SOLUTION BRIEF

The Last Exposed Layer

Dispersive® Stealth Networking for OSINT Operations

Invisible by Design. Safe by Architecture.

Executive Summary

Open source intelligence has emerged as the intelligence community's most consequential collection discipline — the foundation upon which all other intelligence functions increasingly depend. OSINT drives early warning, informs targeting, shapes policy, and fills gaps that classified collection cannot reach at speed or scale. Yet as the discipline has matured, a critical asymmetry has developed. Organizations have invested heavily in protecting the analyst — hardened browsers, managed attribution platforms, isolated workstations — while leaving the network transport between analyst and target almost entirely exposed. The wire itself remains the largest unaddressed vulnerability in the modern OSINT enterprise.

Dispersive® Stealth Networking eliminates that vulnerability through a fundamentally different architecture. Unlike conventional VPNs or tunnel-based approaches, Dispersive fragments data across multiple simultaneous network paths, encrypts each fragment independently, and routes them through constantly shifting channels. There is no tunnel to fingerprint. No source IP to log. No shared infrastructure to compromise. The technology originated in U.S. Department of Defense programs — engineered from inception for environments where detection equals mission failure and attribution means lives at risk.

This brief makes a focused argument: the managed attribution stack that protects OSINT operations today is incomplete. Browser-layer protections address one attack surface while leaving the network layer — the actual transport of data — visible, fingerprintable, and exploitable. Dispersive provides the missing transport layer without displacing existing investments. It sits beneath the browser and above the wire, completing the stack from endpoint to egress.

The OSINT Challenge

Open source intelligence is the practice of deriving actionable insight from publicly or commercially available information — without direct engagement with the target. It is the intelligence discipline that operates in the open, drawing from social media, public records, commercial databases, forums, and the indexed and unindexed web. The IC OSINT Strategy 2024–2026 designates it "the INT of first resort," recognizing that OSINT increasingly provides the initial thread that other collection disciplines then pull. The workflow follows a disciplined arc: identify and collect relevant open source material, interpret that material through analytical tradecraft, and act on the resulting intelligence to inform decisions or enable operations.

But the same openness that makes OSINT valuable also makes it dangerous. The internet is not a passive environment. Adversaries — nation-states, criminal organizations, violent extremist networks — have become highly sophisticated at detecting, tracking, and retaliating against collection activity. Every connection an analyst makes to a target environment is a two-way street. The analyst observes the target; the target's infrastructure

observes the analyst. This reality has transformed OSINT from a low-risk discipline into one where operational security failures carry severe consequences.

Digital Fingerprinting

Modern adversaries do not rely on simple IP blocklists. They deploy machine-learning-driven behavioral analysis that correlates metadata signatures across sessions and platforms — browser configuration, request cadence, navigation patterns, TLS handshake characteristics, and transport-layer timing. Even when an analyst changes IP addresses between sessions, these composite fingerprints can identify repeat visitors with high confidence. The adversary's detection capability has become an arms race: each new obfuscation technique provokes a more sophisticated classifier. Protecting the browser alone is insufficient when the network transport itself produces a consistent, analyzable signature.

Follow Me Home Attacks

When an analyst's identity or organizational affiliation is exposed, the consequences cascade well beyond the compromised operation. Hostile actors routinely instrument forums, paste sites, and infrastructure with logging mechanisms designed to trace visitors back to originating networks. A single exposed connection can reveal agency identity, compromise ongoing investigations, and place individual analysts at personal risk. In adversary environments where the target has the capability and motivation to retaliate — nation-state cyber units, transnational criminal networks, terrorist organizations — follow-me-home exposure is not a theoretical risk. It is an operational certainty that has been documented across multiple domains.

Harvest Now, Decrypt Later

Nation-state intelligence services are capturing encrypted network traffic at scale with the explicit strategy of decrypting it when quantum computing capability matures. The timeline is not abstract: NIST finalized its first set of post-quantum cryptography standards in 2024, and DHS has directed component agencies to begin transitioning now. Any VPN session captured today — regardless of its current encryption strength — is a future intelligence asset for the adversary. For OSINT operations, this means that encrypted traffic captured during collection against a foreign target today could reveal analyst identity, organizational affiliation, and collection priorities years from now.

