

Cisco Secure DDoS Edge Protection: Technical White Paper





Contents

Introduction and the Evolving Threat Landscape 3

Cisco Secure DDoS Edge Protection: An architecture designed for modern, evasive botnets.....5

Innovative features and technologies that make the difference 8

Deployment Use Cases 12

Total Cost of Ownership (TCO) Benefits 14

Summary and conclusion 15

Introduction and the Evolving Threat Landscape

The global Distributed Denial of Service (DDoS) defense paradigm is undergoing a profound architectural shift. For over a decade, the industry relied heavily on centralized scrubbing centers—massive, geographically dispersed facilities that ingest attack traffic after redirecting (“tromboning”) it away from the target network. This model was effective when attacks were predictable in scale, predominantly volumetric, and originated from a limited number of botnet groups.

That era is over

Today’s threat environment is shaped by hyper-distributed botnets, highly dynamic attack vectors, and a tactical shift toward micro-bursts and evasive patterns. Centralized mitigation has become technically insufficient and economically unsustainable, primarily because:

- Traffic patterns now overwhelm redirection thresholds before scrubbing can start.
- Attackers increasingly target access, aggregation, and outbound paths—not only the edge.
- Botnet agility outpaces the response times of out-of-band systems.

These realities have pushed operators toward on-box, inline, hyper-distributed mitigation, where routers themselves act as the enforcement layer and detection occurs as close as possible to the origin of the traffic.

The New DDoS Reality

The rise of next-generation IoT botnets—most notably [Alsuru](#), [ShadowV2](#) and [Kimwolf](#)—has reshaped the global threat model since late 2024. These botnets are vastly larger, more diverse, and significantly more automated than earlier botnet groups. Their attack patterns have evolved to deliberately undermine centralized defenses, forcing operators to rethink where and how mitigation must occur.

Alsuru: Precision at Terabit Scale

Alsuru is distinguished by its massive scale, tight coordination, and its capability to shift attack vectors in seconds. Examples observed globally include:

- **15.7–29.5 Tbps Precision Floods** with coordinated bursts across tens of thousands of nodes, overwhelming links long before traffic can be redirected.
- **Distributed Carpet Bombing:** Rather than targeting a single subnet, Alsuru spreads low-intensity TCP/UDP noise across tens of thousands of destination IPs—each below traditional detection thresholds yet cumulatively devastating.
- **Outbound Weaponization:** Alsuru-infected smart TVs, cameras, and customer-premises routers launching outbound UDP reflection attacks (NTP, CLDAP, WS-Discovery), quietly consuming aggregation bandwidth and triggering upstream congestion.

ShadowV2: Stealth, adaptation and diversified payload attacks

ShadowV2 represents the opposite end of the spectrum: it favors stealth, evasion, and obfuscation. It is designed to blend into normal traffic until the moment of detonation. Notable behaviors include:

- **Carpet Bombing:** ShadowV2 rotates small flows across /28, /27, and /24 ranges, automatically shifting targets when thresholds rise.
- **Multi vector attacks:** It can alternate between UDP, TCP SYN, fragmented UDP, and GRE floods mid-attack, complicating signature-based defenses.
- **Outbound attacks:** Compromised IoT thermostats and smart meters generate thousands of tiny outbound HTTP/2 requests per second—too small to trigger conventional anomaly detectors yet capable of overwhelming remote APIs.

The Consequence: A Scale and Pattern Shift

Whether it is Alsuru's overwhelming terabit-class bursts or ShadowV2's evasive multi-vector patterns, the outcome is the same: centralized, redirect-based mitigation cannot engage quickly or locally enough to stop these attacks.

[The documented peaks](#)—now reaching 31.4 Tbps and growing—highlight the urgent need for a new architectural model: one that detects, verifies, and mitigates attacks at the router itself, where traffic first enters or exits the network.



