LINK11

LINK11 INSIGHTS-FEATURE

Optimizing On-Prem Investments with Hybrid DDoS Protection Solutions

Whitepaper

Table of Contents

Executive Summary

Distributed Denial of Service (DDoS) attacks are becoming increasingly dangerous and an ongoing threat to companies[1]. Many enterprises have taken various steps to protect themselves against this ever-evolving threat. Most commonly, enterprises are purchasing an on-premise solution to protect against DDoS attacks. Strategically, this hides a variety of approaches, such as buying a firewall, or a hardware-based or software-based protection solution. Their commonality is that they are implemented directly on the company‘s own servers or on-site. Although purchasing an on-premise solution is one of the most common strategies for defending against DDoS attacks, organizations often find this solution cannot scale to the size and complexity of these attacks. In the event of an attack, it can then lead to lost revenue, reputational damage or SEO loss, and much more. As a result, the company‘s return on investment and operating costs also decrease.

Hybrid cloud-based DDoS protection services have become popular to achieve a positive cost-benefit effect. On the one hand, these help enterprises maximize their return on investment, and on the other, they help IT decision-makers fully leverage their on-site investments. When companies supplement their IT security strategy with cloud-based DDoS protection, it can unleash their protective capabilities. Large and complex attacks are dynamically mitigated, while the on-premises platform can handle small attacks and provide application-level security. Individual cloud DDoS protection vendors offer different implementations and features of such a hybrid setup. Despite the different options, the Link11 Insights feature criteria provides a solution for large, uncoordinated, and very specialized and coordinated DDoS attacks that can be combined with a wide range of enterprise configurations.

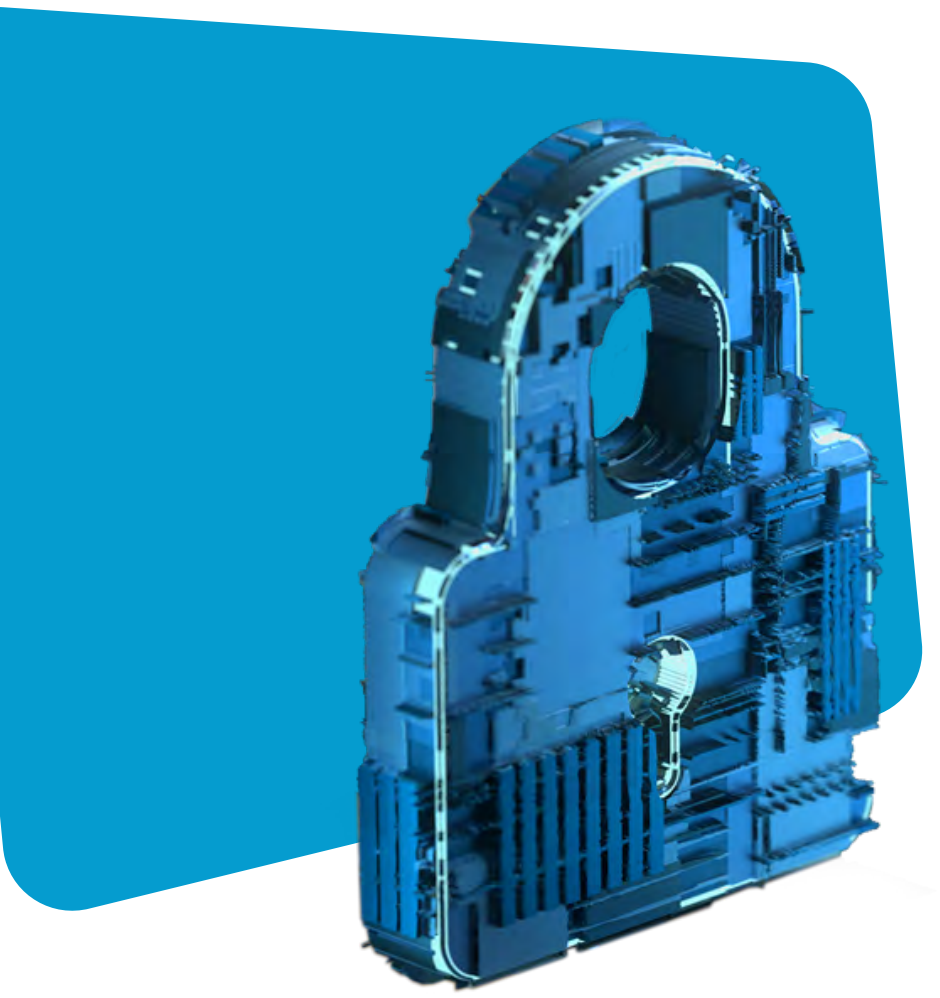

On-premises solutions and their limitations

Increasing business operations and enterprise processes are being conducted online. The Corona pandemic acted as a catalyst for this development. The more everyday business has become digitalized, the greater the need for security measures. Especially for protecting valuable IT resources, companies have increasingly introduced security measures, usually in an on-site edge defence, such as a firewall. These devices successfully defend against low-level DDoS attacks. But, they fail against medium- or large-scale DDoS attacks for the following reasons:

