

An abstract graphic featuring a central, glowing red sphere. Surrounding this sphere are several clear, elongated water droplets or ice shards that appear to be melting or falling towards the center. The background is a light blue gradient.

Red Button **DDoS Threat Pulse**

September 2026

All rights reserved. This document is protected by copyright and any distribution, reproduction, copying or compilation is strictly prohibited without the prior written consent of Red Button. No part of this document may be reproduced in any form or by any means without the prior written authorization of Red Button. This document and information is intended solely for the internal use of authorized Red Button Ltd. customers, for the limited purposes set forth herein. While every precaution has been taken in the preparation of this document, Red Button assumes no responsibility for errors or omissions.

In this Issue

- [DDoS and AI: What the Evidence Shows](#)
- [Norway, Month Two: The Attack Surface Expands](#)
- [When the Attack Looks Like a Routing Fault](#)
- [When the ISP Goes Down, the Town Goes With It](#)
- [Fewer Attacks, Bigger Hits, Less Warning](#)
- [Thirty Dollars and a Grudge: What a DDoS Costs Now](#)
- [The Wiki Its Providers Gave Up On](#)
- [Stay Ahead of the Threat](#)

About the DDoS Threat Pulse

Red Button's Threat Intelligence team continuously researches the evolving DDoS attack ecosystem – tracking new attack techniques, emerging botnets' infrastructures, and the threat actors behind them. DDoS Threat Pulse is our way of sharing that intelligence with you: a monthly briefing on the developments we believe matter most to organizations. Our goal is to help you understand what's changing in the DDoS landscape and use that insight to inform your own security posture – whether that means adjusting or hardening your DDoS protection configuration or simply knowing what to ask your team. We hope you find it useful, and we welcome your feedback.

DDoS and AI: What the Evidence Shows



AI is changing the DDoS landscape — but perhaps not in the way some headlines suggest.

Our new report examines the evidence behind AI and DDoS, separating what we can demonstrate from what remains speculative.

We look at how AI is lowering the barrier to launching and adapting attacks, the reality behind today's hyper-volumetric attacks, the rise of multi-vector campaigns, and what "AI-powered" DDoS protection actually means today.

[Read the full report](#)

Norway, Month Two:

The Attack Surface Expands

[Last month's lead story](#) was a three-day attack on ID-porten, the single sign-in Norwegians use to reach public services. The same day this campaign ended, a pro-Russian group calling itself 'Server Killers' announced it was at ["cyber war" with the country](#).

Its next target was education. [Sikt](#), the government agency that runs shared IT services for Norway's universities and colleges, has been under repeated denial-of-service attacks since 29 August. On that first Saturday, 19 of its services dropped out for periods. On 2 September, the websites of several universities went down or slowed to a crawl, and Server Killers [posted a list of eight institutions](#) it said it had hit. Feide itself was knocked unstable on 4 September, and again on 14 September for around four hours. We found no claim for either of those.

On the afternoon of September 4, the website of [UDI, Norway's immigration directorate](#), went down. A spokesman said that "by all indications" it was a cyberattack and that internal systems were untouched, and Server Killers claimed it. On 14 September the [Norwegian parliament's website](#) turned unstable from around three in the afternoon. Its IT department reported "some instability", a post on Telegram promised the attack would last three hours, and the site was back that evening. Neither organization has confirmed it was a DDoS.

CYJAX, a threat-intelligence company that follows the group's announcements, counted 36 claimed victims that included the national rail operator, the postal service, universities, government offices, and Norway's own national security authority. The national security authority's public line has been [deliberately flat](#). In its view these attacks aren't a threat to national security in themselves. They are "a low-cost tool" for anyone after attention, irritation and weakened trust in Norwegian authorities.

WHAT THIS MEANS FOR YOU

A campaign is not an incident, and it can't be planned for like one. More than a month in, the things that give way are rarely technical. It's the on-call engineer who has been woken up three times. It's the holding statement nobody has touched since the first attack, or the escalation path that depended on one person who is now on leave. If your organization attracts this kind of attention, plan for the long version. That means customer messages already written for the second and third round, and thresholds agreed in advance for switching your protection into a stricter mode.