Cisco Secure DDoS Edge Protection: An architecture designed for modern, evasive botnets

While Edge Protection covers all the legacy DDoS attacks, it was built explicitly to counter Zero-Day threats like the new attack patterns introduced by modern botnets such as Alsuru, ShadowV2 and Kimwolf. These threats demand hyper-distributed, inline mitigation that operates at wire-speed and directly on the routing infrastructure—without depending on the latency, slow activation times, or traffic redirection constraints of legacy centralized scrubbing centers.

The Edge Protection architecture addresses this by using a two-tier, containerized, controller-detector design, where all detection and enforcement occur on the router itself, providing the local reaction needed to stop micro-bursts, carpet-bombing, multi-vector floods, and outbound DDoS traffic before they propagate.




Figure 1. DDoS Edge Protection architecture

Edge Protection Controller

The Controller is the highly available, central management software for the Cisco Secure DDoS Edge Protection solution, serving as the core engine for centralized orchestration of the distributed network defense.

Architecture and Deployment

The Controller is fundamentally a Kubernetes Cluster, built using K3S (a reduced-size distribution), which enables a modular, containerized architecture.

- **Scalability and resilience:** It supports both single and multi-node deployments. Multi-node deployment enables High Availability (HA), fault tolerance, and redundancy. The Kubernetes Control Plane manages the cluster, ensuring that services are automatically managed and redistributed for maximum uptime.
- **System scope:** Typically, a single Controller is deployed to oversee all network sites, managing hundreds of distributed Agents and some 3rd party Scrubbers

Attack Lifecycle Management

The Controller functions as the central coordination layer for all detection & mitigation activities, ensuring that responses across the network are consistent, timely, and aligned with the evolving characteristics of the attack.

Detection correlation

The Controller continuously receives alerts, flow summaries, and event signals from all deployed Detectors. It correlates this information in real time to validate attack conditions, determine the scope of impact, and identify whether the threat is local, distributed, inbound, or outbound.

Mitigation policy selection and deployment

After confirming an attack, the Controller determines the most effective mitigation strategy based on the attack vectors, intensity, and direction. Depending on configuration and integrations, it can automatically deploy:

- Dynamic Access Control List (ACLs)
- FlowSpec rules
- BGP blackholing or (or Remote Triggered Black Hole filtering - RTBH)
- Third-party scrubbing center signaling (when integrated)

Policies are pushed immediately to the relevant routers or external systems, ensuring rapid and consistent enforcement.

Continuous attack monitoring and adaptation

During mitigation, the Controller tracks how the attack evolves across time, vectors, and targeted hosts. If the adversary shifts behavior—for example, changing ports, protocols, or amplification sources—the Controller can adjust the mitigation policy accordingly to maintain effectiveness.

Automated recovery

Once the attack subsides, the Controller automatically withdraws all temporary mitigation artifacts (ACLs, FlowSpec entries, blackhole routes, or external triggers). This ensures that the network quickly returns to its normal operational state without leaving behind stale or overly restrictive rules.

Software and Agent Lifecycle Management

The Controller serves as the central platform for operating, maintaining, and supervising all distributed Agents across the network.

Agent orchestration

The Controller tracks the health, status, and connectivity of every deployed agent. It automatically detects and corrects common operational issues, restarts unhealthy components, and raises alerts when persistent faults require operator attention.

Software lifecycle management

The Controller maintains a full catalogue of software images for both agents and the Controller itself. It supports coordinated, zero-touch upgrades across the entire deployment—rolling out new versions without service interruption and without impacting ongoing detection or mitigation activities.

APIs

The Controller acts as the central intelligence hub, offering comprehensive data viewing and flexible integration capabilities.

- **User interface:** Using a streamlined design avoiding complex windows and offering natural interaction with the events
- **Reporting:** It provides both **real-time and historical event reporting** alongside ongoing attack forensics and threat intelligence analysis.
- **Integration:** It offers robust integration capabilities via APIs, exporting a feature-rich **REST API** based on the **OpenAPI standard** for simple connection with other security management platforms.

