■---
title: "When Independent Constraints Collapse in the Internet of Everything"
subtitle: "Why IoE Needs a Missing Control Layer for Continuous Identity, Policy, and Compute"
author: "Thomas Rocha III"
date: "December 2025"

When Independent Constraints Collapse in the Internet of Everything

Why IoE Needs a Missing Control Layer for Continuous Identity, Policy, and Compute

From fragmented orchestration to Single-Interaction Orchestration and Adaptive Routing (SIOAR/SSOAR-for-IoE)

Author: Thomas Rocha III
Date: December 2025

Abstract

Internet of Everything (IoE) systems were architected as compositions of independently governed layers: identity, transport, policy, orchestration, telemetry, and compute. This fragmentation was rational when constraints evolved independently. It is now failing because five structural pressures (accessibility mandates, Zero Trust validation, data sovereignty, AI-driven compute, and next-generation transport) collide within real-time interaction boundaries. Coordination complexity scales combinatorially; control mechanisms remain linear. The failure mode is not “insufficient tooling.” It is a missing architectural primitive: continuous interaction identity that persists across orchestration changes, policy decisions, and computational routing. Cross-vertical incidents throughout 2025 reveal this pathology. Cloudflare (November 18), Azure Front Door (October 29), Google Cloud IAM (June 12), Heroku (June 10), Microsoft Teams (December). Each experienced control-plane failures where configuration propagation cascaded globally because no interaction-scoped boundary contained the blast radius. The pattern is consistent: the pipes worked; the valve system failed. AI agent deployments expose the same gap. Breakdowns are governance failures, not model failures: context fragments at system boundaries, multi-step workflows collapse. Human operators were absorbing coordination costs the architecture cannot handle; when removed from the loop, the missing primitive becomes visible. Quantitative analysis projects fragmented coordination reaching ~10^16 surfaces by 2040, implying 90-100 GW continuous power for coordination overhead alone, roughly 2× current global data center capacity. Interaction-scoped architectures collapse this to ~10^11 surfaces: a 100,000× reduction that keeps 2040-scale IoE within buildable envelopes. The European Accessibility Act enforced June 28, 2025. Transitional provisions that grandfathered legacy systems are not a reprieve but a strategic trap: vendors can maintain existing fleets or compete for new revenue, but not both. Procurement gatekeeping excludes non-compliant vendors from 15-30% of enterprise ICT spend. The technical requirements (continuous synchronization under 200ms, selective delivery, multi-jurisdiction compliance) necessitate interaction-scoped governance by construction. The path forward is subtractive: collapse fragmented control contexts into unified interaction surfaces. Whether industry must build this capability or can deploy existing solutions will determine IoE’s viability at scale.

1. Executive Summary

1.0 Reader Map: How the Key Terms Fit Together

This paper is written for mixed audiences (operators, standards participants, procurement/legal, and finance). Several communities use different words for the same underlying unit. Here is the crosswalk used throughout: - Interaction (≈ “session”, “workflow”, “conversation”, “transaction”, or “control loop”): the end-to-end, real-time unit of behavior that must remain coherent while people, devices, networks, and services change. - Interaction identity: the durable identifier + state boundary for an interaction. If identity fragments, everything downstream becomes reconciliation. - Control plane: the decision/state layer that governs behavior (as distinct from the data plane that carries media/telemetry and the compute plane that executes workloads). - Orchestration: selecting and assembling participants, services, and resources during an interaction (including handoffs, fan-out, failover, and feature enablement). - Governance (policy enforcement): the constraints that must hold during the interaction (Zero Trust, accessibility/safety, residency/sovereignty, consent, auditability). - Adaptive routing (of transport and compute): dynamic selection of paths and processing locations inside the interaction’s governed boundary. SIOAR/SSOAR-for-IoE names the missing architectural constraint: a session/interaction-scoped control layer that keeps identity, governance, orchestration, and compute routing bound to one continuous interaction context.

1.1 The Thesis

Recent work in real-time communication architecture establishes a previously unarticulated requirement: systems lack a control layer capable of maintaining a single, continuous session identity across orchestration, policy enforcement, and computation (Rocha, December 2025). This requirement does not introduce a new feature or optimization; it defines a boundary condition that existing architectures implicitly violate under convergent pressures. The same analysis identifies a trade-off reversal. Fragmentation, once advantageous for modularity, isolation, and independent scaling, now imposes coordination costs that exceed its operational benefits as accessibility mandates, Zero Trust enforcement, data sovereignty, AI-driven computation, and next-generation transport requirements converge. What previously appeared as independent engineering concerns now collapse into a single architectural constraint: without unified session governance, real-time systems cannot scale coherently under contemporary regulatory and operational demands. This white paper generalizes the same architectural pathology across the Internet of Everything. Telecommunications, cloud infrastructure, healthcare, financial services, smart cities, and industrial automation exhibit the same failure class: fragmented control planes that cannot keep interaction identity, policy state, and compute decisions coherent as constraints collide. An IoE interaction spans device-to-device communication, device-to-cloud orchestration, human-to-AI collaboration, and service-to-service coordination. The missing layer is not RTC-specific; it is the cross-vertical requirement that becomes non-negotiable when regulatory deadlines, scaling limits, and AI automation arrive at the same time.

1.2 The Forcing Function: Independent Constraints Have Collapsed

Five independently evolving pressures now collide within the same interaction, and therefore cannot be satisfied by architectures that keep identity, policy, and computation in separate, asynchronously reconciled contexts: Accessibility and safety obligations: Accessibility has become real-time, not post-hoc. Captions, translation, assistive modalities, and safety constraints must stay synchronized through orchestration changes and device/network handoffs, under the same policy boundary as the primary interaction. Zero Trust as continuous validation: Trust decisions must occur throughout the interaction’s lifetime. Point-in-time authentication at a boundary is insufficient when posture, context, and policy can change mid-interaction. Data residency and sovereignty enforcement: Jurisdictional

boundaries must be enforced deterministically as data and compute move across regions and organizations, not inferred from logs after the fact. AI-driven computation and automation: AI features (routing optimization, security analytics, transcription/translation, agentic workflows) increase the number of interaction steps and boundary crossings. Without interaction-scoped governance, agents either fail on context loss or create shadow state that is expensive to reconcile and risky to audit. Next-generation transport and edge scale: 5G/Wi-Fi 7/satellite and low-latency edge reduce transport bottlenecks, making coordination overhead the visible limiter. Spatial computing adds new, latency-tight streams (depth, gaze, scene context, rendering control) that must remain governed in real time.

1.2.1 Total Conversation: Formal Definition of Accessibility Requirements

ETSI is developing technical specifications (including ES 204 009) to operationalize European Accessibility Act obligations for accessible communication services. The “Total Conversation” definition (ITU-T F.703) treats voice + real-time text + video as one synchronized conversational unit. Two implications matter architecturally. First, accessibility is defined as in-interaction synchronization, not after-the-fact augmentation. Second, the specification explicitly emphasizes interoperable services (avoiding proprietary dependency and user pre-registration), signaling that compliance pressure falls on architectural primitives, not product checklists. The technical needs are the same ones SIOAR targets: continuity of the interaction identity across orchestration/handoffs, policy coherence across all modalities, and selective delivery of assistive streams within the interaction’s governed boundary.

1.2.2 Accessibility as Procurement Gatekeeper

Accessibility enforcement operates primarily through market exclusion, not fines. In the United States, Section 508 requirements flow into federal procurement (via the FAR), making conformance a precondition for many government purchases. In the European Union, EN 301 549 plays the same role in public procurement. The mechanism is binary: vendors either demonstrate conformance and bid, or they do not and are screened out. Public sector and adjacent procurement commonly represent 15-30% of enterprise ICT revenue in major markets; exclusion from this segment is not a marginal compliance cost. Procurement evaluation increasingly demands operational proof under realistic conditions (network transitions, device handoffs, orchestration changes, sustained load), not only static documentation. Architectures that implement accessibility as bolt-on services outside the interaction’s governance boundary struggle to demonstrate deterministic behavior under these conditions.

1.2.3 Defense and Alliance Zero Trust Convergence

A parallel forcing function is emerging from defense and allied security requirements. In November 2025, the U.S. Department of Defense published “Zero Trust for Operational Technology Activities and Outcomes,” specifying 105 zero-trust activities for OT environments: 84 at target level, 21 at advanced level (DoD CIO, November 2025). The guidance explicitly mandates “continuous authentication and fine-grained policy enforcement” across seven pillars, including automation/orchestration and visibility/analytics. The architectural implications map directly to interaction-scoped governance: Continuous verification during resource access: The DoD OT framework requires trust decisions throughout an interaction’s lifetime, not only at connection establishment. Point-in-time authentication is explicitly insufficient; posture, context, and policy must be evaluated as conditions change mid-interaction. Separation of control and data planes: The guidance distinguishes policy decision points from data movement, requiring that governance travel with the interaction rather than being inferred from network topology. This is the session controller pattern: a control plane that binds identity, context, and policy across hops while workloads carry the payload. Time-sequenced capability outcomes: DoD’s “fan chart” roadmap structures controls as staged activities that attach to ongoing interactions. The implicit unit of governance is not the network segment or the device, but the interaction itself. NATO has ratified complementary standards. The Zero Trust Data Format (ZTDF), adopted in 2025 through NATO’s Combined Communications-Electronics Board (CCEB), embeds encryption, access policy, and classification metadata directly into data objects (archTIS, October 2025). Allied Communication Publication 240 (ACP 240), developed under Five Eyes sponsorship and

adopted by NATO and the U.S. Joint Chiefs of Staff, operationalizes this for coalition data sharing (Virtru, October 2025). The core principle: “protection and policy travel with the data wherever it goes.” Operation HIGHMAST, the UK’s 2025 global carrier strike group deployment, validated ACP 240 under operational conditions, demonstrating secure data sharing across multiple command boundaries and allied partners (Virtru, November 2025). The standard explicitly addresses the failure mode of network-centric security: classification determined by which network data resides on, rather than by attributes bound to the data itself. CISA’s updated Zero Trust Maturity Model (January 2025) and OT Asset Inventory Guidance (August 2025, joint with FBI, UK NCSC, and international partners) reinforce the same architectural direction for civilian critical infrastructure. The Connected Communities guidance emphasizes that “traditional perimeter-based security measures are no longer sufficient” for interconnected OT environments (CISA, September 2024). The procurement implications parallel accessibility. CMMC (Cybersecurity Maturity Model Certification) gates defense industrial base contracts. FedRAMP and DoD Impact Levels gate cloud services. NATO interoperability requirements gate allied system integration. Vendors that cannot demonstrate continuous, interaction-scoped governance face exclusion from defense and allied markets, independent of commercial or regulatory pressure. The 2025 NATO Summit formalized cybersecurity as a defined component of defense investment, allocating 1.5% of the new 5% GDP spending target to non-military capabilities including cyber resilience (Federal News Network, December 2025). This is not aspirational guidance; it is funded mandate with procurement consequences.

1.3 The Efficiency Paradox: Why AI Accelerates the Crisis

AI does not merely add requirements; it accelerates interaction density. Each successful automation tends to create:
- More interactions (because friction drops)
- More boundary crossings per interaction (because agents act across services)
- More state that must remain coherent (because multi-step workflows require memory)

As a result, “better models” and “more compute” do not automatically close the gap. They often make the missing governance layer visible sooner. A sensitivity analysis illustrates the acceleration. If participants, modalities, features, authorities, and transports each increase by 20% over a planning horizon:
- Fragmented coordination scales by $(1.2)^5 \approx 2.49\times$
- Interaction-scoped coordination scales by approximately $1.2\times$

1.4 The Coordination Tax: Latency, Energy, and Lost Potential

The converged pressures and the efficiency paradox manifest as one failure mode: the coordination tax imposed by fragmented control planes. Every system boundary crossing tends to require:
- Redundant authentication (identity does not travel with the interaction)
- Redundant policy evaluation (policy state does not persist across boundaries)
- Context reconstruction (state must be reassembled at each hop)
- Asynchronous reconciliation (governance is inferred after the fact rather than enforced by construction)

The tax is paid in three currencies: Latency: Boundary coordination consumes the latency budget of serious applications (safety-critical control, accessibility synchronization, immersive spatial interaction). When tolerances are tight, “eventual coherence” is disqualifying. Energy and compute: The same interaction is evaluated repeatedly across layers and services. At IoE scale, redundant checks,

duplicated pipelines, and reconciliation workloads become a meaningful fraction of infrastructure energy. Lost potential: Capabilities exist (transport, compute, models, sensors), but they do not compose into governed, coherent experiences because the interaction is not treated as the unit of control. The ceiling is coordination, not bandwidth. Section 9 quantifies the infrastructure implications: sustaining fragmented coordination at 2040 scale requires an estimated $2-3 trillion in dedicated coordination infrastructure. Interaction-scoped architectures reduce this to under $10 billion: a 200:1 difference that determines whether IoE is buildable.

1.5 What SIOAR-Aligned Solutions Must Provide (and What They Cannot Be)

SIOAR is not a feature bundle. It is the minimum set of constraints a future-safe IoE stack must satisfy: Continuity: One persistent interaction identity must span orchestration changes, participant/device churn, and policy updates. Transport agnosticism: Identity + policy state remain stable across heterogeneous protocols and network bearers. Computational transparency: All computation that materially affects the interaction (AI inference, transformations, analytics, decision support) executes within the same governed interaction boundary rather than in parallel shadow systems. What SIOAR-aligned approaches cannot be:
- Not a single protocol specification (implementations must work across SIP, MQTT, HTTP, industrial protocols, and future standards)
- Not a vendor platform (this is a constraint, not a product category) - Not “one more orchestration tool” (the move is subtractive, collapsing fragments, rather than additive coordination)

1.6 Business Outcome Summary

When interaction-scoped control replaces fragmented coordination, several effects follow directly from the architecture rather than from operational tuning: - Incident blast radius becomes bounded by interaction scope, because configuration and policy changes cannot propagate beyond defined session boundaries - Compliance enforcement shifts from retrospective to real-time, since policy decisions execute within the interaction rather than being inferred from logs - AI compute overhead declines, as duplicated pipelines and cross-session reconciliation are eliminated by construction • Energy consumption attributable to coordination decreases, because redundant validation, synchronization, and retry loops are no longer required • Interoperability complexity collapses, as policy semantics bind to a shared interaction identity instead of requiring bilateral mappings between systems

1.7 The Post-Enforcement Reality: Transitional Provisions as Strategic Trap

European Accessibility Act requirements applied to new products and services on June 28, 2025. The directive’s transitional provisions (grandfathering certain pre-June installed bases until 2030, and some terminals until 2045) function less as a reprieve than as a strategic trap. Six months post-enforcement, non-compliant vendors are discovering four constraints that “grandfathering” does not solve: The Growth Ban: Grandfathering protects a specific installed base, not a product line. You can operate existing non-compliant units, but you cannot expand capacity with new non-compliant placements. The Replacement Crisis: When a grandfathered unit fails or reaches end-of-life, replacing it with the “same” non-compliant hardware often triggers new-placement compliance rules. The installed base becomes both aging and irreplaceable. The Procurement Lockout: Grandfathering protects past revenue, not future access. Public-sector RFPs issued after June 2025 can require compliance; a vendor may

legally operate legacy fleets yet still be screened out of new bids. The Obsolescence Gap: A 2025-era terminal is unlikely to remain commercially viable until 2045. Long before legal protection expires, the technology becomes economically obsolete. Going into 2026, the strategic result is binary: non-compliant vendors can maintain legacy fleets or compete for new revenue, but they cannot do both. A practical example is the emerging requirement for real-time, AI-assisted accessibility (captions, translation, signing avatars) that must persist through network transitions and orchestration changes while remaining selectively delivered and policy-governed. These are hard real-time coordination problems. If the interaction identity and governance boundary fragment across services, the “accessible experience” becomes non-deterministic under load, exactly what procurement evaluation penalizes.

