler Cellular website or request a demo to begin your journey to a more secure and resilient support and operational system.

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange™ platform protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

© 2026 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at **zscaler.com/legal/trademarks** are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.

+1 408.533.0288 **Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134** **zscaler.com**

zscaler™

Act Fast.
Stay Secure.