Innovative features and technologies that make the difference

Intelligent detection at the Edge

Modern DDoS campaigns no longer rely on brute-force floods. Today's attacks are adaptive, low-rate, and engineered to blend seamlessly into legitimate traffic through short bursts, fragmented vectors, and wide prefix distribution. Systems built around static thresholds or centralized scrubbing often detect these threats only after damage has already begun.

Edge Protection takes a different approach: it brings intelligence directly to the network edge. Machine learning continuously adapts to real-time traffic behavior, enabling immediate, context-aware detection without waiting for anomalies to accumulate elsewhere in the network.

The Shift to Machine Learning Detection Algorithms

Traditional DDoS detection depends on fixed thresholds—attackers can easily evade these by staying just below them or distributing their traffic across many hosts. Edge Protection replaces this rigid model with true behavioral detection, focusing on patterns, relationships, and consistency rather than raw volume.

- For every protected host, the system builds a dynamic baseline that reflects its natural behavior. Instead of treating the network as a monolithic entity, Edge Protection profiles each IP individually and then groups similar profiles using the K-means clustering algorithm. These ML clusters allow the system to understand what “normal” looks like across different segments and to detect meaningful deviations with precision.
- When traffic shifts, detection follows a dual-stage process. First, the system determines whether a host or cluster has exceeded its learned thresholds. Then it validates the anomaly by examining whether the behavior aligns with legitimate patterns—such as typical inbound/outbound ratios or known time-based activity. If behavior diverges sharply from what the system has learned, the event is escalated.

This emphasis on ML integrity enables Edge Protection to uncover even zero-day attacks that bear no resemblance to previously observed vectors.

Detecting What Others Miss: Carpet-Bombing, Pulse Attacks and Outbound attacks

Carpet-bombing illustrates why behavioral detection at the edge is essential. In these attacks, traffic is spread thinly across hundreds or thousands of IPs. Individually, no IP crosses a static threshold—yet together, they create a highly disruptive event.

With per-host baselines, Edge Protection identifies these micro-anomalies immediately. When many hosts in a prefix show small but simultaneous deviations, the system correlates them and elevates the pattern into a single carpet-bombing event. Operators receive one coherent incident, and mitigation is applied precisely across the affected ranges instead of reacting to scattered low-volume alerts.

Pulse attacks—brief spikes designed to slip past averaging algorithms—are handled just as effectively. Because Edge Protection monitors instantaneous behavioral changes rather than long-term aggregates, even second-scale anomalies trigger detection when they contradict learned behavior.

Outbound DDoS attack occurs when compromised devices within a local network are weaponized to flood external targets with malicious traffic. This “inside-out” surge saturates the network’s own upload capacity and aggregation links, often triggering severe upstream congestion. Since Edge Protection inherently monitors bidirectional traffic, it easily detects these anomalous surges and can autonomously block them at the access routers before they impact the broader network.

Multi-tiered mitigation and adaptive Orchestration

Once the system confirms an attack, mitigation begins immediately. Because detection happens at the edge, mitigation can be both instantaneous and highly surgical. Rather than applying a single static policy, the system supports multiple mitigation tiers—ranging from precise on-box filtering to upstream controls—and can combine them dynamically. Central to this flexibility is its scripting engine, which allows each Protected Object (PO) to have a tailored, context-aware response that adapts in real time to the nature of the attack.

On-Box Mitigation: Precision at Line Rate

The fastest and most granular mitigation occurs directly on the router. Using native IOS-XR interfaces, the system programs Access Control Lists (ACLs) and related filtering constructs directly into the forwarding hardware.

- **Precision** – With the ability to match up to 20 different L3/L4 header fields, ACLs provide the **granularity** needed to surgically remove malicious flows without impacting legitimate traffic.
- **Advanced Matching** – Through User-Defined Fields (UDF), the system can match specific byte offsets in headers or payloads, letting operators block new or obscure zero-day vectors without waiting for vendor signatures.
- **Line-Rate Enforcement** – Because the filtering is applied on the line card, mitigation is executed in hardware—maintaining forwarding performance and keeping CPU usage unaffected.