1.8 Spatial Computing: The Discontinuity Ahead

Spatial computing compounds the problem because it is not “more of the same.” It introduces new synchronized streams and tighter latency loops. Identity, consent, safety constraints, and residency policy must apply continuously to gaze, depth, scene context, and environment mapping, often at perceptual latencies. If accessibility is the minimum viable test for interaction-scoped governance, spatial computing is the next discontinuity that multiplies the cost of getting this wrong.

1.9 Regulatory Retreat as Evidence of Architectural Insufficiency

Across 2025, multiple jurisdictions moved toward deferrals, simplifications, or consolidation of overlapping compliance regimes. The pattern is not that requirements disappeared; it is that implementation has outpaced architectural capacity. The signal for IoE is the same as for RTC: when policy frameworks are hard to implement at scale, institutions tend to buy time. That time is an implementation window for architectures that can enforce governance during interactions rather than reconstruct it after. The Austria Meta Decision: Courts Outpace Architecture On December 18, 2025, Austria’s Supreme Court released its judgment in the 11-year Schrems/noyb case against Meta, ruling that GDPR requires continuous consent verification for personalized advertising, not point-in-time consent captured at account creation. The technical implication is immediate: advertising systems must evaluate consent state in real time as user context, preferences, and jurisdictional applicability change. Current adtech architecture precludes compliance, by Meta or anyone else. No deployed system can verify consent continuously because none maintains interaction-scoped governance. There is no alternative architecture that preserves both real-time ad-serving performance and continuous consent verification. Consent state must bind to the interaction identity and evaluate continuously, exactly the pattern SIOAR describes. Fragmented systems that capture consent once and assume it persists cannot comply. The regulatory picture is therefore not simple retreat. It is a pincer movement: legislators buying time while courts advance enforcement. Organizations relying on deferrals may find themselves caught between delayed compliance deadlines and binding judicial mandates that their architectures cannot satisfy.

1.10 The Window

IoE is at the phase where the evidence threshold is met (incidents, scaling math, and procurement gates), but broad institutional recognition is still forming. The window is open: before fragmented architectures harden into another decade of sunk cost, and before procurement and safety regimes make non-compliant architectures unfinanceable. The opportunity is to establish interaction-scoped

governance as a first-class control layer, before coordination failure becomes the default operating condition.

2. Problem Statement: Fragmentation's Trade-Off Reversal in IoE

2.1 Why Fragmentation Used to Work

The original RTC paper articulates the historical logic of fragmentation with precision: layered architectures offered modularity, independent evolution, and operational isolation. Transport could improve without policy changes. Identity systems could upgrade without affecting routing. Computational features could deploy independently without orchestration modifications. This separation worked because coordination costs were manageable relative to the benefits of independent operation. Each layer's evolution proceeded at its own pace, interfaces remained stable enough for loose coupling, and the operational complexity of maintaining consistency across fragments was lower than the complexity that would result from tight integration.

2.2 Why It Fails Now in IoE

The trade-off has reversed. As the original paper states: "Additive layering applies linear control to systems whose coordination complexity scales combinatorially with participants, modalities, and auxiliary computation." In IoE contexts, this combinatorial explosion becomes acute: Participants expand from human users to include devices, sensors, actuators, cloud services, AI agents, and federated organizational identities. Modalities expand from audio/video to include sensor telemetry, actuation commands, spatial data, environmental context, and multi-format content. Features expand from basic communication to include AI-generated content, real-time translation, automated compliance, predictive analytics, and autonomous decision support. Authorities multiply across organizational boundaries, regulatory jurisdictions, safety domains, and operational contexts. Transports diversify across cellular, satellite, enterprise networks, industrial protocols, and edge computing contexts. In fragmented architectures, coordination cost approximates: $C_{frag} = k \times P \times M \times F \times A \times T$ Why the product term appears in fragmented designs: when these concerns live in separate layers, coherence is enforced at combinations of dimensions (participants multiplied by modalities multiplied by features, and so on). Adding one more feature or modality does not simply add work. It multiplies the number of potential mismatch points and therefore the operator effort and failure modes. The constant k aggregates the per-surface cost (latency, CPU cycles, bandwidth, and operational overhead). The thesis does not depend on any particular value of k. The scaling difference between a product and a sum dominates once multiple independent pressures rise together. In SIOAR architecture, coordination cost approximates: $C_{SIOAR} = k \times (P + M + F + A + T)$ Because all dimensions bind to a single interaction identity, producing additive growth. Concrete Example (Healthcare AI Diagnostic Workflow):
- P = 10 participants (patient, 3 clinicians, 2 AI models, lab system, imaging system, EHR, audit system)
- M = 6 modalities (voice consultation, imaging data, lab results, EHR records, AI inference, audit logs)
- F = 8 features (real-time consultation, AI diagnosis assistance, automated lab ordering, clinical decision support, consent verification, compliance logging, multi-site coordination, accessibility accommodations)

- A = 4 policy authorities (patient consent, HIPAA compliance, institutional access control, AI model validation)
- T = 3 transport transitions (hospital WiFi → cellular → cloud → on-prem EHR) Fragmented coordination surfaces: 10 × 6 × 8 × 4 × 3 = 5,760 potential coordination points SIOAR coordination surfaces: 10 + 6 + 8 + 4 + 3 = 31 coordination points The ratio: 186:1 reduction in coordination complexity When each of these dimensions scales independently while governance remains externalized across fragments, “systems begin failing not because any single component breaks, but because coordination across fragments becomes intractable.” The operational burden of maintaining consistency (policy synchronization, identity mapping, state reconciliation, audit correlation) grows faster than the benefits of independent layer evolution.

2.3 Evidence Pattern: "The Pipes Improved; The Valve System Did Not"

The original RTC paper uses a hydraulic metaphor to capture this inversion: “We have the pipes but lack the valve system.” IoE exhibits this pattern with particular clarity because transport capabilities have improved dramatically while control-plane sophistication has not kept pace. More bandwidth and lower latency increase the speed and blast radius of mis-coordination. When configuration changes propagate in milliseconds across globally distributed systems, the time available for human intervention or graceful degradation shrinks to zero. When edge computing enables sub-10ms response times, the coordination overhead of consulting multiple fragmented policy authorities becomes the limiting factor. The incidents examined in Section 6 consistently demonstrate this pattern: the failure is not in the data plane (packets can flow, compute can execute) but in the control plane (configuration artifacts become inconsistent, policy state fragments, observability systems lose coherence under circular dependencies). IoE inherits and amplifies this challenge because it extends control chains into the physical world. In purely digital systems, fragmentation produces service disruptions and compliance violations. In IoE, fragmentation produces unsafe actuation, civic harm, and existential regulatory risk.

2.4 AI Agent Deployment: The Canary in the Control Plane

The same pattern appears with particular clarity in AI agent deployment. Enterprise AI platforms are discovering that models succeed well enough to expose the missing primitive. Contemporary AI agent platforms struggle with complex multi-step tasks in enterprise environments. The reasoning engine executes actions asynchronously and in sequence; if one action in the chain times out, the entire conversation context is lost. The agent does not gracefully degrade. It fails, often with error messages that tell the end user nothing helpful. This is increasingly a governance problem, not a model capability problem. Organizations use 897 applications on average, with only 29% connected, so agents inherit disconnected identity and policy domains. The result is siloed assistance: the CRM agent cannot reliably carry context into the data warehouse, and insights remain trapped inside app-specific boundaries, forcing humans to re-stitch workflow state. Industry practitioners describe discovering that “AI doesn’t speak unstructured data without an activation layer that spans fragmented systems.” The activation layer they describe is session-scoped governance. Without it, AI agents struggle to maintain coherent multi-step workflows because session state does not traverse system boundaries. The efficiency paradox manifests here with particular force. AI deployment was supposed to reduce operational overhead. Instead, successful AI deployment reveals that human operators were absorbing coordination costs that the architecture lacks mechanisms to handle. When the human is removed from the loop, the missing governance layer becomes visible.

2.5 The "Standard Primitives" Fallacy

A common counter-argument posits that existing protocols (QUIC/HTTP/3, service meshes like Istio/Envoy, modern token architecture like OIDC/JWT, workload identity via SPIFFE/SPIRE, and event-driven patterns like Kafka and saga orchestration) already provide the necessary continuity and governance tools. This view conflates transport persistence with interaction governance, request-level policy with interaction-level policy, and event reconstruction with real-time coherence. 1. The QUIC/HTTP/3 Trap: Connection Migration ≠ Session Governance QUIC's Connection ID allows a transport link to survive IP address changes (WiFi to 5G handover). This solves transport continuity. It does not solve policy continuity. The Failure Mode: When a device migrates from a corporate network to public cellular, the transport stays alive but the governance context has changed. Zero Trust mandates policy re-evaluation. In fragmented architecture, this triggers re-authentication, creating latency spikes that break the <200ms accessibility budget. The Verdict: QUIC provides a durable pipe, but a dumb pipe. It lacks semantic awareness to enforce data residency or safety policies dynamically. The connection survives; the governance does not. 2. The Service Mesh Fallacy: Sidecars Don't Scale to the Edge Service meshes excel at East-West traffic governance within a cluster. They rely on attaching a proxy sidecar to every workload. The Failure Mode: You cannot run an Envoy sidecar on a 128KB RAM IoT microcontroller, nor on a battery-sensitive AR headset where every milliwatt matters. Default sidecar requirements (100m CPU, 128Mi memory) exceed the total capacity of most edge devices. Furthermore, service meshes enforce policy at request level, not interaction level. They lack temporal state to know that this request is step 4 of a 10-step workflow requiring consistent audit binding. The Verdict: Service meshes are too heavy for the edge and too stateless for interaction governance. 3. The JWT Gap: Snapshots vs. Continuous State JWTs are point-in-time snapshots of authority. The Failure Mode: A JWT issued with 1-hour expiration remains valid regardless of context changes. Revocation requires stateful blacklists, defeating the purpose of stateless tokens. Revocation in distributed systems requires "non-standardized backend interaction" between authorization servers and resource servers: the very coordination overhead stateless tokens were designed to eliminate. In IoE contexts where safety or sovereignty violations happen in milliseconds, a pre-signed pass is a liability. The Verdict: Compliance in 2026 requires continuous evaluation, not snapshot authorization. Tokens are artifacts of the interaction, not the governance layer. 4. The SPIFFE/SPIRE Gap: Workload Identity ≠ Interaction Identity SPIFFE provides cryptographic workload identity. SPIRE attests and issues those identities. This solves "which service is calling which service." It does not solve "which interaction context governs this call." The Failure Mode: Two workloads with valid SPIFFE identities can communicate. But if call A is part of a healthcare diagnostic interaction requiring HIPAA audit binding, and call B is part of a marketing analytics workflow with different retention rules, SPIFFE cannot distinguish them. The workloads are authenticated; the interaction context is invisible. The Verdict: Workload identity is necessary but not sufficient. Interaction identity must layer above it. 5. The Event Sourcing / Saga Fallacy: Reconstruction ≠ Real-Time Coherence Event sourcing (Kafka, etc.) and saga patterns handle distributed transactions by recording events and enabling reconstruction or compensation. The Failure Mode: These patterns assume you can replay or compensate after the fact. In real-time accessibility (signing avatars must sync within 200ms), safety-critical control (actuation decisions in milliseconds), and Zero Trust (policy must evaluate continuously), "eventually consistent" and "compensating transaction" are not acceptable. You cannot un-actuate a robotic arm. You cannot un-show an unauthorized data field. You cannot compensate for a vestibular mismatch that already caused nausea. The Verdict: Event sourcing solves durability and audit reconstruction. It does not solve real-time governance coherence. Sagas assume compensation is possible; IoE increasingly involves irreversible physical-world effects. The Architectural Void None of these standards provide a single, continuous reference identity that binds user, policy, and compute state together across transitions. They solve connectivity (QUIC), routing (Mesh), authorization snapshots (JWT), workload authentication (SPIFFE), or after-the-fact reconstruction (event sourcing). They do not solve unified interaction coherence. This gap cannot be filled by combining existing primitives. The coordination surface is missing from the architectural vocabulary itself. Engineers who grasp this will recognize that

the path forward requires constructing something not yet in standard toolkits: a session primitive that makes governance continuous rather than intermittent, structural rather than bolted-on.

3. The Missing Control Layer: SSOAR Constraints Generalized to IoE

3.1 The Three Constraints (Explicit Preservation)

The original SSOAR definition established three architectural constraints that must hold simultaneously. These constraints transfer directly to IoE by substituting “interaction” for “session”: 1. Continuity: One persistent interaction identity must exist across all orchestration changes, policy updates, participant additions and departures, transport switches, and computational invocations throughout the interaction’s lifetime. Without continuity, every orchestration change requires re-establishing identity, re-asserting policy, and re-synchronizing state across fragmented systems. This creates coordination overhead that scales with the number of changes, the number of participants, and the number of governance domains involved. 2. Transport agnosticism: The interaction identity and its associated policy state must remain independent of underlying network protocols, physical bearers, and transport mechanisms. Without transport agnosticism, every protocol transition or network handover triggers identity re-mapping and policy re-evaluation. IoE systems routinely transition between 5G, Wi-Fi, satellite, wired connections, and proprietary industrial protocols. If interaction identity is bound to transport specifics, these transitions become governance discontinuities. 3. Computational transparency: All computation affecting the interaction (inference, analytics, transcription, recording, transformation, optimization, security analysis) must execute within the governed interaction boundary. Without computational transparency, AI and automation features operate as parallel “shadow systems” that consume interaction data and influence interaction behavior but exist outside the interaction’s policy and audit scope. This produces the efficiency paradox: features added to improve performance or compliance themselves become sources of coordination overhead and governance gaps.