1. the capacity of these devices is limited. Edge devices with 10 Gbit ports, even if they could handle such a line rate, are easily overwhelmed by the size of current DDoS attacks. Such attacks continue to grow and duration.

2. the activity status (online or standby) of these devices makes them vulnerable to state-exhaustion attacks. A small but sophisticated DDoS attack can easily establish tens of thousands of simultaneous connections. These can quickly overwhelm a firewall's state table and block all new valid connections.

3. firewalls can provide basic protection against DDoS attacks (e.g., protection against SYN, UDP, and ICMP floods). This protection comes at the cost of performance, which impacts the firewall's advanced features (e.g., Layer 7 inspection, SSL decryption, and VPN termination).

4. on-premises devices primarily protect applications and systems within a company's data center. However, they are ineffective at protecting web applications hosted in public cloud infrastructures, a growing trend in enterprises.

> **The risk landscape is becoming increasingly unpredictable. Already, attacks are very precise and coordinated, overwhelming on-premises solutions. A cloud-based protection solution pays off for these very specialized attacks**
>
> **– Jag Bains,**
> Vice President of Solution Engineering, Link11

These functional limitations in implementing an on-premise strategy can create a bottleneck or even a „single point of failure“ in the event of a large DDoS attack. This defeats the very purpose of this solution - to protect valuable IT resources.

The hybrid solution - the best of both worlds

A hybrid solution consists of adding cloud DDoS protection to on-premises devices. As a cloud-based protection solution, this platform can absorb large volumes of volumetric and protocol- and application-level DDoS attacks and filter traffic. Malicious traffic is filtered out automatically, and in real time, so only legitimate traffic reaches the enterprise. By combining an on-premises and a cloud-based solution, enterprises can maximize the benefits of both IT security strategies. On the one hand, this includes the investments already made in the on-premise installation, and on the other hand, its functional disadvantages are minimized.

The hybrid approach means that the on-premises DDoS protection applications mitigate attacks locally until they reach a predefined threshold. Once this is reached, the traffic is redirected to the cloud DDoS protection platform, usually via the BGP protocol. The DDoS protection provider should have a capacity of more than one Tbp/s to defend against an attack. This should detect and block all attack vectors and route legitimate traffic back to the customer network. Once the attack drops below the threshold to a certain level, the traffic is routed right back into the company‘s network.

Such a hybrid approach to the necessary DDoS protection combines the strengths of an on-premises protection application with the multi-terabit scalability of the cloud. This makes it a fast, low-latency solution for most attacks, seamlessly providing additional capacity as needed.

