Nine Internet Solutions AG
Badenerstrasse 47
CH-8004 Zurich

Postmortem: DDoS attack against our network, 11 to 13 August 2026

Summary

On the evening of 11 August 2026, a large DDoS attack targeted the network of one of our customers for whom we provide connectivity. Based on the timeline and pattern of the attack, we believe this customer was the primary target, but cannot fully rule out that others were targeted as well. The resulting traffic degraded our own network as well. A few hours later, the attack was also directed at our own services directly.

- *Duration:* Tuesday 11 August, approx. 19:15 CEST, to Thursday 13 August, approx. 12:51 CEST
- *Affected systems:* every application on deplo.io, our website, Cockpit, and our ticketing system
- *Pattern:* not a continuous outage, but a series of recurring interruptions (The full sequence is on our status page)
- *Peak volume:* an estimated 500 to 600 Gbit/s in aggregate, exceeding our uplink capacity several times over; two upstream providers measured traffic of 260 Gbit/s

There was no unauthorised access to data and no compromised systems. The attack was designed solely to cause overload.

This also exposed gaps in our own safeguards, which we detail further below. If you run your own domains on deplo.io, the final section contains a specific recommendation.

Sequence of events

Period	Events
Tue 11 August, from approx. 19:15	The first wave hits a customer IP and saturates our uplinks. We contain it with a blackhole route at our upstream providers. More waves follow into the night, including against our website.
Wed 12 August, from approx. 07:10	Further waves follow, with targets shifting constantly. We move our exposed applications behind a CDN with DDoS protection and route the affected networks only over paths where our mitigation takes effect. Around 19:15 the attack stops, confirmed by an upstream provider's measurements.
Thu 13 August, until approx. 12:51	We withdraw the mitigations in a controlled manner. All applications and services are back in normal operation.

How the attack worked and how we mitigated it

The attack used UDP amplification. The attacker sends small requests to open services on the internet and enters the victim's IP address as the sender. The responses therefore go to the victim, and they are many times larger than the requests. A small amount of

bandwidth thus produces a very large attack volume. CISA describes the method in detail under UDP-Based Amplification Attacks.

Once the uplink is saturated, filtering inside the network is no longer effective. At that point, legitimate packets are already being dropped because the uplink is overloaded. An attack of this kind can only be mitigated further up, at the upstream providers. The tool there is blackholing: The provider discards traffic destined for a specific IP address instead of delivering it to us. That protects the network and every other customer, but it also affects legitimate traffic to that address. As a consequence, affected IP addresses became unavailable by design while blackholing was active, including legitimate traffic. We temporarily gave up the availability of individual targets to keep the network stable for everyone else.

The attacker adapted. After we moved our website to a new IP address, that address was attacked within minutes as well. The attack was not only a matter of bandwidth. The number of requests was also high enough to overload individual services. While blackholing protected the network from further congestion, it could not keep the targeted services available because it blocks all traffic to the affected IP address. A different mitigation approach was therefore needed to distinguish attack traffic from legitimate requests.

To restore availability, we therefore introduced an additional mitigation layer. On 12 August we moved our exposed applications behind a CDN with DDoS protection, and completed the migration the following day. Our customer also evaluated and implemented additional mitigation measures on their side. Attack traffic now terminates at the protection provider, while legitimate requests reach our applications. Operating our own network allowed us to make routing changes directly during the ongoing attack and coordinate the additional mitigation measures without depending on changes to our underlying network infrastructure. It is why an attack of 500 to 600 Gbit/s did not become a complete outage.

Where our safeguards fell short

Our detection did not cover every network we operate. It was scoped to our own addresses. Customer-owned networks that we announce on their behalf were not included. The first wave hit exactly such a network and had to be contained manually, which cost around 90 minutes. We closed this gap the same night.

The effectiveness of blackholing was not verified per path. This only works for peers with whom blackholing was already agreed in advance, or for traffic reaching us through an internet exchange's route servers. At the internet exchanges we connect to, we had no such agreement with any of our direct peers, so blackholing there only covered traffic sent through the route servers, not traffic exchanged directly. No regular per-path verification was in place, and no alerting either. Routing around it was not an option: withholding a single network from a single provider had not been provided for. We developed that capability during the attack, under full load, and put it into production, as

we did the CDN protection. Both measures worked, but were slower and riskier than necessary.

The attack affected uninvolved customers and our contact channels. We run our website on deplo.io, on the same platform as our customers. The attack on it therefore also affected applications that use the same platform. Cockpit and the ticketing system run on separate infrastructure, but were directly attacked at a later stage as well. This produced the most uncomfortable consequence of the incident: customers temporarily lost their application, its management, and part of their communication route to us at the same time.

Our measures

Every customer network we announce is covered by our automated attack detection; adding it is part of the standard process. Our exposed applications sit behind a CDN with DDoS protection, and migrating further services is available as a prepared procedure. We monitor the effectiveness of our mitigations across all links and work closely with our upstream providers on this. The goal is to mitigate attacks without giving up the availability of the target.

We continue to run our own services on deplo.io, on the same platform we offer our customers. Rather than moving them off it, we are hardening the platform: we are developing the deplo.io architecture further to reduce individual applications' dependence on shared addresses. Additionally, Cockpit and the ticketing system are better protected so that you can reach us during an attack. The scenario of this incident is part of our regular exercises.

Our recommendation: use dynamic DNS record types

Use a CNAME or an ALIAS record. An A record ties your domain to a specific IP address on our platform. That removes the option of switching to a different access path at short notice, and that option was exactly what mattered: where we could change the address, availability could be restored. In addition, further applications are reachable via the same address. That architectural issue is ours, and we are working on it. Until then, a CNAME/ALIAS is the most effective measure you can implement yourself.

Put a CDN with DDoS protection in front. For us, that is what restored availability. For most applications the effort is minimal and the running costs are low.

If you are not sure which option fits your application, contact support@nine.ch. We are happy to work through it with you.

Attacks of this magnitude cannot be prevented. What we can influence is how quickly we react and how many uninvolved parties our mitigations affect. That is what we are working on. We apologize for any inconvenience caused and thank you for your understanding as well as patience during the disruption.