3.2 IoE Terminology: Why "Interaction" Rather Than "Session"

The telecommunications and RTC industries use “session” to describe bounded communication contexts. For IoE, “interaction” better captures the broader coordination requirement while preserving identical architectural constraints. Defining “Interaction” with Precision An interaction is a bounded coordination context where multiple participants (human, device, service, AI agent) exchange data, computation, or control under unified policy scope. Interactions have four definable boundaries that distinguish them from unbounded system operation: Temporal boundary: Explicit start and end events (session initiation, job completion, timeout expiration) Participant boundary: Defined set of entities with explicit join/leave semantics (device provisioning, user authentication, service registration) Policy boundary: Coherent governance domain where rules apply uniformly (consent scope, jurisdictional authority, safety envelope) Scope boundary: Clear delineation of what’s inside versus outside the interaction (which data flows, which computational resources, which audit trails) Domain-Specific Interaction Examples: RTC: Video conference from participant initiation to session termination, spanning all media streams, recording/transcription features, and policy controls (who can join, what can be recorded, where data resides). Telecom: Subscriber session from device network attachment to detachment, encompassing network slice resources, service function chains, mobility events, and QoS policies. Healthcare: Clinical encounter from patient check-in through documentation completion,

unifying patient consent state, clinician access, diagnostic systems, treatment recommendations, AI model invocations, and audit logging. Industrial: Production run from job initialization to completion, binding machine identity, operator credentials, control commands, safety interlocks, quality sensors, and predictive maintenance analytics. Smart City: Traffic signal optimization cycle for a defined corridor over a time window, coordinating vehicle detection, pedestrian signals, transit priority, emergency preemption, accessibility timing, and environmental constraints. Finance: Trading session from participant authentication through settlement, spanning order routing, execution, risk evaluation, compliance checks, and regulatory reporting. What distinguishes an interaction from general system operation is governance unity: all components participate under a single, continuous identity with coherent policy enforcement, not as independently coordinated fragments requiring post-hoc reconciliation. SIOAR as Architectural Pattern We use SIOAR (Single-Interaction Orchestration and Adaptive Routing) to describe the architectural pattern that satisfies the three constraints across these diverse interaction types. Whether implemented through proprietary systems, open standards, or hybrid approaches, SIOAR-aligned architectures collapse fragmented control contexts into unified interaction surfaces.

3.3 The Architectural Invariant

The core invariant for SIOAR can be stated as a “must” requirement: Every packet, policy decision, compute invocation, and audit record that affects an IoE interaction must bind to the same continuous interaction identity. This identity cannot be inferred post-hoc across fragments; it must exist as a first-class architectural primitive. This invariant has immediate architectural consequences: Policy decisions cannot be applied unless attached to a persistent interaction identity that survives orchestration updates. Routing changes must be represented as versioned state transitions of the interaction identity, not as emergent behavior from independent component decisions. Computational resources (AI models, analytics pipelines, transformation functions) must be invoked within the interaction’s policy boundary with explicit bindings to interaction identity, data permissions, and jurisdictional constraints. Audit and compliance evidence must be interaction-scoped by construction, not assembled retrospectively from system-wide telemetry through correlation heuristics.

3.4 Non-Goals (Protecting the Thesis from Implementation Debates)

Following the original paper’s boundary-setting discipline: “No specific mechanism is prescribed. No implementation is implied.” SIOAR is a constraint on architectural coherence, not a mandate for particular technologies. Implementations may vary widely across protocol stacks, cloud platforms, device ecosystems, and organizational contexts. The requirement is that implementations satisfy the three constraints and the architectural invariant, not that they employ specific mechanisms. This boundary is critical for cross-vertical synthesis. IoE encompasses telecommunications, healthcare, finance, civic infrastructure, and industrial automation, domains with vastly different regulatory frameworks, operational requirements, and technology stacks. Defense and national-security systems represent an extreme stress case for these same constraints, but are not treated as a separate architectural category here. SIOAR defines a shared architectural need without prescribing implementation uniformity.

4. "Signals from the Field" Methodology

4.1 Why Use Contemporary Voices

Rather than conducting a 20-year historical analysis of how fragmentation emerged, this paper employs a “signals from the field” approach: contemporary voices from distinct verticals independently describing fragments of the same systemic issue. This methodology serves a “Rosetta Stone” function - demonstrating that telecommunications engineers, cloud architects, healthcare informaticists, financial infrastructure operators, and smart city planners are encountering the same underlying pathology in their own vocabulary and domain contexts. When these independent signals are mapped onto the SIOAR framework, they reveal a cross-vertical architectural requirement that transcends any single industry or technology domain. Each vertical contributes: A stated pain point in domain-specific language from recognized practitioners and thought leaders. A recent incident or policy change demonstrating coordination failure or control-plane fragility. Mapping to the missing control layer requirement, showing how SIOAR constraints would address the root cause rather than symptoms. This approach has three virtues: Temporal precision: Citations from late 2024 and 2025 demonstrate that these are current, active pressures rather than historical artifacts. Independent validation: No single author or organization is positioned to diagnose IoE-wide architectural gaps. The convergence of independent voices provides validation that exceeds any individual perspective. Practical grounding: Each signal comes from operational reality (production incidents, regulatory pressures, deployment challenges) rather than theoretical speculation.

5. Vertical Chapters: Evidence and Mapping to SIOAR

Each vertical chapter follows a consistent template: 1. The vertical’s core “interaction” primitive 2. Which independent constraints dominate in this domain 3. Signals from recognized voices (with compliant short quotations) 4. Recent incidents or policy changes revealing coordination failures 5. Why additive fixes fail (tying back to Section 2’s combinatorial coordination argument) 6. What SIOAR would change at the architectural level 7. Minimal reference architecture for the vertical

5A. Telecom Networks and Telco Cloud

5A.1 Why Telecom Is the "Highest-Pressure" IoE Substrate

Why Telecom Is the “Highest-Pressure” IoE Substrate Telecommunications infrastructure must simultaneously satisfy ultra-scale concurrency (billions of devices), low latency (single-digit milliseconds for 5G URLLC), and regulated trust frameworks (lawful intercept, emergency services, consumer protection). It is the substrate upon which all other IoE verticals depend. Telecom’s interaction primitive is the subscriber session - a continuous relationship between device identity, subscriber identity, service entitlements, and network slice resources that must persist across mobility events, transport handovers, and service function chain updates.

5A.2 Signal: AI Becomes Core Operating Model

Signal: AI Becomes Core Operating Model Ivo Ivanov, CEO of DE-CIX (December 2025), frames the transformation underway: “Smart value creation demands equally smart technologies. AI not only offers support for more intelligent network management, but for operational excellence in all areas of the

telecommunications industry.” This signal matters for SIOAR because AI integration is not additive; it is transformational. When AI models drive network slicing decisions, quality of experience optimization, security threat detection, and customer service orchestration, they must operate within the same governance boundary as the subscriber interaction itself. If AI features exist as parallel systems that ingest telemetry and emit recommendations asynchronously, they create the coordination overhead and governance gaps that SIOAR is designed to collapse.

5A.3 Signal: Standards Gap in Zero Trust Deployment

Signal: Standards Gap in Zero Trust Deployment Xuan Si’s IETF Internet-Draft “Problem Statement of Zero Trust Deployment in Telecom Network Environments” (December 23, 2025, work in progress) explicitly identifies that current standards “fail to establish quantitative constraints on… validation frequency, session persistence mechanisms, and incremental state synchronization.” This signal is architecturally revealing. The draft identifies three SIOAR-adjacent requirements in different terminology: Validation frequency corresponds to the continuous policy evaluation loop rate - how often trust decisions must be reassessed during an interaction’s lifetime. Session persistence mechanisms directly name the continuity constraint - the need for interaction identity to survive network events. Incremental state synchronization addresses the control-plane propagation challenge - how policy and routing updates propagate without triggering full re-initialization. The fact that an IETF working group is explicitly articulating these gaps demonstrates that SIOAR’s problem statement is not theoretical. Telecom operators deploying Zero Trust architectures are encountering the missing control layer in production.

5A.4 Why Additive Fixes Fail in Telecom

Why Additive Fixes Fail in Telecom In telco cloud environments, each new security or policy layer adds coordination requirements across: Subscriber and device identity: Multi-factor authentication, SIM credentials, device fingerprinting, subscriber entitlements. Service chain functions: Network slices, virtual network functions, edge computing resources, content delivery. Regional and residency constraints: Data sovereignty rules, lawful intercept boundaries, regulatory compliance domains. AI inference pipelines: Traffic optimization, fraud detection, quality prediction, automated remediation. This recreates the original RTC paper’s “linear control versus exponential coordination” mismatch at telecommunications scale. When each dimension scales independently while governance remains externalized, operators face the operational burden described in Section 2.2: systems fail not because components break but because coordination becomes intractable. Adding more orchestration layers, more policy synchronization protocols, or more audit correlation tools does not solve this problem - it accelerates it. Each additional layer increases the number of pairwise mappings required for consistency.

5A.5 What SIOAR Looks Like in Telecom

What SIOAR Looks Like in Telecom A SIOAR-aligned telecommunications architecture would implement: Interaction identity: Subscriber, device, and service-chain identity bound to one continuous control surface that persists across mobility events, transport handovers, and service updates. Policy plane: Zero Trust continuous validation operating at the interaction identity scope, with dynamic trust scoring bound to the same identity primitive rather than requiring cross-system synchronization. Compute routing: AI-assisted functions for optimization, security analytics, and service orchestration invoked within the interaction identity boundary with explicit policy bindings rather than as parallel shadow systems. Deterministic compliance: Data residency and regulatory constraints enforced through interaction-scoped routing decisions rather than verified through post-hoc audit. When a subscriber interaction must remain within EU jurisdictional boundaries, that constraint routes computation and storage deterministically rather than being reconciled after violations occur. The architectural consequence is reduced coordination overhead: policy updates, trust decisions, and computational routing all operate against the same continuous interaction identity rather than requiring synchronization across fragmented identity, policy, and orchestration systems.

5B. Real-Time Communication and WebRTC

5B.1 RTC as the "Interaction Laboratory"

RTC as the “Interaction Laboratory” Real-time communication is where “interaction identity” is most intuitively understood because humans notice coherence failures immediately. When a video call experiences a 2-second delay, participants perceive the breakdown. When policy changes (recording starts, transcription activates, a participant is admitted) do not propagate coherently, users experience the system as “broken” even if no packets are lost. RTC’s interaction primitive is the media session - a continuous orchestration of participants, media streams, policy controls, and auxiliary features (transcription, recording, AI summarization, accessibility overlays) that must maintain coherence throughout the session’s lifetime.

5B.2 Signal: Observability as Control-Layer Prerequisite

Signal: Observability as Control-Layer Prerequisite Tsahi Levent-Levi’s work on RTC observability (BlogGeek, November 7, 2025) emphasizes the critical importance of early observability implementation in real-time communication stacks, advocating for observability to be a fundamental architectural consideration rather than an afterthought. This signal reflects a maturation in RTC operational practice: observability is not an optional enhancement but a fundamental control-layer requirement. Without visibility into session state, policy propagation, and computational invocation, operators cannot diagnose coordination failures or implement graceful degradation. For SIOAR, observability becomes interaction-scoped by design. Telemetry, metrics, and audit trails bind to the same continuous interaction identity, enabling coherent diagnosis rather than requiring correlation heuristics across fragmented logging systems.

5B.3 Recent Microsoft Signal

Recent Microsoft Signal: Operational Discipline Acknowledges Control-Plane Reality Microsoft’s official service health documentation establishes that for major incidents with “broad and noticeable impact across a large number of organizations, a preliminary post-incident review (PIR) will be delivered via your Service health within 48 hours of incident resolution, followed by a final PIR within five business days.” This operational commitment reflects an understanding that control-plane failures require rapid diagnosis and communication. The 48-hour preliminary PIR timeline acknowledges that modern RTC systems are mission-critical infrastructure where coordination failures directly impact business continuity.

5B.4 Incident Evidence: Collaboration Platforms Degrade

Incident Evidence: Collaboration Platforms Degrade as Control-Plane Coordination Fails Microsoft Teams messaging delay incident (December 19, 2025) illustrates how “real-time” user experience collapses when coordination layers degrade. Users perceive Teams as a messaging platform, but the underlying architecture involves identity services, presence systems, message routing, policy enforcement, and audit logging - all of which must coordinate for the user experience to feel “instant.” When coordination latency increases, the illusion of real-time breaks. These incidents demonstrate the original RTC paper’s thesis at production scale: fragmentation creates brittleness under load and cascading failures across dependent features.

5B.5 SIOAR Mapping for RTC

SIOAR Mapping for RTC “Session continuity” transitions from operational best practice to architectural requirement once RTC sessions incorporate: Accessibility overlays: Live captions, translation, sign

language interpretation, audio descriptions, and AR accessibility aids must maintain synchronization with the session state. Safety and compliance controls: Content moderation, harassment detection, regulatory compliance all require mid-session enforcement. AI features: Summarization, translation, sentiment analysis, meeting insights must operate within the session's policy boundary. Transcription and recording: Policies governing storage, retention, and access must bind to the session identity. An SIOAR control layer for RTC provides: Single session identity as the routing primitive for participants, media, and features. Policy propagation that occurs internally to the session identity rather than requiring cross-service synchronization. Computational transparency: Features do not instantiate parallel "shadow sessions" that consume media streams outside the governed session boundary. The operational consequence is reduced incident frequency and faster recovery: when policy and features bind to a continuous interaction identity, coordination failures become locally contained rather than cascading across system-wide dependency graphs.

5C. Multi-Cloud, Edge Cloud, and Internet Infrastructure

5C.1 The Key Pattern: Circular Dependencies and "Blindness"

The Key Pattern: Circular Dependencies and "Blindness" Adrian Cockcroft, in a LinkedIn post from October 2025 discussing cloud infrastructure failures, identifies a critical pathology: "This is a circular dependency… the ability to observe and fix the failure… failed… engineers… were completely blind." This signal captures a fundamental control-plane vulnerability: when observability infrastructure depends on the same systems it monitors, failures can create operational blindness. Engineers cannot diagnose or remediate problems because the diagnostic tools themselves are unavailable. For SIOAR, this highlights the importance of interaction-scoped control planes that maintain independence from data-plane infrastructure. Control and observability must be architectural primitives, not layered services that share fate with the systems they govern.

5C.2 Incident Evidence: Configuration Propagation

Incident Evidence: Configuration Propagation and Deployment Safety System Failures Azure Front Door disruption (October 29, 2025): ThousandEyes analysis reports that "A configuration change caused nodes across the global fleet to fail to load properly." This incident exemplifies control-plane fragility at cloud scale. A single configuration artifact, when propagated across a globally distributed system, triggered fleet-wide failures. The failure mode was not resource exhaustion or transport congestion but control-plane state inconsistency. The SIOAR implication: configuration artifacts must be versioned, scoped to interaction domains, and propagated with deterministic rollback boundaries. Fleet-wide propagation of unversioned configuration creates unbounded blast radius.

5C.3 Why This Matters for IoE

Why This Matters for IoE IoE control planes inevitably traverse multiple organizational and technical boundaries: Device edge → Carrier network → Cloud edge → Regional cloud → SaaS platforms → AI endpoints If cloud infrastructure (operating within a single organizational domain and technology stack) can become operationally blind under circular dependencies, IoE will amplify that risk because control chains extend across organizational boundaries, regulatory jurisdictions, and technology heterogeneity. The physical-world coupling makes this risk existential. In pure cloud contexts, configuration failures produce service disruptions. In IoE contexts, configuration failures can produce unsafe actuation (industrial control systems), regulatory violations (healthcare data mishandling), or civic harm (smart city infrastructure failures).

5C.4 SIOAR Mapping for Cloud and Edge Infrastructure

SIOAR Mapping for Cloud and Edge Infrastructure The missing control layer manifests in cloud infrastructure as: Unbounded blast radius for configuration propagation: Changes affect global fleets rather than scoped interaction domains. Policy and identity state not coherently versioned across dependent systems: Each service maintains its own view of identity and policy truth, requiring continuous synchronization. Observability systems that share fate with monitored infrastructure: Circular dependencies create operational blindness. SIOAR requirements for cloud infrastructure: Versioned interaction identity with versioned policy state: Changes are atomic updates to interaction scope rather than eventually consistent synchronization across fragments. Safe rollout boundaries: Configuration and policy changes propagate within interaction-scoped boundaries with deterministic rollback rather than fleet-wide propagation with heuristic rollback. Independent control plane: Observability and governance infrastructure maintain operational independence from data-plane systems they monitor and control.