Fig 1. Typical Hybrid Topology

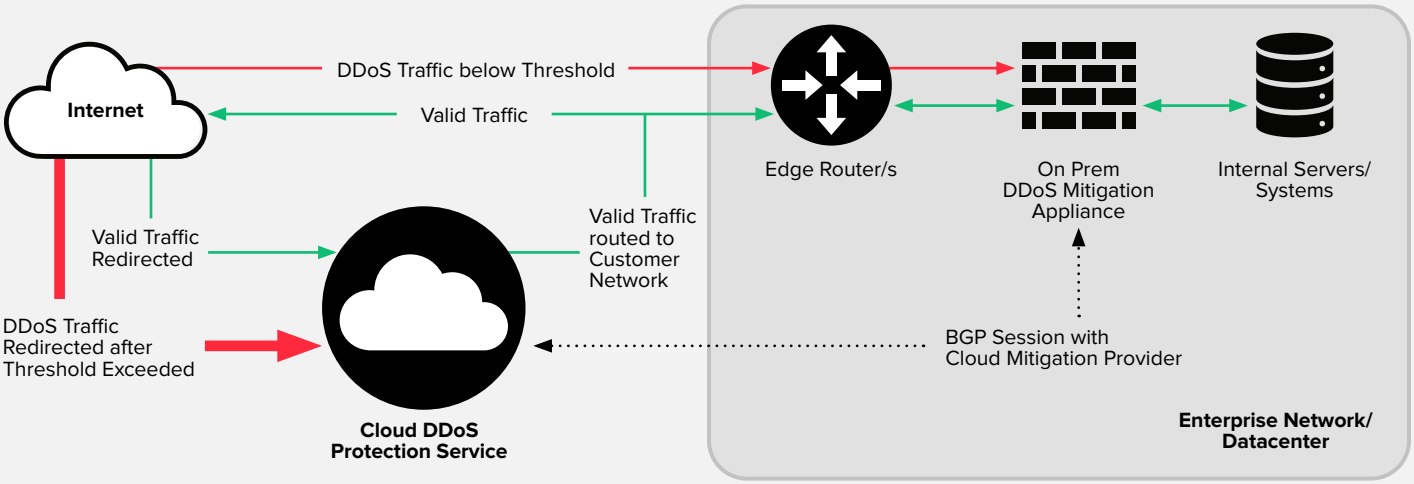

The Link11 Insights approach

A typical hybrid setup has many helpful benefits. At the same time, such a setup often requires a single vendor solution. This can arrange for the company's BGP traffic to be redirected. Especially for companies that do not have a security operations center around the clock, this can be difficult or almost impossible to implement.

The hybrid Link11 product „Insights" counters these limitations and offers decisive advantages:

1. technology independent: The Link11 Insights feature is compatible with hardware from any on-site vendor. The security solution is not limited to specific hardware options.

2. multiple control plane options: Most hybrid solutions require a BGP session between the customer network and the cloud provider. This requires the customer to initiate the redirection of traffic to the cloud provider, an often manual process. With the Link11 Insights feature, organizations can continue to use a standard BGP session and/or set up a Netflow export to the Insights feature. In addition, BGP redirection is automatically performed when a threshold is reached. Because Insights analyses the incoming Netflow stream in real time, the system immediately detects if a threshold is reached. This is followed by automatic BGP redirection without the end customer having to do anything (see Fig. 2).

Fig 2. Link11 Hybrid Topology

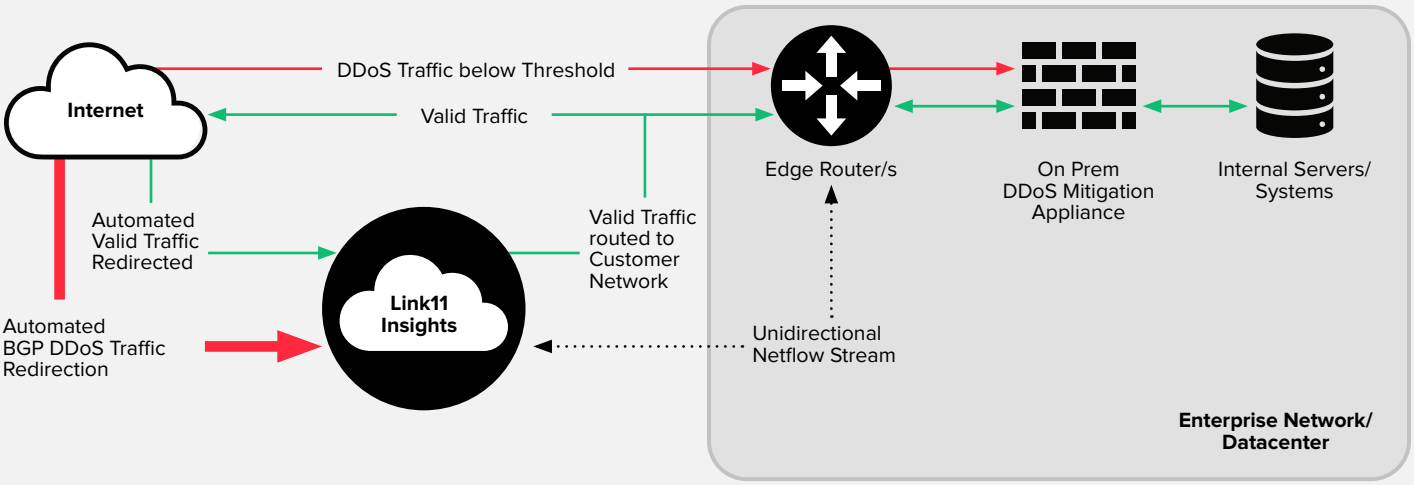

3. multiple options for setting thresholds. Most hybrid solutions look at the total traffic throughput within the customer network on a selected group of interfaces to determine if an anomaly is occurring. With Insights, instead of a selected group, many threshold/condition options can be set, including:

- MB/s per second,
- Packets per second,
- Number of traffic source addresses,
- Number of source countries,
- number of ISPs.

