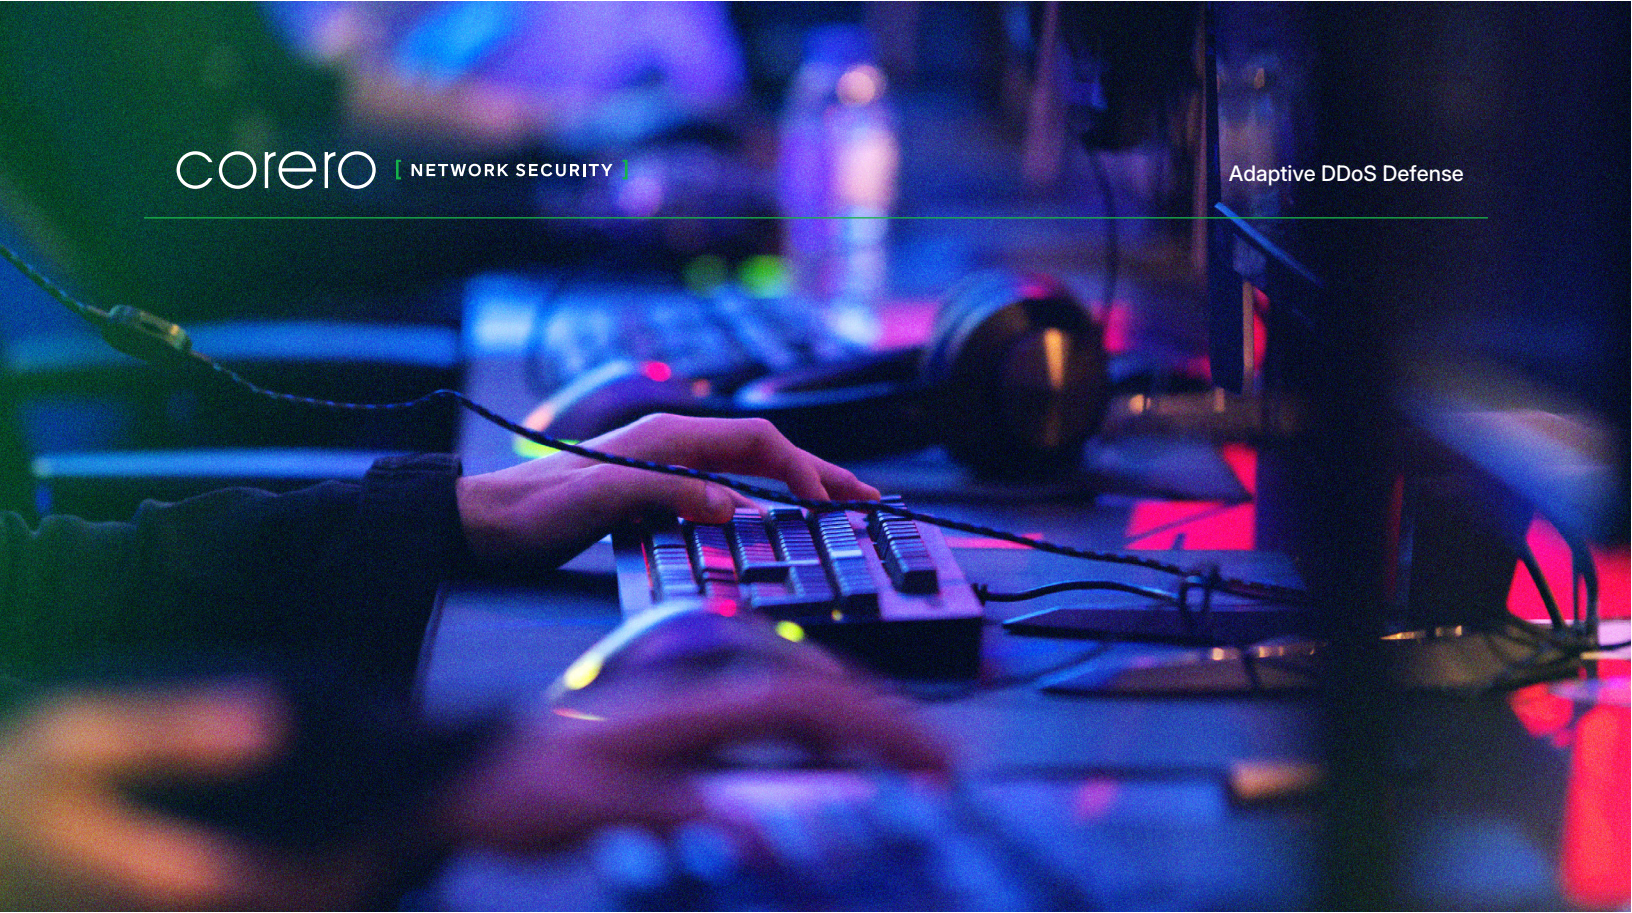


The DDoS Threat to Online Gaming: Challenges and Protection Strategies





Executive Summary

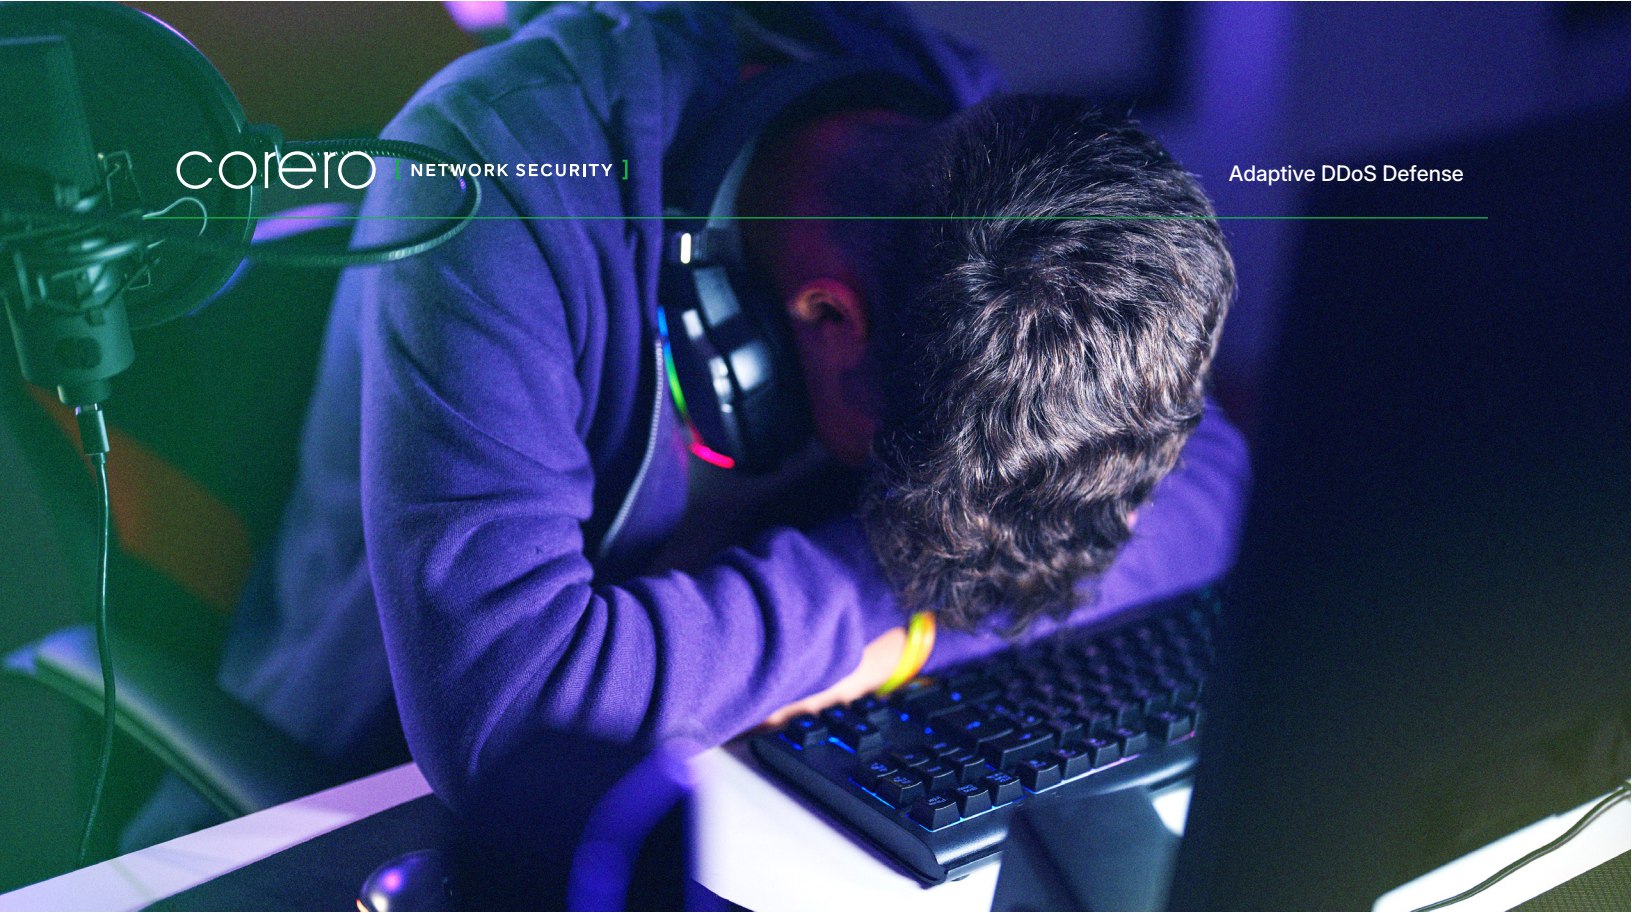
Online gaming platforms face a persistent and evolving threat from distributed denial-of-service (DDoS) attacks. These disruptions strike at the heart of the gaming experience, causing latency, disconnects, and outages that frustrate players and jeopardize business performance. Whether targeting login services, voice chat, matchmaking, or core game servers, DDoS attacks are now frequent, multi-layered, and increasingly sophisticated.