5D. Enterprise SaaS and Salesforce Ecosystem

5D.1 Incident Signal: Immutable Infrastructure

Incident Signal: Immutable Infrastructure as Explicit Control-Plane Lesson Heroku's (Salesforce) June 15, 2025 outage summary states: "The root cause… was an unexpected change to our running environment." This language ("unexpected change to our running environment") reflects the maturation of immutable infrastructure thinking. The operational expectation is that running environments should not change except through controlled deployment processes. When that invariant is violated, the system becomes unpredictable. For SIOAR, this highlights the importance of versioned, interaction-scoped control planes. "Running environment" is not a monolithic global state but a composition of interaction-scoped configurations that must evolve deterministically.

5D.2 Strategic Signal: Enterprise AI Agents

Strategic Signal: Enterprise AI Agents Require Session Continuity at Scale Enterprise software platforms pursuing AI agent deployment face fundamental architectural challenges related to session continuity. AI agents operating at enterprise scale must maintain persistent context across: Multiple application boundaries: Communication platforms, CRM systems, external API integrations, email threads, calendar systems, document repositories Extended time horizons: Complex enterprise workflows span hours, days, or weeks, requiring continuous state preservation Multi-party coordination: Human users, other AI agents, external services, approval workflows must coordinate coherently Dynamic policy enforcement: Data access rules, approval authorities, compliance constraints that change mid-workflow must propagate consistently Why fragmentation creates an insurmountable barrier: Each application maintains independent identity and policy contexts. An AI agent operating across multiple systems must coordinate identity, permissions, and state across fragmented infrastructure. The coordination overhead scales combinatorially with workflow complexity, creating three failure modes: 1. Degraded context operation: The agent forgets previous conversation context, loses workflow state, requires constant user re-authentication - producing poor user experience that undermines the value proposition 2. Parallel shadow state: The agent creates its own coordination layer that bypasses existing governance systems to maintain continuity - creating existential compliance and security risk 3. Fragmentation into app-specific assistants: Separate agents for each application, requiring users to manually coordinate across fragments - defeating the business case for unified AI assistance Of the three coping strategies, two destroy product value and one creates regulatory liability. The root issue is architectural: enterprise AI agents require continuous interaction identity across applications with coherent policy enforcement for the workflow's full lifetime. The 2025 market evidence supports this analysis: organizations announcing aggressive AI agent roadmaps subsequently temper deployment timelines, citing "quality assurance" and "responsible AI practices" -

language that often masks the underlying architectural reality that current fragmented infrastructures cannot support the session continuity these features require.

5D.3 Policy and Control-Plane Signal

Policy and Control-Plane Signal: Data Access Governance Tightens Under AI Pressure Salesforce's June 2025 update to API terms restricting third parties from indexing, copying, or storing Slack messages for long-term use (reported by Reuters) reflects growing AI governance concerns: when third-party AI tools consume messaging data outside the platform's governance boundary, organizations lose control over data lineage, retention compliance, and downstream usage. The platform response is boundary enforcement at APIs. But this is insufficient for IoE contexts because interactions routinely span multiple platforms. SIOAR's requirement is that policy binds to interaction identity across platforms rather than being enforced independently at each platform boundary.

5D.4 Why This Matters for IoE

Why This Matters for IoE IoE encompasses not only devices and networks but enterprise workflow fabric: Identity and policy authority: Employee credentials, device entitlements, application permissions. Audit and compliance systems: Activity logs, data lineage, regulatory evidence. AI and automation integration: Workflow bots, intelligent assistants, automated decision support. As AI features attach to enterprise data exhaust, vendors enforce policy at platform boundaries - often without a unified interaction identity shared across tools. An employee's workflow might span Slack messages, Salesforce records, Google Drive documents, and external API calls. If each system maintains independent identity and policy contexts, coordination overhead grows combinatorially with integration complexity.

5D.5 SIOAR Mapping for Enterprise SaaS

SIOAR Mapping for Enterprise SaaS IoE requires a control layer where: Data access, AI inference, retention rules, and user/device identity bind to a single interaction boundary rather than being independently enforced by each platform. Policy updates propagate coherently across all systems participating in an interaction rather than fragmenting into inconsistent app-by-app enforcement. Auditability is real-time and interaction-scoped: compliance evidence is produced by construction during the interaction rather than assembled retrospectively through log correlation. The business consequence is reduced compliance risk and operational complexity. When policy binds to interactions rather than platforms, organizations can reason about governance at the workflow level rather than managing pairwise integrations between independent systems.

5E. Financial Services: 24/7 Markets and Real-Time Risk

5E.1 Signal: The "Always-On" Interaction Model

Signal: The "Always-On" Interaction Model Becomes Normal Billy Hult, CEO of Tradeweb, stated via Reuters (December 18, 2025): "I have no doubt… round-the-clock trading will become the norm across markets." This signal reflects a fundamental shift in financial infrastructure: the assumption of operational "downtime" is disappearing. When markets operate 24/7, all supporting infrastructure (identity verification, risk controls, compliance surveillance, settlement processes) must operate continuously without scheduled maintenance windows.

5E.2 IoE Analogy

IoE Analogy Financial services infrastructure is a real-time distributed system where: Identity comprises participant credentials, account entitlements, and venue-specific permissions that must remain

coherent across trading sessions. Policy encompasses risk limits, compliance rules, market surveillance, anti-money laundering checks, and algorithmic trading controls. Compute includes routing optimization, pricing engines, fraud detection, and regulatory reporting systems. The “24/7” operational model collapses the old assumption that governance can occur “offline” through batch reconciliation. Risk decisions, compliance checks, and surveillance analytics must occur within the same real-time interaction boundary as trade execution.

5E.3 SIOAR Mapping for Financial Services

SIOAR Mapping for Financial Services Interaction identity in financial services must persist through: Venue hopping: A single order might execute across multiple trading venues, liquidity pools, and settlement systems. Liquidity routing: Algorithmic trading strategies dynamically route orders based on price, volume, and latency - all subject to risk controls and compliance rules. Dynamic risk rules: Credit limits, position limits, and concentration constraints update in real time based on market conditions and portfolio exposure. Computational transparency becomes critical: AI-driven trading algorithms, risk analytics, and surveillance systems must be invoked inside the same governed transaction identity. When these systems operate as parallel infrastructures that consume market data and emit recommendations asynchronously, they create governance gaps: trades execute under one risk profile while risk calculations operate under a different, potentially stale profile. SIOAR enables: Real-time risk enforcement: Risk limits and compliance rules bind to the same interaction identity as trade execution, ensuring policies are evaluated against current positions rather than reconciled after executions. Audit coherence: Surveillance and regulatory reporting become interaction-scoped by construction, producing compliant evidence at execution time rather than reconstructing transaction context from fragmented logs. Graceful degradation: When connectivity to specific venues or liquidity pools degrades, interaction-scoped control allows partial operation rather than global failure.

5F. Healthcare and Digital Health

5F.1 Signal: Orchestration Becomes Explicitly the Battlefield

Signal: Orchestration Becomes Explicitly the Battlefield At HLTH 2025 (November 2025), InterSystems and Google Cloud announced a new integration between InterSystems HealthShare and Google Cloud’s Healthcare API. The collaboration explicitly aims to create “a real-time, harmonised data foundation designed specifically to support generative and agentic AI across healthcare enterprises.” Signify Research’s analysis noted: “data infrastructure is now the strategic bottleneck, and enabler, for AI in healthcare.” This signal directly names the SIOAR thesis: orchestration and trust are not independent concerns but coupled requirements. Without unified orchestration (interaction identity spanning clinical workflows, AI inference, data access, and patient consent) trust becomes unenforceable.

5F.2 Signal: Institutional Recognition

Signal: Institutional Recognition of Orchestration as Strategic Imperative In October 2025, Mayo Clinic launched Platform_Orchestrate, described by Maneesh Goyal (COO, Mayo Clinic Platform) as a program “created to help our life sciences and biopharma partners bring therapies to patients safely and faster.” When a leading clinical institution creates a named platform-level orchestration program, it signals that fragmented systems have become a strategic barrier to safety and speed, not merely an IT inconvenience. The initiative addresses clinical data, research, and deployment workflows; the architectural recognition that unified orchestration is required for governance and safety applies equally to real-time clinical interactions.

5F.3 Why Healthcare Is "Hard Mode" for IoE

Why Healthcare Is “Hard Mode” for IoE Healthcare constraints collide with maximal intensity: Trust: Clinical decisions carry existential risk - diagnostic errors, treatment complications, medication interactions. Privacy: Patient data is protected by HIPAA, GDPR, and domain-specific regulations that require deterministic enforcement. Safety: Medical devices, medication administration, and procedural workflows require humans-in-the-loop for critical decisions. Residency and sovereignty: Cross-border data flows require jurisdictional compliance, while emergency care may require immediate access across boundaries. Interoperability: Clinical workflows span electronic health records, medical imaging systems, laboratory information systems, pharmacy systems, and medical devices - each with independent identity and policy contexts. AI governance: Machine learning models for diagnosis, treatment recommendation, and operational optimization must be traceable, auditable, and subject to clinical validation. Fragmentation in healthcare is not merely inconvenient - it is patient-impacting. When coordination failures occur, the consequences range from degraded care quality to adverse events and regulatory penalties.

5F.4 SIOAR Mapping for Healthcare

SIOAR Mapping for Healthcare Interaction identity in healthcare must unify: Patient consent state: What data can be accessed, by whom, for what purposes, under what conditions. Clinician identity and credentials: Licensing, specialization, institutional privileges, on-call status. Device identity and provenance: Medical device authentication, software version verification, calibration status. Model identity and validation: Which AI models can operate on which patient data, in which clinical contexts, under which regulatory approvals. Data lineage: Complete provenance of diagnostic inputs, treatment recommendations, and clinical documentation. Policy enforcement must be mid-interaction: When a radiologist reviews an imaging study assisted by an AI detection model, consent verification, data access authorization, model validation, and audit logging must occur within the same governed interaction boundary. Post-hoc reconciliation is insufficient for clinical safety and regulatory compliance. Compute routing must be governed: Which AI models run on which patient data, in which geographic region, under which consent profile - these decisions are not operational optimizations but clinical governance requirements. SIOAR enables: Real-time consent enforcement: Patient consent state binds to clinical interaction identity, enabling mid-interaction policy decisions rather than batch access reviews. Deterministic jurisdictional compliance: Data residency requirements route clinical workflows and AI inference deterministically rather than being verified through audit. Clinical safety with AI: Machine learning features operate within governed interaction boundaries where model validation, data lineage, and clinical oversight are structural rather than procedural.

5G. Smart Cities and Civic IoT

5G.1 Signal: Human Intent Versus Technology Optimization

Signal: Human Intent Versus Technology Optimization Kate O’Neill, in Smart City Expo session materials (November 5, 2025), poses the fundamental governance question: “Are we optimizing for technology, or for human flourishing?” This signal reframes smart city discourse from technical optimization to civic governance. The relevant question is not “what can the technology do” but “what should the technology be permitted to do, and how do we enforce those boundaries.”

5G.2 Why Smart Cities Amplify the Missing Layer

Why Smart Cities Amplify the Missing Layer Cities are systems of systems: Mobility infrastructure: Traffic signals, public transit, parking, ride-sharing coordination, electric vehicle charging. Energy systems: Smart grids, distributed generation, demand response, renewable integration. Public safety: Emergency response, surveillance systems, environmental sensing, disaster coordination. Environmental monitoring: Air quality, water quality, noise levels, climate adaptation. Citizen services:

Permitting, payments, communications, civic engagement platforms. Fragmentation in smart cities yields: Policy inconsistency: Privacy rules, surveillance constraints, and data retention requirements vary across departments and vendors. Cross-agency identity incoherence: A resident interacting with multiple city services encounters fragmented identity contexts - parking permits, utility accounts, emergency contacts, civic engagement profiles are independently managed. Brittle vendor integrations: Procurement processes produce heterogeneous systems that integrate through bilateral APIs rather than shared interaction identity.

5G.3 SIOAR Mapping for Smart Cities

SIOAR Mapping for Smart Cities Define civic interaction identity not as a person-tracking primitive but as a governance boundary that binds policy to system interactions. Examples of civic interaction identities: Traffic signal optimization interaction: Coordinates vehicle detection, pedestrian signals, public transit priority, emergency vehicle preemption - all subject to accessibility requirements, environmental constraints, and safety policies. Emergency response interaction: Unifies incident detection, resource dispatch, interdepartmental coordination, public alerting - all subject to privacy protections, jurisdictional authorities, and operational protocols. Environmental sensing interaction: Correlates air quality monitoring, noise measurement, water quality assessment - subject to data retention policies, public transparency requirements, and scientific validation. The SIOAR requirement is that privacy rules, access controls, retention policies, and audit requirements embed into the interaction boundary itself rather than being enforced independently by each participating system. Computational transparency for civic AI: When machine learning models optimize traffic flow, predict maintenance needs, or allocate public services, they must operate within governed interaction boundaries where algorithmic accountability, bias monitoring, and public oversight are structural. SIOAR enables: Consistent civic policy enforcement: Privacy protections, surveillance constraints, and accessibility requirements bind to interaction identity rather than fragmenting across vendor platforms. Transparent algorithmic governance: AI optimization operates within governed interaction boundaries where decision provenance, bias assessment, and human oversight are architecturally guaranteed. Citizen control: Residents can reason about and exercise control over civic interactions (opt-out, data access, consent management) at the interaction level rather than navigating fragmented vendor systems.

5H. Industrial Automation and Industry 4.0

5H.1 Signal: From Isolated Products to Unified Platforms

Signal: From Isolated Products to Unified Platforms Chingpo Lin, VP for Global IoT Automation at Advantech (November 2025), observes: “Customers are no longer satisfied with isolated products. They expect partners who understand both IT and OT, can guide them through cybersecurity and regulatory complexities, and provide scalable edge platforms that deliver real-time intelligence across entire facilities.” This signal captures the architectural shift underway in industrial automation. The coordination problem is no longer a technical inconvenience; it is a market requirement. Vendors who cannot unify IT and OT under coherent governance are losing to those who can.

5H.2 Why OT Environments Expose Control-Plane Requirements

Why OT Environments Expose Control-Plane Requirements Operational Technology (OT) environments (manufacturing facilities, process control systems, industrial automation) impose requirements that expose fragmentation’s costs: Deterministic behavior: Industrial processes require predictable, repeatable control sequences. Non-deterministic coordination introduces safety risk. Safety constraints: Machinery and chemical processes operate within safety envelopes that must be enforced structurally, not procedurally. Long lifecycle compatibility: Industrial equipment operates for decades.

Technology integrations must preserve functionality across multi-decade timescales. IT/OT convergence: Digital transformation brings enterprise IT systems (analytics, optimization, supply chain coordination) into operational environments where they interact with control systems, sensors, and actuators. Fragmented governance between IT and OT produces: Identity gaps: Ambiguity about who or what is authorized to command actuators, update control parameters, or access operational data. Policy drift: Safety rules, operational procedures, and compliance requirements diverge across plants, shifts, and vendors. Compute sprawl: Edge inference for predictive maintenance, cloud analytics for optimization, and on-premises control systems operate as independent silos without unified governance.