Shared Pool Compromise

Most managed attribution platforms operate a shared pool of egress IP addresses. Analysts from multiple organizations — and sometimes multiple agencies — exit through the same set of endpoints. This architecture introduces a systemic, single-point-of-failure risk: when one analyst's activity causes a pool address to be flagged, blocklisted, or actively monitored by an adversary, every past and future analyst who used or will use that address inherits the compromise. The blast radius is not limited to one operation or one analyst. It extends across every mission that transited the burned infrastructure.

The Gap in Legacy Tools

Why Protecting the Browser Is Not Enough

The managed attribution industry has concentrated its engineering effort almost exclusively on the application layer. Browser fingerprint randomization, operating system signature spoofing, language and timezone manipulation, canvas and WebGL obfuscation — these are necessary and often sophisticated capabilities. They address a real attack surface. But they address only half the problem. Below the browser, at the network transport layer, data still travels through identifiable tunnels, exits through logged IP addresses, and produces traffic patterns that modern classifiers can fingerprint with high accuracy.

Traditional VPNs create recognizable tunnel signatures that machine-learning traffic classifiers can identify regardless of payload encryption. Research has demonstrated classification accuracy exceeding 95% for common VPN protocols. A single source IP is visible at every session endpoint, providing adversaries a stable identifier to correlate across time. And when a vendor operates a shared pool of egress addresses, the compromise surface is collective — one burned address exposes every analyst who used it. The browser may be invisible. The tunnel is not.

The following table compares the network-layer capabilities of traditional VPN and browser-isolation managed attribution platforms against Dispersive's stealth networking architecture. The comparison isolates the transport layer — the domain where legacy tools leave the most significant gaps.

Capability	Traditional VPN / Browser-Isolation Managed Attribution (MA)	Dispersive® Stealth Networking
Traffic Appearance	Recognizable TLS/VPN signature, flagged by Machine Learning (ML) classifiers	Indistinguishable from normal traffic. No tunnel to fingerprint.
Source IP Exposure	Single source IP visible at every session	No single IP. Fragments arrive across independent paths.
Shared Pool Risk	One burned pool burns every analyst on it	No pool exists. No shared burn surface.
Quantum Resilience	Captured now, decrypted later	Aligned with NIST PQC guidance; resilient against harvest-now-decrypt-later threats.
Resilience Under Attack	Single tunnel fails as one unit. Session lost.	Self-healing multipath. Sessions persist through attack.
Attribution Control	Limited, Point-of-Presence (POP)-based, finite vendor pool	Policy-controlled. 40+ global egress regions.

The implications for decision-makers are direct. Every row in this table represents a category of risk that legacy managed attribution tools leave unmitigated at the transport layer. Dispersive does not replace application-layer protections — it provides the foundational network layer that those protections were never designed to address. The combination of both layers produces a complete attribution stack. Either layer alone leaves exploitable gaps.

Dispersive Stealth Networking

The Missing Layer in the Attribution Stack

Dispersive is not a replacement for existing managed attribution platforms. It is the foundational transport layer that those platforms were never designed to provide. The architecture is best understood as a stack: at the top, workspace and browser-layer managed attribution tools control the digital persona — fingerprint, language, time zone, behavioral signature. In the middle, Dispersive controls the network transport — how data moves, through which paths, with what visibility. Below both layers sits the public internet. Each layer serves a distinct and non-overlapping function. Removing any layer leaves the analyst partially exposed.

This positioning is deliberate. Organizations have already invested in application-layer managed attribution. Those investments deliver real value and should be preserved. What Dispersive provides is the layer those tools assumed someone else was handling — and no one was. By completing the stack rather than competing with it, Dispersive eliminates the need for rip-and-replace decisions and accelerates time to operational capability.

Dispersive is workspace-agnostic, operating transparently beneath leading Secure Enterprise Browsers (SEB), including Seraphic (CrowdStrike), Kasm, and Tehama, without requiring changes to analyst tools, procedures, or training. Organizations with existing SEB investments preserve those investments while gaining the transport-layer protection those platforms were never designed to provide."