This white paper provides a practical overview of the DDoS threat landscape as it relates to online games and evaluates the three primary defense models available to game hosters: on-premises, hybrid, and cloud-only protection. Each model has merits—and trade-offs—depending on infrastructure design, latency sensitivity, and operational strategy.

Along the way, we highlight why integrated hybrid defenses are gaining traction, how modern on-prem solutions have evolved beyond their legacy limitations, and what cloud-only models can and can't offer in terms of simplicity and scale. Whether operating a single tournament platform or a globally distributed multiplayer network, infrastructure teams must match their DDoS posture to the needs of their player experience—and this paper is designed to help.

Contents

Introduction		04
The DDoS Threat Landscape in Gaming		05
Why Game Platforms Are Prime Targets		06
Key Challenges for Online Game Hosters	<ul style="list-style-type: none"> Real-Time Response Requirements Latency Sensitivity Traffic Complexity and Filtering Accuracy Mitigation System Performance Under Load Scaling and Business Continuity 	07
On-Premises Protection	<ul style="list-style-type: none"> Deployment Flexibility Performance and Control Scalability and Hybrid Readiness Operational Models Summary 	09
Hybrid Protection	<ul style="list-style-type: none"> How It Works Integration Quality Matters Latency and Coverage Considerations Summary 	11
Cloud-Only Protection	<ul style="list-style-type: none"> Key Characteristics Benefits Trade-Offs Summary 	13
Comparative Summary		15
Conclusion		17



Introduction

Few digital services are as demanding—or as unforgiving—as online games. A successful multiplayer game requires more than creativity and gameplay balance; it depends on responsive infrastructure and unbroken availability. Players expect a seamless experience. Even momentary lag can disrupt a match, and sustained outages can drive away users entirely.

The infrastructure that powers these experiences has never been more exposed. DDoS attacks are a constant presence across the internet, and gaming has become one of their most frequent targets. The reasons are varied—ranging from competitive sabotage and trolling to financially motivated extortion—but the outcomes are predictable: service degradation, lost revenue, and reputational harm.

This paper explores the evolving tactics behind these attacks, the specific pressures they place on game hosters, and the strategic options available for defending against them—without sacrificing player experience.

The DDoS Threat Landscape in Gaming

DDoS attacks are no longer rare, and in the gaming industry, they are anything but random. In recent years, gaming platforms have consistently ranked among the most targeted sectors. Industry reports show that gaming accounted for nearly one-third of all DDoS attacks globally in 2024, with some months seeing over 25 billion malicious requests directed at gaming infrastructure.