5H.3 SIOAR Mapping for Industrial Automation

SIOAR Mapping for Industrial Automation Interaction identity in industrial contexts binds: Machine identity: Equipment authentication, firmware version, calibration status, maintenance history. Operator identity: Certifications, shift assignment, authorization scope, training records. Control command provenance: Complete lineage of control decisions including automated optimization, manual overrides, and safety interlocks. Safety envelope: Operating parameters, emergency shutdown conditions, environmental constraints. Transport agnosticism becomes critical: Industrial environments employ diverse protocols - Modbus, OPC UA, Profinet, EtherCAT, proprietary vendor protocols. SIOAR interaction identity must span this heterogeneity without requiring protocol-specific governance mappings. Computational transparency: Predictive maintenance models, production optimization algorithms, quality control analytics, and energy management systems must operate inside governed interaction boundaries where model validation, safety constraints, and operational authority are structurally linked. SIOAR enables: Deterministic safety: Safety rules and operational constraints bind to control interaction identity, ensuring commands are evaluated against current safety state rather than reconciled after execution. Long lifecycle governance: Interaction identity and policy semantics remain stable across technology evolution, enabling graceful integration of new capabilities without disrupting operational continuity. Cross-domain coordination: IT analytics and OT control systems coordinate through shared interaction identity rather than requiring bilateral integration for each new capability.

5I. Defense and Coalition Systems (Stress Case)

(Stress Case)

5I.1 Why Defense Validates Rather Than Expands the Thesis

Why Defense Validates Rather Than Expands the Thesis Defense and coalition systems provide a stress case for interaction-scoped governance. Unlike civilian platforms, these environments assume degraded networks, dynamic trust posture, and mid-interaction policy change as normal operating conditions rather than edge cases. If the architectural constraint holds here, it holds anywhere. Public DoD Zero Trust guidance explicitly articulates the requirement: “The ZT security model transitions away from trusted networks, actors, and devices to an environment with continuous authentication and fine-grained policy enforcement” (DoD CIO, Zero Trust for Operational Technology Activities and Outcomes, November 2025). This is the same architectural pattern described throughout this paper. Authentication and authorization must occur continuously, not solely at session initiation. Policy must travel with the interaction across heterogeneous transport and administrative domains. The fact that defense practitioners arrived at this requirement independently, under operational pressures that exceed civilian contexts, confirms that the constraint is fundamental rather than domain-specific.

5I.2 The Coalition Interoperability Signal

The Coalition Interoperability Signal NATO's adoption of the Zero Trust Data Format (ZTDF) and Allied Communication Publication 240 (ACP 240) reinforces the same architectural direction. These standards embed policy and classification metadata directly into data objects, ensuring that "protection and policy travel with the data wherever it goes" (Virtru, October 2025). Operation HIGHMAST, the UK's 2025 global carrier strike group deployment, validated ACP 240 under operational conditions. The coordination failures practitioners describe in these contexts mirror those observed in cloud and telecom incidents: not transport capacity, but fragmented control contexts that require post-hoc reconciliation of identity, policy, and state.

5I.3 Scope Limitation

Scope Limitation This paper does not propose defense-specific mechanisms. The relevance of defense systems here is purely confirmatory: when constraints are pushed to their limit, fragmented architectures fail first, and interaction-scoped control becomes unavoidable. The significance of these environments lies not in their uniqueness, but in how quickly they encounter the same coordination limits that civilian systems are now reaching at scale.

6. Cross-Vertical Incident Analysis: Control-Plane Fragility in Production

6.1 Why These Incidents Were Selected

The incidents examined in this section were not selected because they represent the worst outages of 2025, or because they affected the most users, or because they produced the largest financial losses. They were selected because they exhibit the same structural pattern: control-plane complexity and configuration propagation failures, not raw transport limitations, are the primary drivers of systemic outage and governance breakdown.

6.2 Incident Analysis

Cloudflare (November 18, 2025) On November 18, 2025, Cloudflare experienced a global outage lasting approximately three hours, affecting major platforms including X, OpenAI, Anthropic, Spotify, and others. The root cause was a database permission change that caused a Bot Management feature file to expand from 60 to over 200 entries. This larger-than-expected file propagated globally to all network machines. The Rust proxy service (FL2) hit a hard-coded memory limit and triggered a panic, crashing edge nodes in a loop. The failure was not transport capacity. The failure was a configuration artifact that exceeded expected bounds and propagated without session or rollout boundaries. A small control-plane change cascaded globally because no interaction-scoped boundary constrained its blast radius. Azure Front Door (October 29, 2025) On October 29, 2025, Azure Front Door experienced an outage lasting approximately eight hours, affecting Microsoft 365, Teams, Outlook, Xbox Live, and third-party services across all Azure regions globally. The root cause was an inadvertent tenant configuration change that bypassed protection mechanisms due to a software bug. Inconsistent configuration propagated across global AFD infrastructure, affecting all Points of Presence. Microsoft acknowledged that a "configuration change propagated through Azure Front Door" caused the disruption. The failure was not transport capacity. The failure was configuration propagation to global

PoPs without rollout boundaries or interaction-scoped containment. Salesforce/Heroku (June 10, 2025) On June 10, 2025, Salesforce's Heroku platform experienced an outage lasting up to 24 hours for many customers. The root cause was an "unintended system update across production infrastructure" where an automated OS update that should have been disabled executed anyway. Heroku's post-incident analysis identified "lack of sufficient immutability controls" that "allowed an automated process to make unplanned changes to our production environment." The failure was not transport capacity. The failure was an automated change that bypassed immutability controls and cascaded through infrastructure. The architectural assumption (that certain changes would not occur) was violated because governance was procedural rather than structural. Google Cloud Platform IAM (June 12, 2025) On June 12, 2025, Google Cloud Platform experienced a global IAM outage lasting approximately 14 hours. The root cause was an invalid automated update to the API management system (Service Control). A dormant code path deployed May 29, 2025 with quota policy checks contained a null pointer vulnerability. A June 12 policy change to regional Spanner tables activated the dormant code path. Critically, "given the global nature of quota management, this metadata was replicated globally within seconds." The null pointer dereference crashed Service Control nodes across regions in a loop, affecting 50+ GCP services across 40+ regions and cascading to Gmail, Drive, Calendar, Meet, Vertex AI, BigQuery, and third-party services including Spotify, Discord, and Cloudflare authentication systems. The failure was not transport capacity. The failure was policy metadata corruption that propagated globally within seconds because no session boundary constrained its blast radius. This is the control-plane fragility pattern in its purest form: a small configuration change, global propagation, cascading failure. Microsoft Teams (December 2025) In December 2025, Microsoft Teams experienced multiple incidents including a December 19 event causing message delays, delivery failures, and presence/notification inconsistencies. Analysis suggests routing misconfiguration that isolated healthy infrastructure. A December 18 incident affecting Japan and China was attributed to "routing misconfiguration that isolated healthy infrastructure." The failure was not transport capacity. The failure was control-plane routing issues affecting global messaging infrastructure without graceful degradation.

6.2.1 Additional 2025 Incidents: Pattern Breadth

The five incidents analyzed above were not isolated events. Cross-vertical outages exhibiting identical structural patterns occurred throughout 2025 across hyperscalers, critical internet intermediaries, collaboration platforms, and developer infrastructure. These incidents can be classified into two categories that support different aspects of the control-plane fragility thesis: Class 1: Configuration Propagation and Global State Failures These incidents directly support the claim that fragmented control planes lack interaction-scoped boundaries, allowing configuration changes to propagate globally with unbounded blast radius: AWS DynamoDB DNS (October 20, 2025): A latent race condition in DynamoDB's automated DNS management system produced an empty DNS record for the service's regional endpoint. The failure cascaded through US-EAST-1, affecting EC2 instance launches, Lambda invocations, Fargate tasks, and dozens of AWS services. Estimated insurance losses reached $581 million. The root cause was control-plane automation interacting with global state without session boundaries. Cloudflare 1.1.1.1 Resolver (July 14, 2025): A dormant configuration error (introduced June 6 during Data Localization Suite preparation) linked 1.1.1.1 resolver IP prefixes to a non-production service topology. When engineers made an unrelated change July 14, the system withdrew BGP routes globally, making the world's most widely used public DNS resolver unreachable for 62 minutes. The configuration sat dormant for 38 days without triggering alerts because the DLS service was not yet active. When activated, it propagated globally through legacy routing systems lacking staged rollout capabilities. Cloudflare Dashboard/API (September 12, 2025): Dashboard and API services experienced outage for up to one hour. While specific triggers differed from the November 18 incident, the failure exhibited the same control-plane fragility pattern: management-plane changes with broad customer visibility and global propagation. Cloudflare Partial Network (December 5, 2025): Portions of Cloudflare's network experienced significant failures attributed to configuration changes involving firewall handling, consistent with configuration propagation without interaction-scoped boundaries. Class 2: Control-Plane Trust Artifacts and Governance Dependencies These incidents

support the broader thesis that control-plane and governance-plane elements (identity, trust fabric, naming systems) dominate outage patterns even when not exhibiting pure configuration propagation: GitHub Internal Certificate Expiration (November 18, 2025): An expired internal TLS certificate broke service-to-service communication across Git operations, affecting all GitHub services. The failure required replacing the certificate and restarting impacted services. This demonstrates governance artifact failure (trust fabric) with global blast radius, even though the root cause was not configuration propagation but certificate lifecycle management. Slack Maintenance Incident (February 26, 2025): A planned maintenance action interacted with a latent caching defect, producing database overload and partial service unavailability. Control-plane operations (maintenance) interacting with latent system behavior produced widespread impact. Zoom DNS-Layer Issue (April 16, 2025): A DNS-layer problem impacted zoom.us connectivity globally. While not configuration propagation specifically, the incident demonstrates that naming and identity layers dominate failure impact at scale. Pattern Validation The breadth of incidents across vendors (AWS, Azure, Cloudflare, GitHub, Slack, Zoom), infrastructure types (hyperscale cloud, DNS resolvers, collaboration platforms, developer tools), and failure modes (DNS automation, BGP routing, certificate management, caching defects) validates that this is not vendor-specific operational failure but a cross-vertical architectural pathology. Class 1 incidents cleanly support the “unbounded propagation” claim: configuration changes lacking session- scoped boundaries cascade globally because no interaction identity constrains their blast radius. Class 2 incidents support the broader claim that control-plane and governance-plane elements have become the dominant source of systemic outages, even when the specific failure mechanism differs from pure configuration propagation. Both classes demonstrate the thesis that fragmented control planes lack architectural primitives for containing operational changes within interaction boundaries.

6.3 The Consistent Pattern

Each incident exhibits the same structural failure: 1. A configuration or policy change occurs 2. The change propagates without session-scoped boundaries 3. The blast radius is global (or region-wide) rather than contained 4. Cascading failures occur because dependent systems lose coherent state 5. Recovery requires manual intervention because automated rollback lacks deterministic semantics The pattern is systemic: Cloudflare, Microsoft, Google, and Salesforce exhibited the same failure class. Scale amplifies impact, but it is not the root cause. The missing primitive is interaction-scoped governance that contains blast radius, enables deterministic rollback, and preserves coherent state across configuration changes.

6.4 Why Governance Being External, Late, and Inferred Fails

, Late, and Inferred Fails This pattern directly validates the original RTC paper’s claim that post-hoc governance does not satisfy regulatory and operational requirements when enforcement must be real-time. When governance is externalized: Policy decisions are made by systems independent from the interaction they govern, requiring synchronization protocols that introduce latency and inconsistency risk. When governance is late: Violations occur first, then detection follows, then remediation attempts. This sequence is insufficient for safety-critical systems, regulatory compliance, and user trust. When governance is inferred: Audit evidence must be reconstructed from logs produced by fragmented systems, creating evidentiary gaps and reconstruction ambiguity. SIOAR’s architectural requirement (that policy binds to continuous interaction identity) transforms governance from externalized/late/inferred to internal/continuous/by-construction.

6.5 The IoE-Specific Hazard: Physical-World Coupling

In digital systems, control-plane failures produce service degradation, user frustration, and compliance violations. In IoE, control-plane failures couple to the physical world: Industrial automation: Equipment damage, production waste, worker safety incidents. Healthcare: Degraded care quality, adverse patient events, clinical safety violations. Smart cities: Traffic disruption, emergency response delays, civic infrastructure failures. Financial services: Market instability, settlement failures, systemic risk propagation. Remote surgery: Under 150ms round-trip with under 10ms jitter, or patients die. Coordination overhead at system boundaries produces jitter that exceeds tolerance thresholds regardless of transport capacity. Vestibular synchronization: Under 20ms motion-to-photon, or users experience physical discomfort. Fragmented control planes introduce coordination latency that produces nausea in spatial computing applications. This physical-world coupling makes SIOAR's requirement existential rather than merely operational. The cost of fragmentation is no longer measured in service-level agreement violations but in civic harm, patient safety, and systemic stability.

7. Proposed SIOAR Reference Architecture for IoE

This section outlines conceptual components and architectural invariants for SIOAR implementation across IoE domains. These descriptions are implementation-neutral by design - specific mechanisms, protocols, and technologies will vary across verticals, but the architectural requirements remain constant.

7.1 Conceptual Components

7.1.1 Interaction Identity Service

Responsibilities: - Issue and maintain continuous interaction identifiers that span all participants, devices, services, and computational resources involved in an interaction. - Bind interaction identity to cryptographic identity roots where available (device certificates, user credentials, service authentication). - Manage interaction lifecycle: initiation, participant join/leave, policy updates, orchestration changes, termination. - Provide versioned interaction state that enables deterministic rollback and audit reconstruction. Architectural requirements: - Interaction identity must be transport-agnostic (survive network handovers, protocol transitions). - Identity must be continuous (persist through orchestration changes, not re-established with each update). - Identity must be first-class (not inferred post-hoc from correlated logs).

7.1.2 Policy Authority and Enforcement Plane

Responsibilities: - Implement Zero Trust continuous evaluation: trust decisions occur throughout interaction lifetime, not only at initiation. - Maintain policy state as versioned, interaction-scoped artifacts (not global, eventually consistent state). - Enforce data residency, retention, accessibility, safety, and regulatory constraints as routing decisions rather than post-hoc verification. - Propagate policy updates within interaction scope with deterministic semantics (atomic updates, rollback guarantees). Architectural requirements: - Policy decisions must bind to interaction identity (not participant identity alone, which may be shared across multiple interactions). • Policy evaluation must be real-time (not batch, not asynchronous reconciliation). • Policy semantics must be portable across technology boundaries (enabling cross-vendor, cross-domain coordination).

7.1.3 Adaptive Routing and Orchestration Plane

Responsibilities: - Route data flows, participant connections, and computational invocations within the interaction boundary according to policy constraints. - Manage participant and device join/leave events, maintaining interaction continuity through topology changes. - Propagate orchestration updates (media additions, feature activations, topology changes) coherently across all interaction components. - Implement graceful degradation: when subsets of participants, devices, or resources become unavailable, maintain partial operation rather than global failure. Architectural requirements: - Routing decisions must be deterministic and policy-bound (not emergent from independent component optimizations). - Orchestration state must be versioned and interaction-scoped (enabling rollback, replay, and audit). - Transport selection must be transparent to policy enforcement (policy binds to interaction identity, not transport specifics).