When the Attack Looks Like a Routing Fault



MEO is one of Portugal's largest telecom operators, and on Monday 14 September its customers lost service three times. The first wave hit just after midnight, the second between half past five and half past seven in the evening, and the third from a quarter to eleven into the early hours of Tuesday. At the evening peak [more than 7,200 people](#) had logged complaints on DOWNDetector, a site that collects outage reports from users.

MEO's first statement described "a mass attack" that had congested and degraded its international network, with no intrusion and no access to data. People outside the company were looking at something else. Internet routing is public: anyone can watch the BGP announcements that networks use to tell each other where to send traffic. A Portuguese tech site counted more than 3,600 routing changes on a single MEO address block overnight, against 17 the night before. [One outlet's first headline](#) said the network data pointed to a failure in how MEO was managing its routes.

On Tuesday MEO published a [fuller explanation](#). It confirmed a distributed denial-of-service attack aimed specifically at its infrastructure, with enough traffic to congest its international links. The routing instability was real, it said, but it came from "the degradation caused by the

attack and the operational response needed to mitigate it". The broken connections with neighboring networks were "a consequence of the instability caused by the attack, not its cause". MEO reported the incident to Portugal's national cybersecurity center. It gave no attack size and no customer count, and nobody has claimed it. Some of the route-leak screenshots passed around on social media that week turned out to date from July and January.

Both readings looked at the same event. The attack was the cause, and the routing churn was what the attack, and MEO's own efforts to fight it, looked like from outside. That matters to anyone who runs a network, because the first version of your incident the public hears may be assembled from your routing data before you've said a word. It's also why [our own guide to handling an attack](#) starts by confirming that it is one.

WHAT THIS MEANS FOR YOU

Mitigation leaves fingerprints. Pushing traffic through scrubbing, withdrawing and re-announcing address blocks, dropping a congested peering link: all of it shows up in public routing data within minutes. And there are people who watch that data closely and post what they see. It helps to know in advance which of your mitigation steps are visible from outside, and what they will look like to someone who has no idea what you're doing. Keep a holding statement ready that tells customers what they'll actually notice, which is more useful to them than being told you're "aware of an issue". And in the first minutes of any incident, settle the question MEO's critics raised before anyone else answers it for you: is this an attack, or something inside your own network?

When the ISP Goes Down, the Town Goes With It

Eagle Mountain is a city in Utah served by a local internet provider, Direct Communications. On Thursday, 17 September, customers started reporting slow speeds and dropped connections. Two days later the company said its network was seeing “unusually high traffic consistent with a DDoS attack”, which had also hit the upstream networks it connects through.

What follows shows how far a local ISP's outage can reach. On Friday the phones at city hall went down, and the rest of city hall's services soon followed. A fix at the upstream provider restored things on Saturday night, then service failed again on Sunday. The company temporarily closed its local office, and its support lines were unreachable throughout. Service came back on [Monday 21 September](#), with outside security specialists brought in. It still isn't clear whether Direct Communications or its upstream provider was the real target.

In Poland, the wave of attacks on small local providers we covered in August hasn't stopped. ClimaxNet, a provider in the south of the country, was hit on 9 September, and card terminals in some of the shops it serves stopped taking payments. Two more operators reported attacks at the end of the month. One of them, Telnet Ropczyce, said it had [fought off a carpet-bombing attack](#) spread across about 3,500 of its addresses, each receiving anything from a few tens of megabits to several gigabits per second. Its owner also publicly denied a rumor that he'd bought vouchers to pay the attackers.

[Carpet-bombing](#) spreads a flood thinly across a whole range of addresses so that no single address looks like it's under attack, and it keeps growing. NETSCOUT reported that the technique now accounts for 4.54% of attacks. On 26 September a German hosting provider, Unesty, logged [a “large carpet bomb DDoS attack against all Unesty subnets”](#). It lasted about an hour, and the filtering needed to stop it affected some customers' UDP applications.