Built for High-Consequence Environments

Dispersive architecture was developed under U.S. Department of Defense programs for environments where the cost of detection is measured in mission failure and human safety. Three core properties define its operational character.

Invisible

Traffic protected by Dispersive cannot be identified as intelligence activity. There is no tunnel signature for classifiers to flag, no consistent source address for adversaries to log, no traffic pattern that distinguishes collection activity from routine web browsing. Invisibility is not a feature layered on top of a conventional architecture — it is the architecture itself. For decision-makers, this means operations can be conducted against adversary-monitored infrastructure without generating the network-layer signatures that trigger detection.

Resilient

Operations continue under active disruption without session loss. Because data transits multiple simultaneous paths, the failure or compromise of any individual path does not interrupt the session. The system autonomously reroutes around degraded or blocked paths in real time. For decision-makers, this means

collection operations persist through network attacks, infrastructure outages, and adversary interference — conditions that would terminate a single-tunnel VPN session immediately.

Sovereign

The organization controls the transport plane with no external dependencies. There is no third-party IP pool to trust. No vendor infrastructure that must remain uncompromised. No shared egress points where one customer's exposure becomes another's liability. For decision-makers, this means operational security is determined by internal policy, not by the security posture of a managed attribution vendor's shared infrastructure.

How It Works

Architecture for Invisible, Resilient Communications

When data leaves an analyst's device, Dispersive intercepts it before it reaches the network. The payload is first fragmented at the origin — broken into discrete segments that are individually meaningless. No single fragment contains enough information to reconstruct the communication's content, source, or destination. This fragmentation occurs at the application layer, before any network transmission begins.

Each fragment is then independently encrypted using multi-layered, NSA-grade encryption aligned with NIST post-quantum cryptography guidance. Critically, each fragment receives its own encryption envelope — there is no shared session key across fragments, and no single decryption event can expose the complete payload. This eliminates the harvest-now-decrypt-later threat at the architectural level: even a future quantum-capable adversary who decrypts one captured fragment obtains only an unintelligible data shard.

The encrypted fragments are then dispersed across multiple independent network paths simultaneously. These paths are selected dynamically and change continuously — different carriers, different routes, different geographies. No two fragments follow the same path. No single network observer, ISP, or nation-state monitoring apparatus ever sees more than one fragment of any given communication. The traffic on each individual path appears as routine, unremarkable internet activity. There is no tunnel to observe. No VPN handshake to fingerprint. No consistent flow to correlate.

At the destination, the fragments are reassembled into the original payload. The reassembly process validates integrity and completeness before delivering data to the application layer. The entire sequence — fragmentation, encryption, dispersal, transport, reassembly — is transparent to the analyst and to the applications running above it. The critical takeaway: no single observer, at any point in the network path, ever has access to complete data. The attack surface is not defended. It is eliminated.

Managed Attribution and Stealth Transport

Controlling What the Adversary Sees

Managed attribution is fundamentally about control — controlling the apparent identity, origin, and organizational affiliation of communications so that the adversary's picture of who is collecting against them remains distorted or blank. Browser-layer tools control the digital persona. Dispersive controls the network identity. Together, they present the adversary with a fabricated surface at every layer of the stack — from the HTTP headers down to the packet routing.

Tunnel-Free Traffic Appearance

Conventional managed attribution platforms route traffic through VPN tunnels that produce identifiable protocol signatures. Dispersive eliminates this exposure entirely. Because traffic is fragmented and dispersed rather than tunneled, there is no VPN handshake, no tunnel encapsulation header, and no consistent encrypted stream for classifiers to flag. To any network observer — including nation-state deep packet inspection systems — Dispersive-protected traffic is indistinguishable from ordinary internet activity. This is not obfuscation layered on top of a tunnel. It is the absence of a tunnel.

Global Egress Control Across 40+ Regions

Dispersive provides policy-controlled egress across more than 40 global regions, enabling analysts to appear to originate from virtually any geography required by the mission. Unlike managed attribution vendors that offer a fixed and finite set of egress points, Dispersive's egress architecture is elastic and policy-driven. Analysts or mission planners select the desired apparent origin through centralized policy — no manual reconfiguration required. This enables rapid pivoting between geographic personas without operational downtime.