7.1.4 Compute Control Plane

Responsibilities: - Treat computation (AI inference, analytics, transformations, decision support) as governed interaction sub-resources rather than independent systems. - Enforce model identity and version tracking: which AI models are authorized for which data, in which contexts. - Implement data permission binding: computational access to interaction data must satisfy the same policy constraints as human access. - Enforce jurisdictional constraints: where computation may execute, where results may be stored, where models may be hosted. Architectural requirements: - Compute invocation must occur within interaction policy boundary (not as parallel shadow systems). • Model provenance must be verifiable and auditable (model identity, version, training data, validation status). • Computational routing must satisfy data residency and sovereignty constraints deterministically.

7.1.5 Observability, Verification, and Audit Fabric

Responsibilities: - Provide interaction-scoped telemetry: metrics, logs, and traces that bind to continuous interaction identity. - Enable deterministic replay: reconstruct policy decisions, routing choices, and computational invocations from versioned interaction state. - Produce audit evidence by construction: compliance artifacts are generated during the interaction, not assembled retrospectively. - Support operational diagnosis: when failures occur, provide coherent visibility into interaction state, policy configuration, and coordination status. Architectural requirements: - Observability must be interaction-scoped (not only system-scoped). - Audit trails must be cryptographically verifiable and tamper-evident. - Observability infrastructure must maintain independence from monitored systems (avoiding circular dependency failures).

7.1.6 Change Safety System

Responsibilities: - Implement interaction-scoped rollout: configuration and policy changes propagate within defined interaction boundaries rather than globally. - Provide versioned policy artifacts: updates are atomic version transitions, not eventually consistent synchronization. - Enable bounded blast radius: failures in one interaction scope do not cascade to unrelated interactions. - Support deterministic rollback: failed updates can be reverted with guaranteed semantics. - Implement kill switches: emergency policy changes (security incidents, compliance violations, safety events) propagate rapidly within interaction scope. Architectural requirements: - Change propagation must be scoped and versioned. - Rollback must be deterministic and interaction-scoped. • Emergency updates must maintain policy coherence (not bypass governance).

7.2 Architectural Invariants

These invariants define what “SIOAR-compliant” means across implementation diversity: Invariant 1: Policy Binding A policy decision cannot be applied unless attached to an interaction identity that

persists through orchestration updates. This invariant prevents policy from becoming detached from the interactions it governs. Invariant 2: Computational Boundary Compute invocation cannot occur outside the interaction's policy boundary. AI models, analytics pipelines, and transformation functions must operate within the governed interaction scope, subject to the same data access, residency, and audit requirements as the interaction itself. Invariant 3: Routing Determinism Routing changes must be represented as versioned interaction state transitions, not implicit emergent behavior from independent component decisions. This invariant ensures that routing is auditable, reproducible, and policy-bound. Invariant 4: Observability Independence Observability and control infrastructure must maintain operational independence from the data-plane systems they monitor. This invariant prevents circular dependency failures that create operational blindness. Invariant 5: Change Atomicity Configuration and policy updates must be atomic transitions within interaction scope, not eventually consistent synchronization across fragmented systems. This invariant enables deterministic rollback and bounded blast radius.

8. Standards and Ecosystem Alignment

8.1 Telecom and Zero Trust Standardization Gap Is Already Being Articulated

The December 23, 2025 IETF Internet-Draft "Problem Statement of Zero Trust Deployment in Telecom Network Environments" explicitly identifies gaps in current standards, stating they "fail to establish quantitative constraints on… validation frequency, session persistence mechanisms, and incremental state synchronization." This work-in-progress draft demonstrates that telecommunications standardization bodies are encountering SIOAR requirements in production deployment contexts. The gaps being articulated (validation frequency, session persistence, state synchronization) map directly to SIOAR's continuity and policy enforcement constraints.

8.2 Where SIOAR Conversations Can Live

SIOAR's architectural requirements span multiple standardization contexts. Rather than attempting to create a new standards body or protocol specification, SIOAR can inform work across existing venues: Internet Engineering Task Force (IETF): Security area (Zero Trust, identity protocols), applications area (real-time communication, web infrastructure), operations area (observability, configuration management). 3rd Generation Partnership Project (3GPP): Mobile telecommunications architecture, network slicing, edge computing integration. Cloud Native Computing Foundation (CNCF): Kubernetes orchestration, WebAssembly runtime governance, service mesh policy enforcement. Industry alliances for IoT interoperability: OPC Foundation (industrial automation), HL7 (healthcare interoperability), MQTT consortium (IoT messaging). Smart city and civic technology working groups: Open & Agile Smart Cities (OASC), Smart Cities Council standards development. The positioning should be as a shared architectural constraint, not a protocol specification or turf claim. SIOAR defines what implementations must accomplish (continuous interaction identity, transport-agnostic policy binding, computational transparency) without prescribing specific mechanisms. This approach enables domain-specific instantiations (telecommunications implementations will differ from healthcare implementations will differ from industrial implementations) while maintaining cross-vertical architectural coherence.

8.3 Cross-Vertical Synthesis Opportunity

The “signals from the field” methodology in Section 5 demonstrates that independent voices across verticals are describing fragments of the same systemic issue. This creates an opportunity for deliberate cross-vertical synthesis: Convene representatives from telecommunications, cloud infrastructure, healthcare, financial services, smart cities, and industrial automation to explicitly map domain-specific pain points onto shared architectural requirements. Frame SIOAR not as a solution imposed from above but as a recognition of convergent bottom-up requirements that deserve explicit architectural attention. Identify low-hanging interoperability fruit: areas where SIOAR alignment would enable immediate cross-domain coordination (e.g., healthcare-telecom for remote patient monitoring, smart city-transportation for autonomous vehicle infrastructure, industrial-cloud for distributed manufacturing orchestration).

8.4 Historical Precedent: When Coordination Became Infrastructure

The closest historical analogy to IoE’s coordination crisis is electrical grid development in the late 19th and early 20th centuries, before coordinated load management, standardized frequency, and centralized dispatch existed.

8.4.1 The Early Configuration

Independent power generators were added opportunistically as electrification expanded. Loads connected locally without global coordination. Each system worked locally and temporarily. Failures appeared random, cascading, and inexplicable. Engineers focused on improving generators and strengthening transmission lines, not coordination.

8.4.2 The Hidden Scaling Problem

As electrification scaled beyond local deployments, load interactions multiplied. Phase mismatches caused destructive oscillations. Frequency drift cascaded across regional boundaries. Local optimizations increased global instability. Critically: adding more generators made failures more likely, not less. At small scale, this was survivable. At municipal scale, it created persistent operational challenges. At regional and national scale, it became existential. The system was not failing because turbines were inadequate or transmission infrastructure was insufficient. It was failing because coordination was implicit rather than structural.

8.4.3 The Recognition Threshold

The breakthrough was not better generators, stronger transmission lines, or increased capacity. It was the recognition that coordination itself is infrastructure requiring first-class architectural attention. This led to: - Standardized frequency across interconnected grids • Synchronized phase relationships • Centralized dispatch and load balancing • Fault isolation zones • Continuous monitoring and control systems The control plane became first-class architecture rather than operational afterthought. Before that recognition, failures looked like technical glitches, operational mistakes, or inadequate equipment. Afterward, they were recognized as architectural inevitabilities resulting from uncoordinated system expansion.

8.4.4 Why Institutional Recognition Lagged Evidence

Engineers and operators saw individual symptoms: generators overheating, inexplicable outages, unpredictable cascading failures. What they initially did not see was that: - Coordination itself had become the dominant system constraint - Adding capacity without unified governance was making

things worse - The original abstraction (“electricity delivery”) had collapsed under coordination complexity - Optimization at component level was driving instability at system level It took repeated regional blackouts, significant economic damage, and visible public impact to force institutional recognition that the problem was architectural rather than operational. The transition from “we need better equipment” to “we need coordination infrastructure” required crossing an evidence threshold where denial became more costly than acknowledgment.

8.4.5 The Structural Parallel to IoE

The mapping between early electrical grids and fragmented IoE control planes is structural, not metaphorical:same failure patterns, same solution class. The pattern is identical: fragmentation was rational at small scale, became problematic at municipal scale, and proved architecturally untenable at national scale. The trade-off reversed as coordination complexity grew faster than the benefits of independent optimization.

8.4.6 Implications for IoE's Current Phase

IoE is at the same inflection point: past the evidence threshold but before broad institutional recognition. The evidence is visible: - Cross-vertical outages increasing in frequency and blast radius - AI agent deployments fragmenting context across system boundaries - Accessibility requirements exceeding bolt-on implementation capacity - Energy consumption curves bending toward infrastructure limits - Regulatory timelines extending as compliance mechanisms prove inadequate What is not yet widely recognized is that coordination is the architectural primitive requiring first-class attention, not transport capacity or computational efficiency. The same cognitive pattern that delayed electrical grid coordination (“we just need better components”) is delaying IoE coordination infrastructure recognition. Historically, this recognition phase always follows rather than precedes the evidence accumulation. The power grid transition required visible failures and economic consequences before coordination became a funded, standardized, first-class infrastructure concern. The regulatory retreats documented in Section 1.9 (EU Omnibus delays, US OBBBA lobbying, extended compliance timelines) signal that the recognition threshold is being crossed. Regulators have discovered that existing architectures lack mechanisms to satisfy mandated requirements. Industry has requested timeline extensions because they do not yet have the architectural primitive. This is less a softening of requirements than institutional recognition of an implementation bottleneck, and it opens an architecture window. The question is not whether coordination will become first-class infrastructure for IoE. Historical precedent demonstrates it must. The question is whether this transition occurs through deliberate architecture or through accumulated failure costs.

9. Migration and Adoption Strategy

Building on the electrical-grid parallel developed in Section 8.4, this section describes how interaction-scoped governance is likely to migrate into production ecosystems. Large-scale control abstractions rarely emerge through voluntary market coordination alone. Historically, when fragmented systems reach a point where failure becomes unacceptable, adoption is led by actors for whom reliability, determinism, and governance are not optional. The electrical grid followed this pattern: early experimentation and private enterprise existed from the outset, but coherent, large-scale control only emerged when military installations and public infrastructure demanded predictable behavior under load, safety constraints, and failure conditions. Civilian authorities subsequently codified these requirements through regulation and procurement, while private industry supplied, refined, and scaled the resulting systems. This sequence is not a judgment on private enterprise. It is a structural pattern of infrastructure adoption. Markets innovate continuously, but authority acts when fragmentation becomes a systemic risk. Once authority establishes requirements, suppliers align, and diffusion accelerates

across civilian domains. The adoption of interaction-scoped governance follows the same trajectory. The question is not whether private platforms are capable of building such systems, but where the requirement first becomes unavoidable.

9.1 Principles (Recognition to Realization)

(Recognition to Realization) The original RTC paper's "recognition to realization" framing applies directly to IoE contexts: "Migration must preserve continuity and interoperate with fragmented architectures during transition." Migration strategies differ fundamentally based on ecosystem control. Civilian IoE contexts face legacy interoperability requirements. Existing systems represent massive capital investment, operational expertise, and regulatory certification. Migration in these contexts must: Preserve continuity: Interaction identity can be overlaid on existing systems without requiring immediate replacement of fragmented components. Interoperate during transition: SIOAR-aware components must coordinate with non-SIOAR legacy systems during multi-year migration timelines. Provide incremental value: Each migration step must deliver measurable operational or compliance improvements, not require complete transformation before benefits accrue. Maintain safety and compliance: Migration cannot degrade existing safety controls or regulatory compliance; these must be preserved or improved at each step. Controlled-ecosystem contexts face different constraints. Defense, coalition, and sovereignty-sensitive deployments often control their technology stack, can mandate architectural requirements through procurement, and can implement greenfield systems without legacy interoperability burdens. For these contexts, the migration question is not "how do we overlay SIOAR on fragmented legacy?" but "how quickly can we deploy native implementations?" The distinction matters because controlled ecosystems can move faster and serve as proving grounds for patterns that later propagate to civilian contexts.

9.2 Adoption Wedges

Adoption does not proceed uniformly by industry. It proceeds by authority over architecture. Practical adoption begins where fragmentation's costs are most acute and where governance is already enforced structurally. Ordered by ability to mandate: 1. Military and coalition systems (mandate by force): These environments encounter interaction-scoped governance requirements earliest because failure is unacceptable. Coalition interoperability and Zero Trust OT requirements force implementations that satisfy SIOAR constraints under operational pressure, even without civilian IoE vocabulary. 2. Sovereignty-sensitive and regulated deployments (mandate by law): Civilian authorities codify proven patterns into regulation and procurement. When policy must be enforced in real time (not reconciled after violations), and when enforcement boundaries align with jurisdictional rather than platform boundaries, interaction-scoped governance becomes necessary by construction. 3. Hyperscale platforms (economic compulsion): Large platforms are embedded suppliers to defense and sovereign customers. They cannot lag mandated requirements without fragmenting their own architectures or losing contracts. Adoption here is accelerated by contractual and compliance coupling to the first two classes. 4. Downstream civilian adoption (inheritance through suppliers and standards): Enterprises and vertical platforms adopt last, inheriting the architecture through vendor convergence, standards development, and procurement requirements once upstream alignment has occurred. Within civilian ecosystems, early wedges are domains where fragmentation's costs are most visible and SIOAR's benefits are immediate: RTC sessions with AI features: Real-time communication platforms adding transcription, summarization, translation, accessibility features, or safety controls. These systems already experience fragmentation pain (features exist as parallel systems requiring coordination) and have user-facing quality metrics that make improvements visible. Telecom Zero Trust

validation loops: Mobile network operators implementing continuous validation for subscriber sessions across network slices, service function chains, and edge computing resources. Regulatory pressure and operational complexity create urgency; standards work is already identifying SIOAR-adjacent requirements. Multi-cloud control-plane governance: Organizations operating across multiple cloud providers encountering configuration propagation failures, circular dependency blindness, and unbounded blast radius from policy changes. Recent cross-vertical incidents demonstrate the operational pain. Digital health AI workflow orchestration: Healthcare institutions integrating AI diagnostics, treatment recommendations, or operational optimization while maintaining clinical safety, privacy compliance, and jurisdictional constraints. Regulatory scrutiny and patient safety risk create forcing functions. Each wedge provides: Clear value proposition: Reduced incident frequency, simpler compliance, lower AI compute waste, faster feature velocity. Contained scope: Initial implementation in bounded interaction domains rather than requiring system-wide transformation. Demonstrable results: Metrics (incident blast radius, policy propagation latency, audit reconstruction time, computational efficiency) that validate the approach.

