COMMON OPERATIONAL DATA LAYER

A Unified Data Foundation for Defense, Security, Intelligence, and Law Enforcement

Executive Summary

Defense and security organizations don't fail because they lack data. They fail because their data doesn't reach the right people, in the right form, in time to act. Military commands, intelligence services, homeland authorities, and law enforcement are frequently tracking the same threats — but drawing on disconnected systems with incompatible data structures and no reliable mechanism for structured exchange.

The Common Operational Data Layer (**CODL**) addresses that directly. It's a secure, governed, standards-based data architecture that allows authorized organizations to publish, discover, and act on operational data in real time — without surrendering ownership of their systems or requiring a large-scale consolidation program.

CODL can be deployed within a single agency, federated across multiple authorities, or adopted as a national data standard. At every level, the outcome is the same: faster decisions, more complete operational awareness, and a defensible architecture for AI.

> *The constraint isn't sensors, bandwidth, or analytic tools. It's the architecture connecting them.*

Orbis Catalyst is the implementation framework that makes **CODL** operational. Together, they replace integration as an ongoing cost with coherence as a designed property of the enterprise.

1. The Problem That Investment Hasn't Solved

Security organizations have spent decades acquiring capable systems: surveillance platforms, biometric repositories, intelligence analysis tools, border monitoring networks, command-and-control environments. The investment is real. The problem is that each system was procured independently — different directorate, different time, different vendor — and connects to others through custom interfaces built one at a time.

Point-to-point integration multiplies fragility. An enterprise with twenty systems and forty custom links doesn't become more capable when it adds a twenty-first system — it becomes harder to maintain. Every modernization effort risks breaking existing connections. Every new acquisition creates more.

Classification rules and jurisdictional boundaries make the problem worse. Without structured mechanisms for controlled sharing, data moves informally: intelligence is manually downgraded, reports are duplicated, and the operational picture decision-makers receive reflects yesterday's information. Policy designed to protect sensitive data functions instead as a barrier to using it.

AI tools have been deployed as a partial response, and many are capable. But a model trained on one agency's data can't correlate across another's. An inference service built on inconsistent inputs produces unreliable outputs. AI doesn't fix an architecture problem — it inherits one.

2. What **CODL** Is — and What It Isn't

CODL is a secure, governed, standards-based data architecture. It defines how operational data is structured, how it's labeled, how security attributes are enforced, and how authorized parties discover and access it. It sits between existing systems and the services that depend on them.

It's not a database. It's not a platform. It's not a centralized repository. This distinction matters because prior integration efforts have often attempted consolidation — pulling data into a shared system, or imposing a single platform on a diverse enterprise. Those approaches fail predictably: agencies resist losing data ownership, vendors resist giving up integration relationships, and migration costs rarely justify the disruption.

CODL avoids that failure mode. It imposes a common standard, not a common system. Source systems — surveillance networks, case management platforms, border databases, intelligence repositories — continue to operate as they do today. At the point of integration, they map their data to published schemas, apply security labels, and publish through the **CODL** fabric. Consuming systems draw from that fabric. The interface is built once and maintained in one place, not replicated across dozens of bilateral connections.

Five properties distinguish **CODL** from conventional approaches. It's scalable from a single agency to a national backbone. It's secure by design, with cryptographic identity and attribute-based access control embedded at the data level. It's policy-governed, so what flows between whom and under what conditions is documented and enforced. It's AI-ready, giving models the structured, consistent, auditable inputs they need to perform reliably. And it's vendor-neutral — it can't be locked to a single commercial provider.

3. Deployable at Any Scale

Agency Level

A single organization — a military service, intelligence directorate, border authority, or national police force — can adopt **CODL** internally without any interagency agreement or national coordination. At the agency level, **CODL** standardizes internal integration, enforces schema compliance in new procurement, and creates the data foundation on which agency-wide AI services can operate reliably. It's a practical starting point with immediate operational return.

Federated Multi-Agency

Where two or more agencies choose to exchange data, their **CODL** implementations federate through shared identity trust, cross-domain filtering, and policy-controlled release. Each agency retains full ownership of its data and defines what it makes available to others. Interoperability comes from shared standards and explicit policy — not from transferring data ownership or ceding operational control.

National Implementation

At the national level, **CODL** becomes the operational data standard against which new systems are procured, within which civil and military authorities exchange information, and on which national AI capability is built. It publishes authoritative schemas, establishes cryptographic trust anchors, and enforces compliance through acquisition policy. It doesn’t require replacing existing agency implementations — it grows from them. National-scale coherence is reachable through a sequence of achievable steps.

