Darwinium

SOLUTION BRIEF — SECURITY & CISO

The cyberfraud platform built for security teams

The traditional bot playbook is breaking down.

Security teams continue to face familiar threats, including credential stuffing, account takeover, automated abuse, API attacks, and scraping, at relentless scale. But AI agents are changing the nature of the challenge.

Legitimate agents now act on behalf of real customers, while malicious automation increasingly blends in with trusted traffic. As a result, the traditional "human or bot?" test is becoming less useful.

What matters now is intent: who or what is acting, on whose behalf, and whether that activity introduces risk.

Darwinium closes this visibility gap at the edge. Its security and fraud prevention platform helps security and fraud teams work from the same intelligence layer. A single edge deployment provides visibility across human and agentic interactions while supporting fraud prevention use cases—without requiring changes to application code or introducing additional infrastructure for your team to operate.

The visibility gap your current stack cannot close

Logs, WAFs, and traditional bot-management tools were designed for a web dominated by humans and relatively simple automation.

- A legitimate customer
- An undeclared or compromised agent

That gap is widening. Much of today's AI-driven traffic is undeclared and arrives without reliable identification. At the same time, Model Context Protocol (MCP) servers are emerging as a new interaction surface that browser- and network-layer tools may not see in full journey context.

They can tell you that a request arrived. They often cannot tell you whether the actor behind it is:

- A trusted AI agent acting on that customer's behalf
- Malicious automation operating with stolen or hijacked intent

Authentication helps establish identity. Intent helps explain behavior. Darwinium continuously evaluates behavioral signals, device intelligence, network context, and journey activity at the edge. This makes it possible to distinguish trusted activity from compromised or malicious behavior, not simply identify that an agent or automated system is present.

What you can see from the edge

Every step in one view

Gain continuous journey visibility across web, mobile, API, and agentic traffic; beginning with the first pre-authentication interaction, rather than only at login or checkout.

Human and agentic activity assessed by intent

Evaluate declared and undeclared agents according to what they are doing, how they are behaving, and the context of the journey, not through a simple "bot or human" verdict.

Greater visibility into MCP traffic

Bring MCP activity into the wider customer journey, helping security teams understand interactions that point-in-time API or browser tools may assess only in isolation.

A replayable decision trail

Preserve the evidence behind each risk decision and stream relevant outcomes to your SIEM to support investigations, audit, and incident response.

Threats identified before they reach the application

Surface account takeover, credential stuffing, automated abuse, and suspicious API activity at the perimeter, reducing the volume of malicious traffic reaching your origin infrastructure.

Signals traditional logs do not capture

Combine edge and API visibility with device, behavioral, identity, and network context that may not be available through server-side logging alone.

One deployment. Three teams.

The same edge deployment that provides security teams with greater visibility can support fraud use cases and reduce operational demands on platform teams.

SECURITY AND CISO TEAMS: SEE MORE

Gain edge and API visibility beyond traditional application logs • Continuously assess human and agentic traffic as journeys develop • Detect business-logic abuse using the full context of user interactions • Identify and prevent account takeover throughout the journey, not only at login • Detect sophisticated automation, social engineering, and coercion indicators • Analyze API content and context, rather than relying only on request volume • Support privacy and data-residency requirements through edge encryption and customer-controlled storage options • Stream explainable, replayable decision evidence to your SIEM

FRAUD AND RISK TEAMS: CATCH MORE

Use continuous signals across UI, API, mobile, and agentic journeys • Reduce losses from account takeover, scams, bonus abuse, mule activity, and bot-driven fraud • Make faster, more defensible intervention decisions • Establish a shared trust layer across brands and channels • Create and update rules, features, and models without relying on application release cycles

PLATFORM AND DEVOPS TEAMS: OPERATE MORE EFFICIENTLY

Avoid introducing another always-on application stack to run, scale, and patch • Deploy through infrastructure-as-code and controlled release processes • Reduce or eliminate changes to application code • Support bot and fraud controls through a common edge layer • Deflect unwanted traffic before it reaches origin infrastructure • Suspend, roll back, or undeploy through a controlled deployment process

Designed for your security posture

Keep sensitive data under your control

Sensitive personally identifiable information (PII) is encrypted at the edge before it leaves your environment. Darwinium can use anonymized data and hashed identifiers for matching and analysis, while access to decrypted information remains controlled through customer-managed permissions and keys.

You can also choose to store encrypted event data in your own AWS S3 environment to support privacy, security, and data-residency requirements.

Reduce third-party JavaScript supply-chain exposure

Traditional profiling approaches may load JavaScript directly from a third-party domain, creating additional supply-chain and change-control considerations.