Source IP Concealment

In a traditional VPN architecture, a single source IP address is visible at the session endpoint — providing adversaries a stable, correlatable identifier. Dispersive eliminates this by dispersing fragments across independent paths with independent addressing. There is no single IP address for the adversary to log, blacklist, or use as a starting point for attribution. The analyst's true network origin is architecturally hidden, not merely obscured.

Transparent and Non-Transparent Operating Modes

Dispersive supports both transparent mode — where the stealth transport layer operates invisibly beneath existing applications and browser-based MA platforms — and non-transparent mode, where the organization exercises direct control over traffic routing and egress selection. This flexibility allows Dispersive to integrate with any existing managed attribution workflow, including leading Secure Enterprise Browsers such as Seraphic (CrowdStrike), Kasm, and Tehama, without requiring changes to analyst tools, procedures, or training.

Managed attribution controls what the target sees in the browser.
Dispersive controls what the adversary sees on the wire.

Elimination of Shared Pool Risk

Because Dispersive does not operate a shared pool of egress IP addresses, there is no pool to burn. Each organization's traffic is sovereign — routed through infrastructure controlled by policy, not shared with other customers. When one analyst's activity triggers adversary detection, the blast radius is contained to that operation. There is no systemic contagion to other analysts, other missions, or other organizations. This architectural isolation eliminates the single largest systemic risk in pooled managed attribution.

Outbound-Only Gateway Configuration

Dispersive gateways can be configured in outbound-only mode, meaning they initiate connections but do not accept inbound traffic. This eliminates an entire class of attack vectors — port scanning, service enumeration, inbound exploitation — that adversaries use to probe and compromise egress infrastructure. The gateway is invisible from the outside. It cannot be discovered, probed, or targeted because it presents no listening surface to the internet.

Operational Use Cases

Mission Scenarios Where the Network Layer Decides the Outcome

1. Persistent Collection Against a Denied-Area Forum

An intelligence team conducts sustained collection against a foreign-language forum operated by a hostile non-state actor. The forum administrator maintains sophisticated blocklists, tracks visitor metadata across sessions, and feeds fabricated content to identified collectors. Legacy managed attribution tools rotate browser fingerprints and egress IPs — but the VPN tunnel signature remains constant, and the administrator's Machine Learning (ML) classifier flags repeat visitors based on transport-layer patterns. Collection degrades as the team burns through the vendor's IP pool.

With Dispersive, the transport layer produces no consistent signature. Each session arrives via different paths with different network characteristics. There is no tunnel to fingerprint and no IP pool to exhaust. The forum administrator's classifiers find nothing to correlate. Persistent collection continues without detection, deception, or pool degradation — because the network layer never presents a stable surface to analyze.

2. In-Region Collection Where the Tunnel Is the Signal

An analyst needs to collect against open source targets within a nation-state that operates pervasive network surveillance. The country's monitoring apparatus does not need to decrypt VPN traffic to act on it — the mere presence of an encrypted tunnel from a foreign IP range triggers flagging, deeper inspection, and potential blocking. Using a conventional VPN in this environment does not protect the analyst; it identifies them. The tunnel itself is the signal.

Dispersive eliminates this exposure. Traffic exits through in-region infrastructure and appears as routine local internet activity. There is no tunnel for surveillance systems to detect. No foreign IP address to flag. No encrypted stream anomaly to trigger automated response. The analyst operates within the adversary's network environment without producing the transport-layer artifacts that surveillance systems are tuned to detect.

3. Tactical Edge Operations in Contested Environments

A forward-deployed team requires OSINT collection capability in an environment with degraded communications infrastructure — intermittent satellite connectivity, contested cellular networks, and active electronic warfare. A single-path VPN connection fails repeatedly as transport links degrade or are jammed. Each reconnection cycle exposes the team to fingerprinting, interrupts collection sessions, and creates detectable network events.

Dispersive's multipath architecture distributes traffic across every available transport — satellite, cellular, tactical radio, local wireless — simultaneously. When one path degrades or is denied, traffic autonomously shifts to surviving paths without session interruption. The team maintains persistent collection capability across a degraded, multi-transport environment without manual intervention, reconnection events, or the detectable network signatures that single-tunnel failovers produce.