Coalition and Multinational

Where allied or partner nations choose to collaborate, whether in a coalition operation, a bilateral intelligence-sharing arrangement, or a regional security framework, **CODL** brings attribute-based access controls and cryptographic identity models to allow data to flow across national boundaries without surrendering sovereignty over release decisions. Each nation maintains its own **CODL** instance, its own classification authority, and its own policy enforcement. Federation between partners is governed by explicit trust agreements mapped to technical controls: a nation sharing with a Five Eyes ally operates under one release policy; the same nation sharing with a newer bilateral partner operates under another — and both coexist within the same architecture. Data isn't downgraded or manually sanitized for export. It's tagged at origin with security attributes that downstream access controls enforce automatically, at machine speed, across every trust tier, even with Cross-Domain Solutions inserted at key boundaries. The result is that multinational sharing stops being a diplomatic workaround and becomes an architectural property. Specifically, one that scales from a two-nation partnership to a theater-wide coalition without redesigning the data layer each time.

4. How It’s Built

CODL is layered to separate governance from enforcement, and enforcement from consumption. At the base are source systems — the surveillance platforms, intelligence repositories, biometric databases, law enforcement systems, and infrastructure monitoring networks that already exist. These aren’t replaced. They connect to **CODL** through a compliance and adapter layer, which requires conformance to published interface standards, mapping to standardized schemas, and support for cryptographic enforcement. The compliance burden rests with vendors, not with the operating agency.

Above that sits the **CODL** core: schema registry, identity federation, access control enforcement, encryption orchestration, event streaming, and data lineage tracking. This is where policy is enforced and where the audit record lives. Above it, shared services — entity resolution, threat scoring, geospatial fusion, AI inference — operate on data that’s structured, consistent, and properly labeled. Mission applications consume from those services: operations centers, analysis platforms, border control systems, and law enforcement command environments, all drawing from the same governed foundation.

Resilience is designed in. Edge nodes operate autonomously under degraded connectivity, synchronize securely when connectivity is restored, and maintain compartment isolation throughout. Tamper detection and controlled rejoin protocols ensure that no disconnected or compromised node can corrupt the integrity of the broader enterprise on reconnection.

5. **Orbis Catalyst**: From Standard to Running System

CODL defines what the architecture must achieve. **Orbis Catalyst** is the implementation framework that delivers it — the deployment architecture, cryptographic trust model, hardware security integration, PKI governance, compliance testing environment, and edge-node deployment model that translate the standard into a functioning system.

Without **Orbis Catalyst**, **CODL** is a specification with no delivery path. With **Orbis Orbis Catalyst**, an agency has both the standard and the means to meet it. The two are designed to work together: **CODL** sets the requirements, **Orbis Catalyst** operationalizes them at agency or national scale.

6. The Strategic Case

The argument for **CODL** isn't primarily technical. It's about decision speed. Organizations that move faster from information to action than their adversaries hold a structural advantage — and that speed is a function of how well their data is governed, not how powerful their individual systems are. A fragmented, point-to-point integration architecture can't sustain the tempo that modern security environments require, regardless of what sits at either end of the connection.

CODL reduces integration cost by replacing dozens of custom interfaces with a single governed standard. It reduces vendor dependency by eliminating the lock-in that comes from proprietary integration relationships. It enables structured collaboration between agencies that currently share data informally or not at all. It creates the conditions under which AI can be deployed at scale, with consistent inputs, auditable lineage, and controlled model deployment.

These outcomes aren't theoretical. They follow directly from replacing architectural disorder with architectural discipline. Every organization that has standardized its data layer has reduced its integration maintenance burden, improved its operational visibility, and created a more defensible position for future procurement.

Conclusion

CODL isn't a technology acquisition. It's a decision about how a security enterprise governs its most operationally significant asset: information.

It can start within a single agency, expand across federated authorities, and grow into a national operational backbone — each step building on the last, each step delivering measurable value before the next begins. It doesn't require a large-scale consolidation program, a forced migration, or the displacement of existing systems. It requires adopting a standard and holding to it.

Orbis Catalyst makes it deployable. **CODL** makes it durable. Together, they give organizations the architectural foundation to act on what they know — when it matters.

The organization that governs its operational data governs its decision speed. In security environments, decision speed is the advantage.