With an edge deployment, Darwinium profiling JavaScript can be served through your own CDN and domain. Updates can follow your established review, approval, and deployment controls before being promoted into production.

Deploy through the CDN you already use

Darwinium supports edge deployment through Cloudflare, Amazon CloudFront, and Akamai. Deployment is primarily configuration-led and can be managed through Darwinium's deployment manager or incorporated into your existing CI/CD and infrastructure-as-code processes.

The exact deployment effort depends on your CDN environment, internal security review, traffic architecture, and chosen implementation.

Maintain availability through controlled fail-open behavior

Darwinium can be configured so that temporary unavailability of a supporting service does not unnecessarily interrupt legitimate customer traffic. Risk logic deployed at the edge can continue to evaluate available signals and apply configured policy.

Where a required dependency is unavailable, requests can be allowed to proceed according to your chosen fail-open controls rather than introducing an unintended availability risk. Fail-open behavior does not mean losing control. The trade-off is deliberate, configurable, and visible.

Edge-based decisioning

A lightweight Darwinium worker runs within your CDN environment and evaluates real-time signals and policy logic close to the interaction. This reduces the need to make every transaction dependent on a round-trip to Darwinium's cloud platform.

Continuous protection

Darwinium evaluates risk throughout the customer journey, allowing controls to respond as intent and behavior change, not only at a single authentication or transaction event. Depending on policy, suspicious sessions can be challenged, redirected, monitored, or terminated before they reach sensitive application functions.

Logging continuity

Event data can be written to infrastructure you control, including a customer-managed AWS S3 bucket, supporting operational resilience and data-governance requirements.

Reversible by design

Darwinium edge deployments are version-controlled and can be suspended, rolled back, or undeployed through a managed release process. Undeployment removes the relevant Darwinium workers or functions while preserving the customer's underlying CDN behavior.

Certifications and compliance

Darwinium maintains a security and compliance program designed to support enterprise requirements, including SOC 2 Type II, ISO 27001:2022, and PCI DSS obligations. Current certification reports and compliance documentation are available to qualified customers as part of the security-review process.

SOC 2 Type II | ISO 27001:2022 | PCI DSS

Dispelling the Myths About CDN Deployment

Myth: Deploying Darwinium at the CDN means losing control or visibility

False.

You retain control over where and how Darwinium operates. You decide:

- Which journeys and endpoints are monitored
- What information is collected
- When profiling and decisioning occur
- How risk outcomes are used
- How deployment changes are reviewed and released

Data, analytics, decision logic, models, and investigation tools remain visible through the Darwinium Portal. Customers can also use their own deployment and change-control processes where required.

Myth: Deploying Darwinium through a CDN is complex

False.

Most Darwinium CDN deployments can be completed in hours rather than weeks.

Deployment is primarily configuration-led, reducing the need for application code changes or dedicated development resources. Darwinium can also integrate with existing CI/CD and infrastructure-as-code processes.

Once deployed, journey and risk-logic updates can be built, reviewed, and released through a controlled deployment process—helping teams adapt protection without waiting for lengthy application release cycles.

Myth: Darwinium will disrupt or slow existing web traffic

False.

Darwinium is designed to operate at the edge with minimal impact on the customer experience.

Profiling and risk evaluation can be applied only to the pages, endpoints, and interactions you choose to protect, such as account creation, login, password recovery, or payment.

Risk processing takes place alongside the customer journey, allowing real-time decisions without introducing unnecessary friction across the rest of the website or application.

Myth: A CDN deployment creates privacy concerns

False.

Darwinium is designed with privacy and data protection at its core.

Sensitive personally identifiable information is encrypted at the edge before it leaves the customer environment. Darwinium can use anonymized data and hashed identifiers for analysis and matching, while access to decrypted information remains customer-controlled.

Customers can also store encrypted event data in their own AWS S3 environment to support privacy and data-residency requirements.

Myth: A CDN deployment will be difficult to remove

False.

Darwinium can be suspended, rolled back, or undeployed through a controlled deployment process.

Edge deployments are versioned, providing a clear path to revert changes. Undeployment removes the relevant Darwinium workers or functions while preserving the customer's existing CDN behavior.

Where security and fraud finally share one view

Because Darwinium operates at the edge across key digital touchpoints, security and fraud teams can draw intelligence from a shared source.

Security gains device, network, behavioral, and journey context that traditional controls may not provide. Fraud teams gain continuous visibility beyond point-in-time risk checks.

The result is one deployment and one decision layer for understanding human and agentic risk across the customer journey.

A detailed technical deployment guide is available for engineering, platform, and security teams.