Attackers have moved beyond simple volumetric floods. While massive traffic surges remain common—such as the 2 Tbps and 5.6 Tbps attacks reported in late 2024—what’s increasingly prevalent is the use of multi-vector tactics.

A single attack might involve:



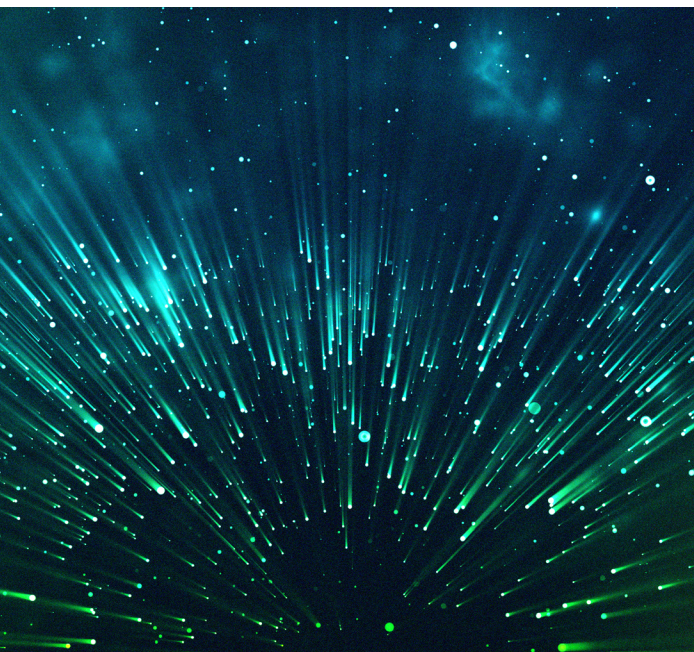
High-bandwidth UDP floods



Application-layer HTTP abuse



Stateful protocol exploitation all at once



Tactics like “carpet bombing,” where traffic is spread thinly across many IPs to evade detection, have become more widespread.

This evolution means that DDoS mitigation is no longer just about scale. It’s about adaptability, protocol awareness, and the ability to respond at multiple layers—network, transport, and application—simultaneously. For game hosters, this raises the bar for what constitutes “effective” protection.

Why Game Platforms are **Prime Targets**

Games are attacked more often than many other services, and not simply because of their popularity. The motivations are as layered as the attacks themselves.



Competitive interference:

Players or teams may use DDoS attacks to knock opponents offline during tournaments or ranked matches.



Trolling and disruption:

Some attackers seek attention by spoiling major releases or online events.



Extortion (RDoS):

Criminal groups threaten sustained outages unless a ransom is paid—often targeting game companies during launches, holidays, or promotional events.



Low barrier to entry:

With DDoS-for-hire services easily available on underground markets, even non-technical actors can launch impactful attacks for a few dollars.



These dynamics make gaming an especially attractive target: high visibility, low cost to disrupt, and the potential for emotional or financial leverage.

Any title with a loyal player base, competitive ladder, or monetized in-game economy is at risk.

Key Challenges for Online Game Hosters

Defending against DDoS attacks in gaming requires more than perimeter filtering. It demands real-time responsiveness, accuracy under load, and architectural choices that preserve gameplay quality—even during an attack.



Real-Time Response Requirements:

DDoS mitigation in gaming operates on a much tighter clock than in other industries.

A five-minute outage in a SaaS application may be recoverable. In gaming, five seconds of lag can prompt mass disconnects, negative social media exposure, and damage to user trust.

Mitigation needs to begin within seconds of attack onset—ideally without waiting for human intervention. Legacy response models that rely on manual policy adjustments or delayed telemetry often react too slowly to be effective.



Latency Sensitivity:

Gaming—especially competitive genres like shooters or battle royale—is highly sensitive to latency and jitter. Any added network hops or delay during mitigation must be minimized.

Protection systems that redirect traffic through distant scrubbing centers can introduce additional milliseconds that players will notice. As such, mitigation strategies must be architected close to the player base or the game servers themselves, with as few detours as possible.