WHAT THIS MEANS FOR YOU

Ask your connectivity provider the questions you'd ask of a cloud vendor. Does its DDoS protection prepared to subnet attacks? How much upstream capacity does it have, and what happens to your circuits if it must blackhole traffic to protect the rest of its network? Most organizations never ask, because the ISP is invisible until the day it isn't. Two practical things come out of Eagle Mountain and Poland. Anything that takes money, from a shop's card terminal to a branch payment system, needs a second path that doesn't share the same provider. And if you are the provider, your support line is part of the service. When the internet goes down, it's the only way your customers have left to reach you

Fewer Attacks, Bigger Hits, Less Warning

Two large half-year reports came out in September and agree on the shape of things. [Link11](#), a European DDoS-protection provider, said the number of attacks on its network fell 42% in the first half of 2026. Almost everything else went up. The largest attack it saw reached 2.3 Tbps, 85% above last year's peak, the highest packet rate climbed to 322 million packets per second, and the total volume of attack traffic rose 61%.

NETSCOUT's [report](#) counted more than 9.1 million attacks worldwide in the same period. Attacks reaching a terabit per second or a billion packets per second rose 1,286% year on year. It also recorded its first attack over a terabit per second carried over IPv6, the newer version of the internet's addressing system. One flood against a hosting provider reached 3.5 Tbps and lasted two minutes. The share of attacks lasting a minute or less has roughly doubled since late 2023.

At the end of the month Link11 also published [a case study](#) that shows what this looks like up close. An attack on one customer climbed to 370 Gbps in under 60 seconds, from roughly 8,500 source addresses, mostly in Brazil, India, Venezuela and Azerbaijan. Almost every packet was a fragment, one piece of a larger packet that has been split up for transmission, and there was "not a single complete packet" among them. Link11's best guess is that the attackers sent only the first piece of each packet and never the rest. Whatever receives fragments has to hold each piece in memory until the others arrive, and if they never come, that memory stays tied up until a timer runs out. Fill those buffers fast enough and there's nothing left for real traffic. This one was filtered in Link11's scrubbing network before it reached the customer in any meaningful volume.

Put together, the numbers say attacks are getting rarer, bigger and shorter, and that a growing share of them will be over before anyone has finished reading the alert. That is the whole idea behind what are known as [burst attacks](#).

WHAT THIS MEANS FOR YOU

Walk through your mitigation and mark every step that needs a person: someone to notice the alert, open a ticket, call the provider, approve a rule change. Each of those is a delay, and against a flood that reaches full size in under a minute, any delay is longer than the attack's ramp. The steps that can't be automated should at least be approved in advance. While you're there, look at the stateful devices sitting in your traffic path, firewalls, load balancers and VPN gateways especially. Their connection tables and fragment buffers can run out long before your bandwidth does, and a device that fails closed takes everything behind it down too. Then plan for the second round: fewer than half of Link11's targeted customers went a month without another attack after the first wave.

Thirty Dollars and a Grudge: What a DDoS Costs Now

The security company [Kaspersky](#) read through more than 256,000 DDoS-related posts on dark-web and deep-web forums and found 148 DDoS-for-hire actors active since 2024. A one-off attack arranged through a broker costs anywhere from \$30 to \$2,000. Subscriptions usually run for 30 days and cost between \$1,000 and \$30,000 a month, for attacks lasting anything from a minute to two days. Application-layer floods, the kind aimed at websites and APIs, were on offer from 92% of sellers.

NETSCOUT's half-year report is pointedly titled "Attacks on Demand". It describes a market with a free tier at the bottom and, at the top, several terabits per second advertised behind subscriptions of about \$20,000 a month. Lower down the supply chain, a seller calling itself WraithTools spent August advertising a Windows botnet kit called x47.c for \$200, plus \$150 for the DDoS module. At that price, plenty of people will try it anyway.