This on-box method is ideal for targeted attacks such as SYN floods, application-specific anomalies, burst attacks, and outbound abuse originating from compromised hosts.

BGP-Based Mitigation for Large-Scale Events

When volumetric attacks threaten upstream capacity, the system escalates mitigation using BGP signaling.

- **BGP Flowspec** – The system injects granular rate limits or filtering rules upstream, containing attacks before they saturate links or impact multiple services.
- **BGP RTBH** – For catastrophic volumetric floods, the system can push a targeted blackhole route upstream to immediately eliminate traffic to a specific PO until conditions stabilize.

These upstream actions provide essential leverage when the attack volume exceeds what can be efficiently handled on-box.

Dynamic Mitigation Through the Scripting Engine

The system's most innovative capability is its **mitigation scripting framework**, which allows each PO to have a unique, adaptive policy. Instead of relying on a fixed sequence of actions, the system evaluates real-time attack telemetry and executes logic that matches the event's scale, complexity, and risk.

This scripting engine makes mitigation **conditional, contextual, and per-object specific**, enabling an almost "self-orchestrating" response. Operators can define logic that behaves differently for critical services, large prefixes, or latency-sensitive applications.

Conceptual example of script-driven logic:

- 1. Magnitude check** – If total attack bandwidth exceeds a defined threshold (e.g., 50 Gbps), automatically initiate RTBH to protect upstream links.
- 2. Vector complexity** – If the signature contains multiple parameters or resembles a multi-vector pattern, apply a precise ACL block at the edge to maximize good-traffic preservation.
- 3. Low-intensity anomaly** – Otherwise, apply a BGP Flowspec rate limit to contain the event without disrupting service.

By evaluating the attack in real time, the system ensures each PO receives the right mitigation—never too aggressive, never too passive.

Integration with External Mitigators

For organizations that want to apply hybrid defense strategies, the system provides a flexible integration framework to coordinate with external scrubbing centers or security appliances. Using a templated approach, vendor-specific differences are abstracted and mapped into the system's mitigation script actions.

Adding support for a new appliance or cloud-based solution is straightforward: a few short Python scripts are all that's needed to translate the system's script actions into the corresponding REST API calls for the external mitigator.

Link Congestion Monitoring and Forensics

A critical operational requirement is the ability to monitor the network infrastructure itself. The Edge Protection Controller offers comprehensive interface congestion monitoring.

Alerting Requirements

The user must be able to select specific interfaces (e.g., Internet Ingress or PO source interfaces) to monitor. The alert condition is defined by a threshold.

Percentage from the Link capacity

- **Hold Time:** To prevent jittering, the congestion condition must be consistent for a configurable Hold Time (default **60 seconds**) before an alert is created.
- **Release Time:** To clear the alert, the interface traffic must drop below the congestion level for a configurable Release Time (default **60 seconds**).

Forensic Data Capture

When an alert is triggered, the system automatically collects a “Forensic Snapshot.” This data provides immediate operational context, detailing the **Top 10** applications, talkers, and destinations contributing to the congestion.

MSSP-Optimized Architecture

Cisco Secure DDoS Edge Protection is built with Managed Security Service Providers in mind, offering native multi-tenancy, scalable onboarding, and flexible service-tiering to help providers deliver differentiated DDoS protection to wholesale and enterprise customers.

Massive Multi-Tenancy at Scale

The system supports onboarding more than 10,000 customers, each represented through customizable Protected Objects (POs) that define the customer’s IP ranges or AS number. Enabling a self-service Tenant Portal transforms a PO into a managed asset, the MSSPs are customers quickly onboarded while maintaining strict isolation.

Tiered Service Models

To match commercial requirements, the platform flexible detection and mitigation policies could be tailored into different service tiers (Bronze, Silver, Gold and Platinum), or fully customized profiles. Providers can define policies that govern detection sensitivity, mitigation actions, and reporting levels. Customers with portal access can adjust elements of their assigned tier while remaining within the limits set by the MSSP.