Traffic Complexity and Filtering Accuracy:

Game traffic is not uniform. Many titles use proprietary protocols over UDP or TCP, sometimes on custom ports, with bursts of activity tied to in-game events. These patterns are difficult to differentiate from attack traffic using volumetric or application-layer methods. A poorly tuned filter may block real players (false positives) or allow malicious traffic through (false negatives). This is a tightrope, and one that demands deep protocol awareness and context-sensitive filtering strategies.



Scaling and Business Continuity:

Attackers often time their activity to coincide with known traffic surges: major updates, promotional events, or peak evening hours. This means that mitigation solutions need to scale not just to absorb the attack itself, but to handle legitimate player spikes as well. In some cases, a lack of headroom during these critical windows can result in self-inflicted degradation if protection systems become overwhelmed.



Mitigation System Performance Under Load:

The mitigation solution itself must be resilient under duress. If an inline appliance becomes a bottleneck during an attack, it may compound the problem it was meant to solve. Systems must handle high packet-per-second rates, sustain performance under multi-vector load, and fail open rather than fail closed in worst-case scenarios. High-throughput hardware, scalable cloud platforms, or layered approaches may be required depending on the environment.

These challenges aren't theoretical—they're central to the architectural decisions game infrastructure teams must make when choosing a DDoS mitigation model.



On-Premises Protection

On-premises DDoS protection refers to mitigation that takes place within the hoster's infrastructure—typically inside the data centers or hosting environments where game servers reside. This model has long been valued for its low-latency performance and direct control, but it has also carried a reputation for being resource-intensive or inflexible. That reputation is increasingly outdated.

Modern on-prem solutions have evolved to offer deployment flexibility, integration with hybrid models, and managed service options that reduce operational overhead. When properly architected, they can serve as a highly effective first line of defense, mitigating attacks locally and instantly without introducing the performance or visibility trade-offs associated with cloud-only approaches.

Deployment Flexibility

Contemporary on-prem offerings support a range of deployment models, including:



Inline and virtual inline modes for direct traffic inspection



Edge mitigation using flow telemetry (e.g., NetFlow, sFlow), which allows the device to inject filtering policies without being in the path



Detect-and-redirect architectures that reroute only suspect traffic to a local scrubbing instance

This flexibility allows game hosters to design a DDoS strategy that matches their network and business architecture, rather than retrofitting around a rigid hardware stack.

Performance and Control

On-prem solutions provide immediate, localized mitigation that avoids the need to reroute traffic through external scrubbing centers. This is essential for latency-sensitive applications like competitive gaming, where even a few milliseconds can affect the user experience.

In addition to performance, on-prem models offer granular control over detection logic and filtering policy. Game-specific protocols—often running over custom UDP/TCP formats—require fine-tuned inspection that can't always be achieved with cloud-based or generic filters. An on-prem appliance, managed directly by the security or infrastructure team, can apply context-specific rules tailored to each title.

Scalability and Hybrid Readiness

While on-prem systems are naturally bound by uplink and appliance capacity, that doesn't mean they can't scale. Many modern solutions:



Support high-throughput appliances (e.g., 400G interfaces)



Use telemetry and sampling to cover traffic at scale



Integrate with cloud scrubbing for volumetric overflow scenarios

Some even support hybrid failover as a built-in feature, using threshold detection or telemetry to trigger upstream mitigation. In this way, on-prem no longer means isolated—it can serve as the intelligent front end to a larger, adaptive architecture.

Operational Models

Historically, on-prem required deep expertise and 24/7 monitoring. Today, many solutions offer:



Managed services:

Where vendors provide tuning, threat updates, incident support, and even total management of the deployment



Automated mitigation:

With pre-defined policies that activate instantly upon attack detection



High availability (HA):

Including fail-open behavior, configuration sync, and hardware redundancy

This makes on-prem viable even for smaller teams—as long as the right operational model is selected.

Summary

On-prem protection remains a compelling option for organizations that prioritize latency, control, and precision. It's no longer a heavyweight solution reserved only for enterprises with large network teams; today's on-prem systems can be lightweight, smart, and fully integrated into hybrid strategies.

Hybrid Protection

Hybrid DDoS protection combines local mitigation (typically via on-prem appliances) with on-demand cloud scrubbing capacity. In this model, the on-prem system handles day-to-day attacks and latency-sensitive filtering, while the cloud provider stands by to absorb large-scale volumetric floods when needed. The hybrid approach is popular with mid-to-large-scale gaming platforms that need both fine-grained control and elastic capacity. When implemented well, hybrid solutions offer the best of both worlds: performance under normal conditions, and resilience under extreme stress.