4. Law Enforcement Research into Criminal Networks

Federal analysts research hostile criminal infrastructure — dark web marketplaces, encrypted communication platforms, fraud networks — to support active investigations. The operational requirement is absolute: the agency's identity and investigative interest must remain completely hidden. A single network-layer attribution event could compromise the investigation, alert subjects, and trigger evidence destruction. The stakes extend beyond operational security to prosecutorial integrity and officer safety.

Dispersive ensures that no network artifact connects the analyst's activity to the agency. Fragmented, independently encrypted traffic traverses multiple paths with no common origin. The criminal infrastructure's logging and counter-surveillance mechanisms capture fragments from different network paths — none of which can be correlated to a single source, a single agency, or a single investigation. The network layer is not merely obscured; it is architecturally decoupled from the analyst's identity.

Deployment Models

Flexible Architecture — From Managed Cloud to Air-Gapped Sovereignty

Dispersive provides three deployment models designed to match the operational, security, and sovereignty requirements of different organizational environments. Each model delivers the same core stealth networking capability — multipath fragmentation, independent encryption, tunnel-free transport — through a different operational and infrastructure model.

DispersiveCloud™ — Managed SaaS

DispersiveCloud is the fastest path to operational capability. Delivered as a fully managed SaaS platform, it provides stealth networking with no infrastructure deployment required by the customer. The platform is SOC 2 Type II certified, managed via REST API, and maintained by Dispersive's operations team. DispersiveCloud is designed for organizations that need rapid deployment, minimal operational overhead, and the ability to scale collection capacity without capital infrastructure investment. It is the right choice for teams that want to add stealth transport to their existing managed attribution stack without building or managing network infrastructure.

DispersiveFabric™ — Self-Managed / On-Premises

DispersiveFabric provides the same stealth networking architecture in a self-managed model that can be deployed on-premises, in private cloud environments, or in air-gapped facilities. The organization owns and controls the entire transport plane — no Dispersive-operated infrastructure is involved. DispersiveFabric is designed for intelligence community organizations, defense agencies, and enterprises with sovereignty requirements that preclude any external cloud dependency. It delivers maximum operational control for environments where data residency, infrastructure sovereignty, and air-gap compliance are non-negotiable.

Dispersive Client — Endpoint Integration

The Dispersive Client operates at the endpoint — installed on the analyst's device to intercept and protect traffic before it reaches the network. It supports cross-platform deployment across Windows, macOS, Linux, iOS, and Android. The client integrates transparently with existing applications and managed attribution platforms, requiring no changes to analyst workflows. It is designed for organizations that need per-device stealth transport capability, whether connecting to DispersiveCloud, DispersiveFabric, or a hybrid of both.

Strategic Alignment

Meeting the Intelligence Community Where It's Headed

The Intelligence Community OSINT Strategy 2024–2026, released jointly by ODNI and CIA in March 2024, establishes OSINT as "the INT of first resort" and directs the IC to professionalize, modernize, and scale the discipline. The strategy is not aspirational guidance — it is a directive with budgetary implications and implementation timelines. Dispersive's architecture aligns directly with three of the strategy's core focus areas.

Tradecraft Modernization

The IC OSINT Strategy calls for developing "next-generation OSINT workforce and tradecraft" — recognizing that collection methods must evolve as the threat environment changes. Network-layer stealth is the logical next evolution of managed attribution tradecraft. Browser-layer obfuscation was the first generation. Transport-layer invisibility is the second. Dispersive provides the technology foundation for this tradecraft evolution, enabling analysts to operate against sophisticated adversaries who have already developed the capability to defeat first-generation managed attribution through transport-layer analysis. Adopting Dispersive is not adding a tool — it is advancing the tradecraft baseline.

Innovation on Unclassified Systems

The strategy emphasizes driving "OSINT innovation to deliver new capabilities" — with particular focus on unclassified systems where the majority of OSINT collection occurs. Dispersive operates at the unclassified level, integrating with commercial internet infrastructure and existing analyst workstations. It does not require classified network access, special facilities, or dedicated infrastructure. This positions it squarely within the strategy's emphasis on delivering capability improvements to the unclassified environments where OSINT analysts actually work — not within classified enclaves where OSINT is consumed but not collected.