Secure Role-Based Access

A dedicated tenant portal enforces strong segmentation through role-based access control. Each tenant is associated with a single PO and can be assigned roles such as Tenant Admin, Tenant Analyst, or read-only Tenant User. A notification-only role supports customers who require alerts but no portal access.

Customer-Facing Reporting

The system includes reporting designed specifically for downstream customers. Tenants can generate detailed attack reports and period trend summaries in PDF or HTML, while scheduled daily digest emails provide ongoing visibility. All reports can be customized with provider-branded templates.

A Platform for Service Providers

Combined, these capabilities allow MSSPs to operationalize DDoS protection as a scalable, revenue-generating service. With multi-tenancy, automated onboarding, policy tiering, and customer-facing tools built in, the platform enables providers to offer differentiated, high-value protection across thousands of tenants with minimal operational overhead.

Deployment Use Cases

Inbound Attack Protection

For datacenter operators and connectivity service providers, inbound DDoS attacks are a constant operational threat. These attacks originate on the public internet and target both the provider's own network and the customers hosted within it.

Cisco Secure DDoS Edge Protection stops these attacks before they enter the provider's environment. Detection runs directly on the edge and peering routers, eliminating reliance on centralized monitoring or scrubbing centers and avoiding the blind spots created by asymmetric traffic paths.

The platform reacts rapidly—typically achieving mitigation within **~30 seconds**—a critical capability for maintaining SLAs and supporting premium DDoS protection tiers.

If an attack exceeds local capacity or requires broader network actions, the system can automatically apply routing-based methods such as blackholing or redirection into cloud-based mitigation services. This allows providers to combine ultra-fast edge protection with their existing regional or cloud scrubbing solutions for seamless resilience during large-scale events.

For service providers, the outcome is straightforward: **Inbound attacks are neutralized at the network edge, preventing disruption to shared infrastructure, protecting customer experience, and avoiding costly backhaul to centralized scrubbing points.**

Outbound Attack Protection

For Communication Service Providers (CSPs), outbound DDoS activity has become a critical operational threat. Modern botnets—like **Alsuru**, Kimwolfand **ShadowV2**—can generate massive, short-burst attacks directly from infected subscriber devices. These attacks not only damage external targets but also jeopardize the CSP's reputation, trigger abuse escalations, and can lead to upstream filtering or blacklisting of entire IP ranges.

Cisco Secure DDoS Edge Protection is designed to stop these attacks at their source. Detection happens on the CSP's own access and aggregation routers, where subscriber traffic is monitored. The system uses advanced machine learning algorithms to learn "Large Crowds" behavior and finds even small anomalies that indicates outbound attacks on specific targets, the algorithm is designed to differentiate a legitimate surge from an attack.

Mitigation in access networks is performed using **on-box ACLs**, which are pushed directly to the router. This allows the CSP to surgically block or rate-limit only the malicious flows from the infected device, without interrupting legitimate customer traffic. Reaction time is fast—typically under ~30 seconds—preventing the attack from ever leaving the provider's infrastructure.

For CSPs, outbound protection ensures clean address space, fewer abuse cases, and resilience against the new generation of hyper-aggressive botnets. Infected hosts are contained instantly, protecting both the wider internet and the CSP's business reputation.

East-West Attack Protection

Inside a CSP network, **East-West attacks**—malicious traffic moving between subscriber networks, tenants, or data-center segments—pose a growing risk. These attacks never traverse the public internet, so traditional perimeter defenses and cloud scrubbing services cannot see them. Lateral botnet spread, tenant-to-tenant floods, and compromised VMs attacking nearby infrastructure are now common in broadband, enterprise, and data-center environments.

Cisco Secure DDoS Edge Protection addresses this gap by detecting and stopping East-West threats directly on the routers that interconnect customer segments, VRFs, or internal network zones. The system continuously learns normal traffic patterns between internal ranges and flags sudden anomalies—unexpected spikes, port-agnostic floods, or lateral scanning bursts.