9.3 Delivery Mechanisms

SIOAR adoption begins differently in mandated ecosystems versus civilian ecosystems. Native implementation (controlled ecosystems): For greenfield deployments or major architecture refreshes in controlled ecosystems, SIOAR can be implemented as a first-class architectural primitive from initial design: - Interaction identity as foundation: System architecture begins with continuous interaction identity as the primary governance unit, not as an overlay on request-response or connection-based primitives. - Policy-by-construction: Residency, access control, retention, and audit requirements are enforced through interaction-scoped routing decisions, not verified through post-hoc reconciliation. - Compute governance native: AI inference, analytics, and transformation services operate inside interaction policy boundaries by design, with explicit provenance and authorization. - Deterministic rollout: Configuration and policy changes propagate within interaction-scoped boundaries with versioned state and rollback guarantees. Overlay migration (civilian contexts): For civilian contexts with legacy interoperability requirements, SIOAR adoption can begin with overlays and incrementally progress toward native implementations: Phase 1: Interaction identity overlay Implement interaction identity as a metadata layer that spans existing fragmented systems without replacing them. This provides audit coherence, policy correlation, and observability unification before modifying core systems. Phase 2: Policy versioning and propagation constraints Introduce versioned policy artifacts and interaction-scoped rollout boundaries. This reduces blast radius and enables deterministic rollback without requiring full SIOAR implementation. Phase 3: Compute routing governance Bring AI and analytics features inside interaction policy boundaries. This addresses the efficiency paradox and enables governance of computational resources. Phase 4: Native SIOAR implementation For new systems or major architecture refreshes, implement continuous interaction identity, transport-agnostic policy binding, and computational transparency as first-class architectural primitives. Migration timelines vary by ecosystem control and legacy burden: Defense and coalition systems: 2-5 years for new programs; 5-10 years for legacy modernization. Procurement authority enables faster adoption than civilian consensus processes. Sovereignty-sensitive deployments: 3-7 years, driven by regulatory codification and procurement gates. Hyperscale platforms: 2-5 years, accelerated by defense and sovereign requirements they must support. Telecommunications: 3- 7 years (infrastructure refresh cycles, standards development timelines) Healthcare: 5-10 years (regulatory certification, clinical validation, interoperability complexity) Financial services: 3-6 years (regulatory approval, systemic risk management, vendor ecosystem coordination) Smart cities: 7-15 years (public procurement cycles, cross-agency coordination, civic governance processes) Industrial automation: 5-15 years (equipment lifecycle, safety certification, operational risk management) The timeline differential between controlled and civilian ecosystems is structural, not merely political. Defense can mandate; civilian contexts must

persuade. This makes defense and sovereignty-sensitive deployments natural proving grounds for patterns that later propagate through standards development and demonstrated operational value.

9.4 Event Horizon Analysis: The 2025-2040 Infrastructure Requirement

9.4.1 IoT Doesn't Just Need Frameworks

- It Needs Pipes and Power Discussions of IoE architecture often focus on software frameworks, protocols, and orchestration platforms while ignoring the physical infrastructure reality: IoT requires massive deployment of actual network connectivity and electrical power to billions of edge devices. Current estimates project 50-75 billion connected IoT devices by 2030, scaling to 125-150 billion by 2040. These devices require: Network connectivity: 5G small cells, Wi-Fi 7 access points, satellite ground stations, fiber backhaul, edge data centers Electrical power: Grid connections, renewable generation, battery systems, power distribution infrastructure Real estate and installation: Cell towers, equipment shelters, trenching/conduit for fiber and power, facility cooling This is infrastructure build-out, not a software rollout, closer to electrification or highway construction than a typical IT migration.

9.4.2 Additive vs. Subtractive: Moon Launch Economics

The question is not whether this infrastructure will be built, but what coordination architecture it must support. Additive Approach: Fragmented Coordination at Scale If IoE deploys with fragmented architectures where identity, policy, orchestration, and compute remain independently coordinated: Coordination overhead scales multiplicatively with device count, as demonstrated in Section 2.2:
- 2025: 30B devices $\rightarrow$ ~10^12 coordination surfaces (manageable with current approaches)
- 2030: 75B devices $\rightarrow$ ~10^14 coordination surfaces (approaching operational limits)
- 2040: 150B devices $\rightarrow$ ~10^16 coordination surfaces (fundamentally unmanageable)

The infrastructure investment required to sustain fragmented coordination: For 150 billion devices with multiplicative coordination:
- Processing overhead: Each coordination requires CPU cycles, memory, network bandwidth
- Network capacity: Policy synchronization, identity verification, state reconciliation traffic scales with coordination surfaces, not device count
- Energy consumption: Coordination overhead translates directly to computational load and heat dissipation
- Operational complexity: Human operators must manage systems whose coordination surfaces exceed human cognitive capacity Quantifying the “Moon Launch” Investment: Assume each coordination consumes on average:
- 10ms CPU time (across distributed systems) • 1KB network transfer (for policy/state sync) • 0.1 joules (watt-seconds) of energy To translate per-coordination costs into continuous infrastructure requirements, assume a conservative cadence of one coordination transaction per surface every 3 hours (τ = 10,800 seconds). For 10^16 coordination surfaces operating continuously: • Processing: 10^16 surfaces × (10ms / 3h) ≈ 9×10^9 CPU-core equivalents (≈9 billion CPU-years/year) • Network: 10^16 surfaces × (1KB / 3h) ≈ 0.9 PB/sec coordination traffic • Energy: 10^16 surfaces × (0.1J / 3h) ≈ 9×10^10 W (≈90-100 GW) continuous power for coordination overhead alone Reading the arithmetic (units and conversions) Each line has the same structure: (coordination surfaces) × (per-transaction cost) ÷ (seconds between transactions). CPU time yields CPU-seconds per second, which is directly interpretable as an average core-equivalent requirement. Bytes per second is network throughput. Joules per second is watts, which is continuous power. Two conservative choices are embedded: (i) the cadence assumes one coordination transaction per surface every three hours, which is slower than what continuous trust evaluation, live captions, and AR synchronization imply; (ii) the model counts only machine coordination overhead and omits the human and organizational coordination burden that appears as incident response effort and governance delay once state becomes inconsistent. For

comparison:
- Global data centre electricity consumption (2024): ~415 TWh/year (≈47 GW average, IEA estimate)
- Fragmented IoE coordination alone would require 2x current global data center capacity This does not include the actual IoT workloads - only the overhead of coordinating fragmented governance. The investment implication: To sustain fragmented IoE at 2040 scale requires data center capacity equivalent to building 1,000 new hyperscale facilities (100MW each) dedicated solely to coordination overhead. At ~$2-3 billion per hyperscale facility: $2-3 trillion in coordination infrastructure alone. SIOAR Approach: Interaction-Scoped Coordination With continuous interaction identity, coordination scales additively with participants:
- 2025: 30B devices → ~10^10 coordination surfaces
- 2030: 75B devices → ~10^10.5 coordination surfaces
- 2040: 150B devices → ~10^11 coordination surfaces Same device count, 100,000x fewer coordination surfaces at 2040 scale. Using the same per-coordination costs:
- Processing: 10^11 surfaces × (10ms / 3h) ≈ 9×10^4 CPU-core equivalents (≈90k CPU-years/year) (vs ≈9×10^9)
- Network: 10^11 surfaces × (1KB / 3h) ≈ 9 GB/sec coordination traffic (vs ≈0.9 PB/sec)
- Energy: 10^11 surfaces × (0.1J / 3h) ≈ 0.9 MW continuous power (vs ≈90-100 GW) Infrastructure investment: Standard cloud capacity growth, no dedicated coordination infrastructure required. Incremental cost: <$10 billion (vs $2-3 trillion).

9.4.3 The Green Imperative

The energy differential is not merely economic - it is environmental: Fragmented approach: ~100GW coordination overhead = annual CO2 emissions equivalent to ~390 million metric tons (assuming ~445 gCO2/kWh global average electricity carbon intensity, IEA estimate) SIOAR approach: ~1MW coordination overhead = annual CO2 emissions equivalent to ~3,900 metric tons The ratio: 100,000:1 reduction in coordination-related carbon emissions This differential matters because: 1. Regulatory pressure: EU taxonomy, corporate ESG reporting, carbon pricing mechanisms increasingly penalize inefficient architectures 2. Operational cost: Energy is 40-60% of data center OpEx; coordination overhead directly impacts profitability 3. Social license: Climate commitments require demonstrable efficiency; fragmented IoE is incompatible with net-zero targets

9.4.4 Why "Additive" Cannot Scale

The 2025-2040 event horizon analysis reveals why incremental patching and additive architectures are not merely suboptimal - they converge on an envelope where coordination overhead exceeds plausible physical and economic constraints: Network physics: ~0.9PB/sec coordination traffic (under a conservative 3-hour cadence) is already on the order of today’s total global internet traffic; tighter cadences push coordination overhead into multi-PB/sec regimes. Energy physics: ~90-100GW continuous load would be dedicated to coordination overhead alone (IEA order-of-magnitude comparisons place global data-centre average power in the tens of gigawatts); tighter cadences push this toward the terawatt scale. Economic physics: Multi-trillion-dollar coordination infrastructure (on the order of ~1,000 additional 100MW facilities in the conservative case) is an economic non-starter for coordination overhead, not payload. The path that continues fragmentation does not lead to degraded performance; it leads to an event horizon where coordination overhead dominates and governance becomes infeasible. The system cannot be built or operated sustainably at 2040 scale. SIOAR is not an optimization; it is the difference between possible and impossible at 2040 scale.

9.5 What SIOAR Does Not Claim: Limitations and Non-Applicability

To maintain intellectual honesty and prevent misapplication, it is critical to delineate what SIOAR does not claim and where it does not apply.

9.5.1 Not a Complete Solution

SIOAR addresses coordination and control-plane fragmentation specifically. It does not solve: Data-plane performance optimization: Network bandwidth, latency, jitter, packet loss require transport-layer solutions (better protocols, more capacity, improved routing algorithms). Resource efficiency: CPU utilization, memory management, power consumption require operating system and runtime optimization. Feature functionality: What capabilities a system offers (video quality, AI model accuracy, sensor precision) remain separate engineering challenges. SIOAR's claim is narrow: when coordination complexity becomes the limiting factor (when the system can technically perform the required functions but cannot coordinate them coherently) interaction identity provides the missing architectural constraint.

9.5.2 Not Universally Applicable

Simple systems with minimal coordination requirements may not benefit from SIOAR's architectural overhead. When SIOAR is not needed: - Single-participant interactions (standalone device operation, offline applications) - Trivial coordination (request-response with no policy, no state, no multi-step workflows) - Static systems (no policy changes, no orchestration updates, no dynamic routing) - Short-lived, fire-and-forget interactions (metrics collection, log shipping, batch uploads) When SIOAR becomes necessary: - Multi-participant coordination across organizational or technical boundaries - Dynamic policy enforcement (continuous trust evaluation, mid-interaction policy changes) - Long-lived interactions with state continuity requirements - Compliance-critical or safety-critical contexts where post-hoc reconciliation is insufficient - AI and automation features that must operate within governed interaction boundaries The framework applies when coordination costs exceed fragmentation benefits - not to all systems universally.

9.5.3 Not Immediate

Migration timelines are measured in years to decades, not quarters. The constraint is not awareness; it is the time required for procurement cycles, certification regimes, installed-base interoperability, and operational retraining. Defense and coalition systems: 2-5 years for new programs (mandate authority, greenfield deployment); 5-10 years for legacy modernization Sovereignty-sensitive deployments: 3-7 years (regulatory codification, procurement gates, jurisdictional enforcement requirements) Telecommunications: 3-7 years (infrastructure refresh cycles, standards development timelines) Healthcare: 5-10 years (clinical validation, regulatory certification, interoperability complexity) Financial services: 3-6 years (regulatory approval, systemic risk management, vendor ecosystem coordination, legacy system integration) Smart cities: 7-15 years (public procurement, civic governance, cross-agency coordination, infrastructure deployment) Industrial automation: 5-15 years (equipment lifecycle, safety certification, operational risk management, workforce training) SIOAR is architectural direction for next-generation systems, not a tactical fix for current production problems. Organizations should begin alignment now for systems that will operate in 2030-2040, not attempt to retrofit every existing system immediately.

9.5.4 Not Vendor-Specific or Implementation-Prescriptive

or Implementation-Prescriptive No single implementation validates or invalidates SIOAR's architectural claims because SIOAR is a constraint, not a specification. This means: - Different implementations will emerge across domains (telecom SIOAR ≠ healthcare SIOAR in mechanism) - No reference implementation can prove or disprove the framework (success is measured by whether implementations satisfy the three constraints) - Vendor-specific solutions can be SIOAR-aligned without being identical - Open-source and proprietary implementations can coexist Evaluation criteria: Does an implementation provide continuous interaction identity, transport-agnostic policy binding, and computational transparency? If yes, it is SIOAR-aligned regardless of technical mechanism. If no, it remains fragmented regardless of marketing claims.

9.5.5 Not Regulation-Replacement

SIOAR enables structural compliance enforcement but does not define what should be enforced. What SIOAR provides: Architectural mechanisms to enforce policy deterministically within interaction boundaries What SIOAR does not provide: The policies themselves (privacy rules, safety constraints, accessibility requirements, security controls) Regulators, standards bodies, and organizational governance processes define requirements. SIOAR provides the architectural foundation to enforce those requirements structurally rather than procedurally. The distinction matters: SIOAR does not make claims about what privacy policies should be, what accessibility thresholds are appropriate, or what security controls are necessary. It claims only that whatever policies are defined, they can be enforced more coherently and deterministically through interaction-scoped architecture than through fragmented post-hoc reconciliation.

9.5.6 Not Physics Violation

SIOAR does not eliminate fundamental limits: Speed of light: Latency for global coordination is bounded by physics; SIOAR reduces coordination surfaces, not photons' travel time Computational complexity: NP-hard problems remain NP-hard; SIOAR does not make intractable algorithms tractable Heisenberg uncertainty: Observability/control trade-offs in quantum systems are unchanged What SIOAR addresses is coordination complexity that is architecturally induced, not physically fundamental. When systems coordinate poorly because of fragmented identity and policy contexts, that is solvable through architecture. When systems coordinate poorly because they are physically distant, that requires transport solutions, not architectural changes.

9.6 Emerging Implementations: Solutions Likely Already Exist

This paper is implementation-neutral. But the forcing functions are strong enough that a question arises: has anyone already built this? The answer is almost certainly yes, in partial form, in high-pressure domains, without the general vocabulary this paper provides. What is missing is not invention but recognition.

9.6.1 Why partial implementations are plausible

Three forces make partial implementations plausible even without an industry-wide vocabulary: - Operational survival: Large platforms already treat control-plane blast radius as an existential risk. Interaction-scoped rollout, bounded state, and deterministic rollback are the durable answers to the incident class surveyed in Section 6. • Feature coupling: AI-generated captions, translation, summarization, and accessibility overlays are not independent add-ons. They consume the same streams, must maintain the same timing, and must obey the same governance envelope as the interaction itself. • Regulatory determinism: Residency, retention, and accessibility obligations increasingly require real-time enforcement. In systems with tight latency budgets, post-hoc reconciliation is not a compliance strategy.

9.6.2 What partial implementations would look like

If SIOAR-aligned systems exist in development or early deployment, they would tend to exhibit the same observable characteristics, regardless of mechanism: - A stable interaction identifier that survives participant changes and orchestration updates. - Policy state that is versioned and bound to the interaction, not reconstructed from global logs. - Compute invocations (AI inference, analytics, transformations) executed inside the same governed scope as the interaction, with explicit provenance. - Rollouts and rollbacks scoped to bounded interaction domains rather than global fleet state.