How It Works

Hybrid configurations typically operate with:



Always-on local protection:

Catch most attacks instantly and with minimal latency



Threshold-based cloud activation:

Using triggers such as link saturation or packet rate anomalies to initiate failover



Routing changes via BGP or DNS:

Redirect traffic through the cloud provider when thresholds are breached

When the attack subsides, traffic can return to normal routing—often with no impact on players.

Integration Quality Matters

Not all hybrid solutions are created equal. Some are loosely coupled—requiring teams to manage separate systems, coordinate policy updates, and manually handle cloud diversion events. Others are fully integrated, offering:



Shared management and telemetry between on-prem and cloud layers



Unified policy enforcement



Automatic failover and rollback procedures

The tighter the integration, the less operational complexity and risk of inconsistent filtering during transition.

Latency and Coverage Considerations

Under normal conditions, traffic never leaves the hoster's network—ensuring low latency. During cloud failover, some latency is introduced as traffic is routed through scrubbing centers. If those centers are well-distributed and close to the player base, this latency may be negligible. Poor routing, however, can introduce noticeable performance degradation, particularly in real-time games.

Still, hybrid solutions ensure coverage across the full spectrum of DDoS threats—from volumetric floods to application-layer abuse—without permanently routing all traffic through a third party.



Summary

A well-integrated hybrid model gives game platforms the flexibility to handle routine attack traffic with precision, while maintaining the safety net of cloud-scale protection when needed. It allows teams to balance performance, cost, and control—without overcommitting to any one strategy.

Cloud-Only Protection

In a cloud-only model, all incoming traffic is routed through a third-party scrubbing network before reaching the game servers. These services operate at enormous scale and are designed to handle even the largest volumetric attacks without requiring on-prem equipment or dedicated staff.

This model is particularly attractive to cloud-native game platforms, small or fast-growing studios, and companies that prioritize **ease of deployment over fine-tuned control**.

Key Characteristics

Cloud DDoS services are typically deployed in one of two ways:



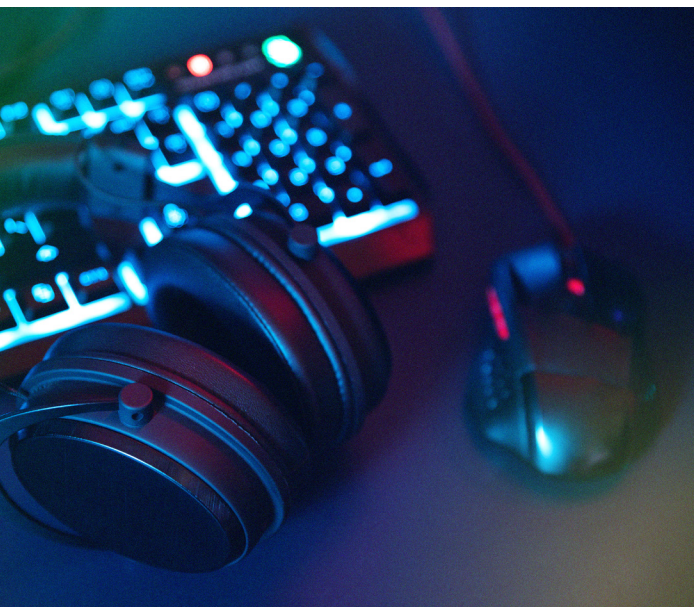
Always-on:

All player traffic is permanently routed through the provider's network (e.g., via anycast IPs)



On-demand:

The hoster only redirects traffic to the provider when an attack is detected (via BGP or DNS)



Always-on provides the fastest mitigation but introduces persistent third-party routing, which may affect latency or traffic transparency.

On-demand can reduce latency under normal conditions but risks a short window of exposure during failover.



Benefits

Massive scale: Top providers can absorb multi-terabit attacks with global reach

Operational simplicity: Little to no on-site infrastructure required

Flexible pricing: Often subscription-based, pay-per-use, or bundled with hosting platforms

Cloud services are often maintained by expert teams with round-the-clock coverage, ongoing threat intelligence updates, and dynamic rule sets that evolve based on global threat behavior.