Mitigation is executed **on-box using ACLs**, allowing the CSP to isolate only the malicious flows between specific internal ranges while keeping all legitimate inter-customer traffic fully operational. Because mitigation happens at line-rate on the internal routers, the response is extremely fast—typically under ~30 seconds—preventing lateral attacks from propagating or causing cross-tenant outages.

For CSPs, East-West protection delivers a simple outcome: infected subscribers or compromised workloads are contained before they can damage adjacent customers, disrupt shared platforms, or escalate into large-scale internal incidents.

Total Cost of Ownership (TCO) Benefits

Cisco Secure DDoS Edge Protection is explicitly designed to deliver substantial **TCO** improvements for operators by moving detection and mitigation directly to the network edge. By leveraging existing router infrastructure, the platform reduces both capital and operational costs while improving network security and resilience.

CapEx Reduction

The distributed design eliminates the need for costly centralized scrubbing hardware:

- **Leverages Existing Routers:** Detection and mitigation run on Cisco IOS-XR routers, using available CPU and memory.
- **No Additional Hardware:** Docker Agents operate without extra footprint, avoiding incremental capital expenditure.
- **Reduced Scrubbing Center Demand:** Most L3/L4 attacks are blocked at the edge, reducing the need to deploy or scale central scrubbing centers.
- **Lower Environmental Costs:** Less hardware means savings on power, cooling, and data-center floor space, supporting both cost reduction and sustainability goals.

Operational and Network Expenditure (OpEx) Savings

Edge Protection also cuts ongoing costs associated with traditional mitigation workflows:

- **Eliminates Backhauling:** Malicious traffic is dropped at the ingress point, avoiding extra bandwidth provisioning and reducing recurring network costs.
- **Predictable, Stable TCO:** The platform provides a fixed cost model compared to variable costs of centralized scrubbing.
- **Simplified Deployment and Maintenance:** Containerized, zero-touch upgrades reduce operational overhead and lifecycle management expenses.

Overall TCO Impact

By combining existing hardware use, on-box mitigation, and reduced backhaul, operators can achieve **up to 60% reduction in TCO¹** compared to traditional scrubbing models.

For CSPs and data-center operators, Edge Protection is not only a high-performance security solution—it is a **strategic investment** that lowers cost, reduces risk, and improves financial predictability in an era of increasingly aggressive and sophisticated DDoS threats.

More than that, the licensing scheme of the product which is attached to the router FCM (Flexible Consumption Model) created a predictable cost structure without any unexpected surprises.

¹ Based on Cisco data and analysis

Summary and conclusion

Cisco Secure DDoS Edge Protection represents a fundamental shift from the legacy centralized scrubbing model. It counters the latest threat environment—characterized by the scale of Alsuru and the stealth of ShadowV2—by moving **intelligent detection and line-rate mitigation directly to the network edge. The use of Machine Learning (ML)** for behavioral baselines and specialized algorithms for **Carpet-Bombing** and **Geo Anomalies** ensures that evasive, low-volume, and micro-burst attacks are detected instantly. The multi-tiered mitigation, driven by an adaptive scripting engine, ensures a surgical and fast response—typically under 30 seconds—for all attack types, whether inbound, outbound, or East-West.

Cisco Secure DDoS Edge Protection is a strategic investment for service providers, utilities, large enterprises and data-center operators. By leveraging existing Cisco IOS-XR router infrastructure and eliminating the need for traffic backhauling to centralized scrubbers, the platform achieves significant **Total Cost of Ownership (TCO) benefits**, with operators potentially seeing up to a **60% reduction** compared to traditional models. It provides a high-performance security solution that reduces cost, lowers risk, and ensures predictable expenditure in the face of increasingly sophisticated DDoS threats.

Glossary:

- **NTP** – Network Time Protocol, a protocol used to synchronize clocks across networked devices. It can be abused in reflection attacks to amplify traffic.
- **CLDAP** – Connectionless Lightweight Directory Access Protocol, a directory-service protocol that can be abused for reflection and amplification attacks.
- **WS-Discovery** – Web Services Dynamic Discovery, a protocol used by devices to discover services on a local network. Exposed devices can be abused in UDP reflection attacks.