External Partnerships

The IC OSINT Strategy explicitly calls for "partnering with brilliant American innovators" to accelerate capability development. Dispersive is a U.S. company with technology lineage from DoD programs, operational deployments across defense and intelligence, and SOC 2 Type II certification. It represents precisely the type of innovative American technology partner the strategy envisions — a commercial entity delivering defense-grade capability through commercial mechanisms, with the operational validation and security posture required for IC adoption.

Proof Points

Validated Where It Matters Most

In high-consequence environments, vendor claims are worthless without operational validation. Defense and intelligence organizations do not adopt technology based on marketing collateral or benchmark slides. They adopt technology that has been tested under operational conditions, deployed in mission environments, and proven against real-world threat scenarios. Dispersive's credibility rests not on what it promises, but on where it has already delivered.

U.S. Government and Defense

Dispersive has been deployed within U.S. Government and defense environments, including the Department of Homeland Security, to provide stealth transport for sensitive communications across contested and high-threat networks. These deployments validated Dispersive's ability to deliver invisible, resilient connectivity in environments where conventional VPN and SD-WAN solutions produced detectable signatures or failed under adversary interference. The operational requirement was unambiguous: communications must not be identifiable as government traffic, and sessions must persist through active network disruption.

Intelligence Community and DoD Programs

Dispersive's core technology architecture was originally developed under Department of Defense programs designed for mission-critical defense communications. This lineage is not incidental — it means the architecture was engineered from first principles for environments where traffic detection results in mission compromise and where network resilience must be maintained under active electronic warfare conditions. Subsequent deployments within intelligence community programs have validated the technology against the specific operational requirements of classified and sensitive collection missions.

American Tower — Global Edge Operations

In a multi-month pilot with American Tower and Vertical Data at the American Tower Edge Data Center in Raleigh, North Carolina, Dispersive demonstrated measurable performance and security gains for high-throughput workloads at the network edge. Benchmark testing against conventional secure networking approaches showed up to 30% throughput improvement, sub-second autonomous failover, and enhanced traffic segmentation. This deployment validated Dispersive's capability beyond government environments — demonstrating that stealth networking delivers performance and resilience advantages for commercial critical infrastructure at global scale.

Getting Started

From Assessment to Operational Scale

Dispersive engages through a structured, three-step process designed to move from initial assessment to operational deployment with minimal organizational friction. Each step is scoped to deliver concrete value and inform the next decision point.

Step 1: Mission and Attribution Assessment

Dispersive conducts a focused assessment of the organization's current OSINT collection architecture, managed attribution stack, and operational security requirements. The assessment identifies specific transport-layer gaps, quantifies risk exposure from current approaches, and maps Dispersive's capabilities against mission requirements. This step typically completes in two to three weeks and produces a written assessment with specific architectural recommendations.

Step 2: Technical Proof of Value

Based on assessment findings, Dispersive deploys a scoped proof-of-value engagement that demonstrates stealth transport capability against the organization's actual operational environment and threat profile. The proof of value validates traffic invisibility, resilience under disruption, managed attribution integration, and egress control across target geographies. Organizations see Dispersive operating in their environment, against their use cases, before making a deployment decision. This step typically runs four to six weeks.

Step 3: Deploy and Scale

Following successful proof of value, Dispersive deploys the selected architecture — DispersiveCloud, DispersiveFabric, or a hybrid model — and scales to operational coverage. Deployment includes integration with existing managed attribution platforms, policy configuration for egress control and geographic attribution, endpoint client deployment, and operational training. Dispersive provides ongoing operational support and architecture optimization as the organization's OSINT mission evolves and expands.

To initiate the process, contact Dispersive at https://dispersive.io/company/contact-us.

Others assume the network can be trusted.
Dispersive® Stealth Networking operates when it's not.

Dispersive Holdings, Inc.
300 Colonial Center Pkwy Suite 100 | Roswell, GA 30076 USA
dispersive.io | 1.844.403.5850

dispersive