On the enforcement side, the FBI's Anchorage field office and Canada's federal police [seized the domains of NightmareStresser](#) on 15 September, as part of Operation PowerOFF, the long-running international crackdown on [DDoS-for-hire services](#). The US Justice Department called it one of the world's longest-running services of its kind, used for "hundreds of thousands of actual or attempted DDoS attacks" since 2022. It's a real win, and a familiar one: one of NightmareStresser's domains was among 48 seized in December 2022, and the service simply carried on elsewhere. No arrests have been announced.

Cheap attacks change who gets attacked. Late on 24 September, [Photopea](#), a free browser-based image editor made by a single developer, went down under a DDoS attack that lasted for hours. A few hours in, the developer posted that someone was "asking me for \$10,000". The demand came from a persona calling itself Lumen Cyber, which presents itself on Telegram as an "ethical hacker".

WHAT THIS MEANS FOR YOU

When an attack costs less than dinner out, "who would bother attacking us?" stops being a question that protects anyone. A former customer with a grudge, or a rival with a free afternoon, can now buy what used to need a botnet of their own. Two things follow from that. Plan on being attacked at some point, and size your protection for the common case first, since application-layer floods are what nine in ten sellers offer. And decide in advance what you would do with a demand like Photopea's. It's priced to look cheaper than the outage, which is exactly the point, and paying it tells the attacker the method works.

The Wiki Its Providers Gave Up On

The Cutting Room Floor is a long-running wiki that documents unused content hidden inside video games. From 27 August to 9 September, it was under attack for 12 days and 16 hours, mostly raw network traffic peaking at around 2 Gbps. In a [postmortem](#) published on 23 September, its owner describes what each of its providers did once he became the target.

His hosting provider, Linode, null-routed the server the first night because the attack was affecting its other customers, later told him the attack had stopped, and then admitted the block had never been lifted. A backup copy on a fresh address went behind a second provider with a \$10 spending limit. Within 14 hours the attack had pushed the bill to \$21 from nearly 700,000 requests, and the provider deactivated the account without warning for “bandwidth abuse”. The attacker had found the backup because, as the owner admits, it accepted requests that hadn't come through the protective proxy. A third provider's under-attack mode finally held, including against a single HTTP flood of 87 million requests, and the site came back on 1 September. His conclusion is simple: you need a big enough network in front of you, and [an origin that refuses everything that didn't come through it](#).

WHAT THIS MEANS FOR YOU

There's a question most hosting customers never ask until it's too late: what does my provider do to me when I'm the one being attacked? For a lot of providers, the honest answer is that they protect everyone else first, by null-routing you, suspending you or billing you. None of that is unreasonable, but it's far better learned from the contract than in the middle of an incident, so check the thresholds, the notice you'd get and the route back to service. Look at spending caps with fresh eyes as well. A cap that switches your service off is an availability control, and an attacker can trip it on purpose.

Stay Ahead of the Threat

Red Button helps organizations understand and validate their DDoS resilience through authorized, real-world attack simulations, technical advisory, and continuous threat intelligence. If anything in this issue raised a question about your own environment – the shared services sitting in front of your login, how your provider behaves when filtering turns into blackholing, or what your systems do to everyone else while they recover – your Red Button team is here to help.

“The DDoS threat landscape evolves constantly. Our job is to make sense of it before it becomes your problem.” – Red Button Red Team & Threat Intelligence

If you'd like to discuss how these trends apply to your environment, contacting your technical manager is the best place to start.

DDoS Threat Pulse is produced by Red Button's Red Team & Threat Intelligence function and compiled from publicly available research and vendor reporting. It is intended as general awareness content for Red Button customers and does not constitute a personalized risk assessment of your environment



Confidentiality & Proprietary

This document contains confidential and proprietary information of Red Button Ltd. ("Red Button") and may not be disclosed outside your organization, transmitted, duplicated, or used in whole or in part for any purpose other than its intended purpose. Any use or disclosure, in whole or in part, of this information without explicit written permission of Red Button is prohibited. Red Button makes no warranty that the information contained in this document is complete or error-free.