Trade-Offs

Latency overhead: Depending on routing architecture, cloud scrubbing can add delay—particularly for players far from scrubbing centers

Reduced control: Policies and tuning parameters may be limited or managed by the provider

Cost: Large volumes of traffic can drive higher costs

Protocol challenges: Custom or stateful game traffic may not be handled gracefully by generalized mitigation engines

Visibility gaps: Access to packet-level telemetry, real-time filtering data, and log detail can vary by provider





Cloud-only protection is ideal when simplicity and scale outweigh the need for deep customization. But for latency-sensitive applications or teams that require hands-on control, it may feel restrictive.

Summary

Cloud-only models offer a fast, scalable way to onboard DDoS protection—especially for organizations without dedicated network teams. However, they also introduce trade-offs in latency, visibility, cost, and policy flexibility that game infrastructure teams must weigh carefully.

Comparative Summary

There is no single “best” DDoS protection model—only the model that best fits a given platform’s architecture, player expectations, and operational goals. Each approach offers a distinct balance of latency, control, scalability, and complexity.

 Dimension	 On-Premises	 Hybrid	 Cloud-Only
Latency	Lowest; mitigation occurs locally	Low during normal operations; may rise slightly during cloud failover	Variable; depends on routing and scrubbing center proximity
Control	Full control over filtering, policies, and response logic	High local control with partial cloud delegation	Limited; governed by provider capabilities and policies
Scalability	Bounded by appliance and link size; extendable with cloud	Elastic; handles local and overflow traffic seamlessly	High; designed to absorb massive volumetric attacks
Operational Model	Self-managed or vendor-managed; predictable once deployed	Moderate complexity; best when tightly integrated	Minimal overhead; mostly provider-driven
Best Fit	Low-latency, high-control environments (e.g., esports)	Balanced resilience for medium to large-scale platforms	Simple, scalable protection for cloud-native or smaller studios