9.6.3 Where to look first

The earliest emergence is most likely in domains where users notice coherence failures immediately and where compliance demands are already embedded: RTC contexts - Conferencing and contact-center stacks where captions, translation, and recording behave as first-class session resources, not parallel services that attach late. IoT and OT contexts - Industrial platforms where machine identity, operator authority, and control provenance are bound to a production interaction that persists across protocol heterogeneity. Enterprise AI agents - Workflow assistants that maintain continuous context across multiple application boundaries while remaining auditable and policy-governed. Data sovereignty and Zero Trust • Continuous authorization systems that adjust permissions mid-interaction based on device posture and network context without forcing session termination and re-initiation.

9.6.4 The accessibility forcing function, revisited

, revisited Accessibility overlays (captions, signing avatars, audio description) are a sharp forcing function because they combine: - tight latency budgets, - continuous synchronization, - selective delivery (only to those who request it), - multi-jurisdiction governance, - and explicit audit obligations. Meeting these requirements through fragmented coordination adds not just latency but governance ambiguity. If an interaction changes topology mid-stream, the accessibility computation cannot lose context and re-bootstrap with a new identity or a new policy context. This is precisely the continuity constraint.

9.6.5 Why this matters for ecosystem planning

Whether implementations already exist or must be built, the ecosystem outcome depends on recognition and interoperability: - If interaction-scoped approaches emerge as interoperable patterns, standards bodies can encode common semantics while allowing domain-specific mechanisms. - If they emerge as closed, platform-specific solutions, the risk is lock-in for governments and enterprises and an increase in cross-platform coordination cost. - If they do not emerge quickly enough, the industry will continue to pay the coordination tax in outages, compliance friction, and energy overhead.

9.6.6 Preparing to evaluate solutions

Organizations can prepare without betting on any vendor or mechanism: - Document domain-specific requirements against the three constraints: continuity, transport agnosticism, and computational transparency. - Evaluate platforms and proposals against the architectural invariants in Section 7.2. - Engage in standards venues already discussing adjacent gaps (for example, validation frequency, session persistence, and incremental state synchronization in telecom contexts). • Pilot interaction-scoped governance in bounded domains where the cost of fragmentation is already measurable. The practical objective is to make interaction identity as a first-class primitive an evaluatable criterion, so that when implementations surface (or are built), the ecosystem can recognize conformance rather than reverse-engineering it after failures.

10. Conclusion

10.1 The Core Claim: IoE Version

Internet of Everything systems are experiencing the same architectural trade-off reversal documented in “When Independent Constraints Collapse: Why Real-Time Systems Need a Missing Control Layer”: fragmentation’s coordination costs now exceed its benefits under converged pressures. When Zero

Trust continuous validation, AI-driven compute routing, accessibility and safety mandates, data sovereignty enforcement, and next-generation transport capabilities must all be satisfied simultaneously within real-time interaction boundaries, fragmented architectures hit coordination limits. They hit these limits not because any component breaks but because coordination across fragments scales combinatorially while control mechanisms remain linear. The efficiency paradox ensures that technology improvements accelerate rather than resolve this gap. AI deployment scales interaction volume and boundary crossings faster than efficiency gains. Transport improvements expose coordination overhead that was previously masked. Each technology generation amplifies the crisis rather than resolving it. The coordination tax is the unifying failure mode. Latency is a tax on user experience. Energy consumption is a tax on sustainability. Lost potential is a tax on human capability. Fragmented architectures pay this tax at every system boundary. Session-scoped governance pays it once. The ceiling is coordination, not transport. We are not being held back by physics. We are being held back by architecture that treats system boundaries as the unit of governance instead of sessions. Therefore, IoE requires the same missing architectural constraint identified for RTC: a control layer capable of preserving continuous interaction identity across orchestration changes, policy enforcement decisions, and computational routing. This is an architectural constraint, not a protocol draft or vendor platform: interaction-scoped governance that preserves continuous identity across orchestration, policy enforcement, and compute.

10.2 The Evidence Pattern

Cross-vertical incidents from 2025 demonstrate this pattern consistently: Cloudflare (November 18, 2025): Oversized control-plane configuration artifact triggered fleet-wide disruption. Azure Front Door (October 29, 2025): Configuration change caused fleet-wide node failures. Google Cloud IAM (June 12, 2025): Policy metadata corruption propagated globally within seconds, crashed Service Control nodes in loop. Heroku (June 2025): Unexpected environment change violated deployment control expectations. Microsoft Teams (December 2025): Routing misconfiguration isolated healthy infrastructure. AI agent deployment exhibits the same pattern: context fragments across system boundaries, multi-step workflows fail, enterprises discover that “AI doesn’t speak unstructured data without an activation layer that spans fragmented systems.” Regulatory retreat confirms the pattern: the EU Digital Omnibus delays high-risk AI compliance to December 2027; the U.S. One Big Beautiful Bill’s AI moratorium was stripped by 99-1 Senate vote; the industry spent $50 million lobbying for time because they do not have the primitive. In each case, the failure was control-plane fragility, not data-plane capacity. The pipes worked; the valve system failed. Contemporary voices from telecommunications, real-time communication, cloud infrastructure, healthcare, financial services, smart cities, and industrial automation independently describe fragments of the same systemic issue. These signals converge: what telecommunications calls “session persistence” and “incremental state synchronization,” what healthcare calls “orchestration and trust,” what cloud infrastructure calls “configuration safety” and “circular dependency elimination,” what AI practitioners call “the activation layer that spans fragmented systems” are different names for the same missing control layer.

10.3 The Path Forward Is Subtractive

The original RTC paper’s key insight remains definitive: “The path forward is subtractive, not additive.” More orchestration tools, more policy synchronization protocols, more audit correlation systems will not solve a problem rooted in missing architectural constraint. Additive approaches increase the number of fragments requiring coordination, accelerating the very pathology they attempt to address. SIOAR (Single-Interaction Orchestration and Adaptive Routing) proposes the subtractive move: collapse

fragmented control contexts into a unified interaction surface. Make interaction identity a first-class architectural primitive. Bind policy to continuous interaction identity rather than externalizing it across fragments. Bring computation inside interaction policy boundaries rather than treating it as parallel infrastructure requiring asynchronous coordination. This is an architectural constraint, not an implementation prescription. Telecommunications, healthcare, finance, smart cities, and industrial automation will instantiate SIOAR differently according to their domain-specific requirements, regulatory frameworks, and technology ecosystems. But they share the same requirement: governance must be structural, not procedural; continuous, not eventual; by-construction, not post-hoc. The regulatory pause is not a reprieve. It is the final window. Organizations can implement unified session governance while compliance timelines permit iteration, or attempt retrofit when enforcement resumes under deadline pressure. The ceiling is coordination, not transport. We have the pipes. We have the compute. We have the models. We have the displays. We lack the session primitive that lets them compose into governed, coherent experiences. The sky has a ceiling we built. We know how to remove it. The architecture the industry needs exists, in partial form, in production stacks, waiting to be recognized and generalized. The next 12-18 months will determine whether we articulate and deploy it, or continue paying the coordination tax until the infrastructure fails.

Appendix A: Named Citations Bibliography

Primary architectural reference: Thomas Rocha III, “When Independent Constraints Collapse: Why Real- Time Systems Need a Missing Control Layer” (December 2025). Available at: https://doi.org/10.5281/zenodo.18001608 Telecommunications and network infrastructure: Ivo Ivanov, CEO of DE-CIX, quoted in TelecomsTech News, “Telecoms in 2025: AI and next-gen infrastructure among key trends” (December 2025). Xuan Si, “Problem Statement of Zero Trust Deployment in Telecom Network Environments,” IETF Internet-Draft (December 23, 2025; work in progress). IETF Datatracker. Real-time communication: Tsahi Levent-Levi, “Automatically collecting webrtc-internals with rtcstats,” BlogGeek (November 7, 2025). Covers RTC observability best practices. Microsoft Learn, “Service health and continuity,” Microsoft 365 Service Descriptions (accessed December 2025). Cloud and infrastructure incidents: Adrian Cockcroft, AWS outage commentary on circular dependency and operational blindness, LinkedIn (October 2025, surfaced in December 2025). ThousandEyes, “AWS Outage Analysis: October 20, 2025” (October 2025). Analysis of DynamoDB DNS automation race condition and cascading failures. AWS, “Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-EAST-1) Region” (October 2025). Official post-event summary. Cloudflare, “Cloudflare incident on November 18, 2025” (November 2025). The Cloudflare Blog. Cloudflare, “Cloudflare 1.1.1.1 incident on July 14, 2025” (July 2025). The Cloudflare Blog. Cloudflare, “Dashboard and API outage” (September 12, 2025). The Cloudflare Blog. Cloudflare, “Partial network failures tied to combined factors and configuration logic” (December 5, 2025). The Cloudflare Blog. ThousandEyes, “Azure Front Door disruption analysis” (October 29, 2025). Google Cloud, “Google Cloud Service Health: June 12, 2025 incident” (June 2025). Heroku (Salesforce), “June 10 outage summary” (June 15, 2025). GitHub, “GitHub Availability Report: November 2025” (November 2025). The GitHub Blog. Slack, “Slack Status: February 26, 2025 incident” (February 2025). Network World, “Zoom global outage attributed to DNS-layer issue” (April 16, 2025). AI agent deployment: MuleSoft, “2025 Connectivity Benchmark Report” (February 2025). Reports that organizations use 897 apps on average, with 29% connected. Confluent, “The AI Silo Problem: How Data Streaming Can Unify Enterprise AI Agents” (2025). Salesforce engineering documentation on Agentforce Atlas reasoning engine behavior. Energy and infrastructure: International Energy Agency (IEA), “Energy and AI” report (2025). Provides data center electricity consumption projections. Pew Research Center, “What we know about energy use at U.S. data centers amid the AI boom” (October 2025). Bloomberg, “How AI Data Centers Are Sending Your Power Bill Soaring” (September 2025). Regulatory and policy: European Commission, “Digital Omnibus on AI Regulation Proposal” (November 19, 2025). Jones Day,

“EU Digital Omnibus: How EU Data, Cyber, and AI Rules Will Shift” (December 2025). Issue One, “As Washington Debates Major Tech and AI Policy Changes, Big Tech’s Lobbying is Relentless” (July 2025). Provides $50 million lobbying figure. Wikipedia, “One Big Beautiful Bill Act” (December 2025). Provides legislative history and timeline. The White House, “Ensuring a National Policy Framework for Artificial Intelligence” Executive Order 14365 (December 11, 2025). Austria Supreme Court (Oberster Gerichtshof), Schrems/noyb v. Meta judgment (November 26, 2025; publicly released December 18, 2025). Rules that GDPR requires continuous consent verification for personalized advertising, not point-in-time consent at account creation. Defense and alliance zero trust: U.S. Department of Defense Chief Information Officer, “Zero Trust for Operational Technology Activities and Outcomes” (November 2025). Specifies 105 ZT activities for OT: 84 target-level, 21 advanced-level; seven pillars including automation/orchestration. DefenseScoop, “Pentagon posts guidance on implementing zero trust for operational technology” (December 1, 2025). Covers DoD OT guidance. DefenseScoop, “Pentagon plans to publish zero trust strategy 2.0 in early 2026” (December 9, 2025). Features Randy Resnick remarks at DefenseTalks. archTIS, “What is Zero Trust Data Format (ZTDF)?” (October 2025). Describes ZTDF ratified by NATO CCEB as a data-centric security wrapper. Virtru, “Zero Trust Data Format (ZTDF): The Clear Path to Data Interoperability” (February 2025). Explains how ZTDF bridges US IC/DoD and NATO STANAGs. Virtru, “Operation HIGHMAST Validates ACP 240: The New Standard Powering Coalition Data Sharing” (November 2025). Documents UK carrier strike group validation of ACP 240 under operational conditions. GlobeNewswire/Virtru, “Virtru Data Security Platform Enables ACP 240 Adoption Among Allied Nations and Technology Partners” (October 15, 2025). Announces ACP 240 adoption by NATO and US Joint Chiefs of Staff. CISA, “Zero Trust Architecture Implementation” (January 29, 2025). Provides Federal Zero Trust Maturity Model progress report. CISA, “Foundations for OT Cybersecurity: Asset Inventory Guidance for Owners and Operators” (August 2025). Joint guidance with FBI, UK NCSC, and international partners. CISA, “Connected Communities Guidance: Zero Trust to Protect Interconnected Systems” (September 2024). Addresses zero trust for smart cities and critical infrastructure. Federal News Network, “Why deeper defense collaboration demands a zero trust approach to cybersecurity” (December 2025). Covers NATO 2025 Summit cybersecurity spending commitment. Cyber Security District, “Key Cybersecurity Takeaways from the 2025 NATO Summit” (July 3, 2025). Details 1.5% of 5% GDP target allocated to cyber resilience. Accessibility standards and enforcement: ETSI, “TC HF Activity Report 2023” (2023). ETSI Technical Committee on Human Factors. Covers ES 204 009 development for European Accessibility Act Total Conversation requirements. ITU-T Recommendation F.703, “Multimedia conversational services” (2000). Provides Total Conversation definition: voice, real-time text, and video synchronized in single session. ETSI EN 301 549 V3.2.1, “Accessibility requirements for ICT products and services” (March 2021). Harmonized European Standard for Web Accessibility Directive and EAA. Section508.gov, “Section 508 Standards” (revised 2017). US General Services Administration. Federal accessibility requirements integrated into Federal Acquisition Regulation. Section508.gov, “Buy Accessible Products and Services” (2025). Covers procurement requirements and Best Meets exception process. US General Services Administration, “Federal Acquisition Regulation (FAR)” (2025). Legal framework for accessibility requirements in federal procurement. Enterprise SaaS governance: Reuters, “Salesforce restrictions on Slack data access for third-party AI tools” (June 11, 2025). Financial services: Billy Hult (Tradeweb CEO), remarks via Reuters on 24/7 trading norm (December 18, 2025). Healthcare and digital health: Maneesh Goyal (COO, Mayo Clinic Platform), quoted in Digital Health News, “Mayo Clinic Launches Platform_Orchestrate to Speed Therapies to Patients” (October 2025). InterSystems and Google Cloud, HLTH 2025 announcement, covered in Signify Research, “From Ambient Notes to Data Fabric: What October Revealed About the Future of Enterprise Health IT” (November 6, 2025). Smart cities: Kate O’Neill, Smart City Expo keynote framing (November 5, 2025). Tomorrow Mobility. Industrial automation: Chingpo Lin, VP for Global IoT Automation, Advantech, quoted in IoT Analytics, “SPS 2025: How Siemens, Beckhoff, Rockwell, ABB, and peers are positioning for industrial automation’s future” (December 9, 2025). Document metadata: Title: When Independent Constraints Collapse in the Internet of Everything Subtitle: Why IoE Needs a Missing Control Layer for Continuous Identity, Policy, and Compute Author: Thomas Rocha III Date: December 2025 (Six months post-EAA enforcement) Version: 6.3 (All field signals verified against primary sources) Word count:

Approximately 20,100 words Citation style: In-line attribution with compliant sourcing Scope and positioning: This white paper analyzes cross-vertical coordination failures in IoE systems and establishes architectural requirements for SIOAR-aligned solutions. It does not prescribe a specific protocol, implementation mechanism, or vendor platform, and it does not disclose any proprietary implementation details. The intent is to provide a coherent problem statement, a scaling argument, and evaluation criteria that can be applied across domains.