What Game Hosters Should Consider When Choosing a Strategy



How sensitive their user experience is to latency



What level of control and customization is required



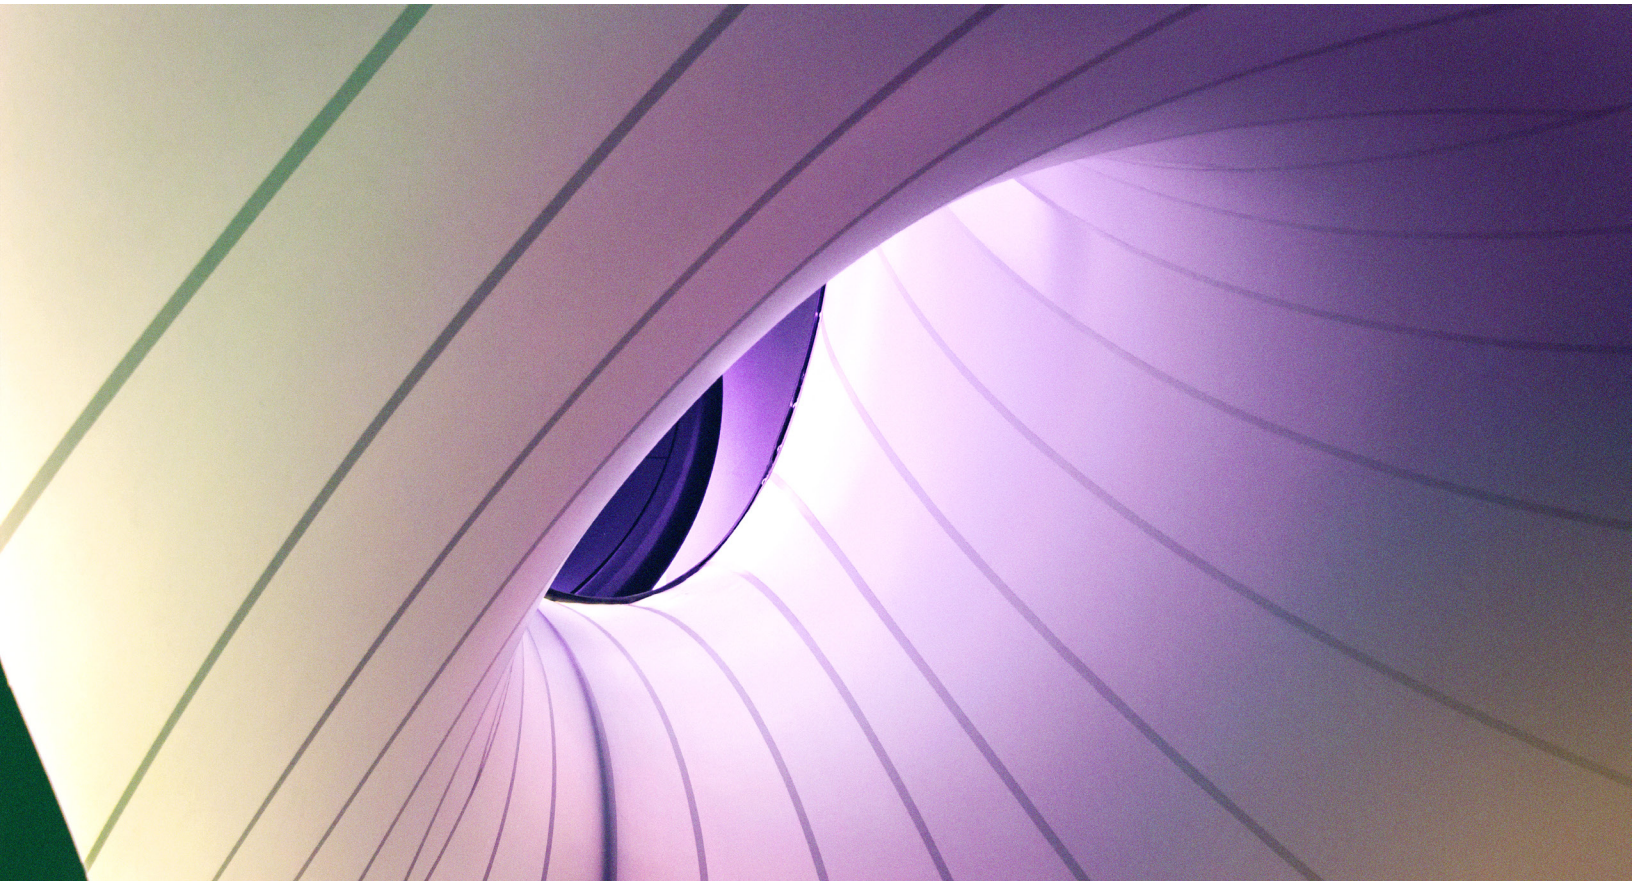
How predictable or elastic their traffic patterns are



What internal resources they have to manage infrastructure



Whether data sovereignty or regulatory constraints influence traffic handling



Many platforms ultimately adopt a layered approach—starting with on-prem for speed and control, then building out hybrid capabilities to ensure resilience at scale.

Others begin in the cloud for speed of deployment, then add localized defenses as infrastructure matures.

Conclusion

DDoS attacks are no longer speculative—they are a daily concern for online gaming platforms.

Whether aimed at degrading service during a major release, extorting a studio mid-event, or targeting key game systems out of malice or competition, these attacks test the infrastructure—and the reputation—of every game hoster.

As the tactics behind these attacks continue to evolve, so too must the defense strategies. Volumetric floods are still common, but today's attacks often include subtle application-layer components, protocol abuse, or distributed probing that demands nuanced, real-time defense.



On-premises solutions

Provide low-latency precision and complete control.



Hybrid models

Add elasticity and cloud capacity for large-scale events.



Cloud-only

Platforms offer immediate scalability with minimal operational overhead.

Each model has its place.

But beneath the labels, implementation matters. A tightly integrated hybrid model performs very differently than two loosely connected systems. A modern on-prem solution with HA, telemetry, and hybrid failover capability is not the same as a legacy inline box. And cloud services that understand gaming protocols will outperform generic scrubbing services that don't.

The right approach balances player experience, operational fit, and long-term resilience. With informed planning, game companies can meet the DDoS challenge head-on—ensuring uptime, protecting brand equity, and keeping players engaged even when under attack.

That's the real measure of success: not just surviving an attack, but maintaining trust in the moments **when it matters most.**

SPEAK WITH A SPECIALIST

corero

[NETWORK SECURITY]