4. „Infrastructure Protection Engine“: At the heart of the Link11 Insights feature is the „Infrastructure Protection Engine“, which employs multiple strategies to defend against and mitigate DDoS attacks, including:

- Fingerprinting: incoming traffic is analyzed, and a unique „fingerprint“ is assigned to each client. Each fingerprint includes hundreds of unique characteristics and is far more sophisticated than an IP address. This ensures that legitimate users can access resources while blocking fingerprints that contain known attack patterns.
- AI analysis: the platform‘s self-learning AI module analyses traffic for malicious activity and can even warn of AI-driven attacks. The module actively disrupts attacks by sending false information to attackers. The attacker is led to believe that an asset has been „shut down“ even though it remains accessible to legitimate users.
- Cross-checking traffic against threat data: The platform cross-checks all traffic against real-time threat data. This determines if the traffic corresponds to known malicious activity. This allows the platform to block malicious activity while accepting legitimate requests.

Each time the platform identifies a new threat, the attack sequence is stored in a database for later retrieval. It is blocked immediately if the module detects the same attack sequence again.

The result of these combined strategies: The shortest time to mitigation compared to leading international security vendors, as shown in an independent **study by Frost & Sullivan**[2]. Figure 3 shows the time-to-mitigate (TTM) between Link11 and other leading cloud-based protection vendors:

Fig 3: comparison chart Time-to-Mitigate

	Link11	Arbor Cloud/ Neustar	Cloudflare Magic Transit	Imperva Incapsula	Akamai Prolexic
UDP bandwidth flood	21 secs.	4 mins.	fail	1 mins.	up to 59 secs.
UDP DNS refl/amp bandwidth flood	18 secs.	4 mins.	4 mins.	1 mins.	up to 59 secs.
Fragmented IP bandwidth flood	0 secs.	4 mins.	4 mins.	2 mins.	up to 59 secs.
TCP SYN flood	0 secs.	4 mins.	2 mins.	1 mins.	5 mins.
TCP SYN-ACK flood	12 secs.	4 mins.	up to 59 secs.	2 mins.	8 mins.
TCP RST „carped bomb“	1 min 20 secs.	fail	fail	1 mins.	5 mins.
HTTPS GET request flood-404	18 secs.	4 mins.	fail	fail	fail
Average TTM	21 secs.	240 secs.	150 secs.	80 secs.	180 secs.

Options and considerations for tech leaders

Depending on how much organizations have invested in their on-premise strategy and what stage of the hardware lifecycle it is in, the challenges and considerations required to address them will vary. If the company has a new implementation, the IT decision maker is incentivized to maximize their investment. In this way, a connection to the Link11 Insights feature will ensure this.

If the contract is about to expire and needs to be renewed, the CISO should rethink their planning and consider a cloud-only solution with Insights. With Insights, organizations gain tremendous capacity and the ability to thwart volumetric and protocol-based attacks as well as complex attack vectors with the advanced „infrastructure protection engine."

In addition, the CISO should consider the Always-On mode when the organization's tolerance for downtime is very low. In Always-On mode, customer traffic is always routed through the Link11 cloud, where it is inspected, filtered and malicious traffic is thwarted. In addition, this mode eliminates the time required to update the global BGP route table, which can take between 45 and 90 seconds[3]. With Link11's short time-to-mitigate (TTM), downtime is minimized, reducing costs and operational burden within the enterprise.

The following graphic is intended to illustrate the importance of TTM. The graphic shows the TTM in mitigating a 15-minute, medium-scale attack (e.g., 20 GB/s) that can exhaust the capacity of an enterprise network. In addition, the TTM is compared in four modes:

a) Always On. In this mode, traffic always passes through the Link11 cloud DDoS platform.
b) Hybrid Insights. Here, Link11 automatically takes care of BGP redirection.
c) Hybrid Manual. Here, the enterprise itself must make the necessary BGP changes to redirect traffic to the cloud DDoS protection provider.
d) On-Premise. The company uses its on-premise solution exclusively.

> **The faster and more accurately DDoS attacks are mitigated, the greater the positive cost-benefit effect of the Insights feature will be for enterprises.**
>
> **– Marc Wilczek,**
> Chief Operating Officer, Link11

Impact by Solution

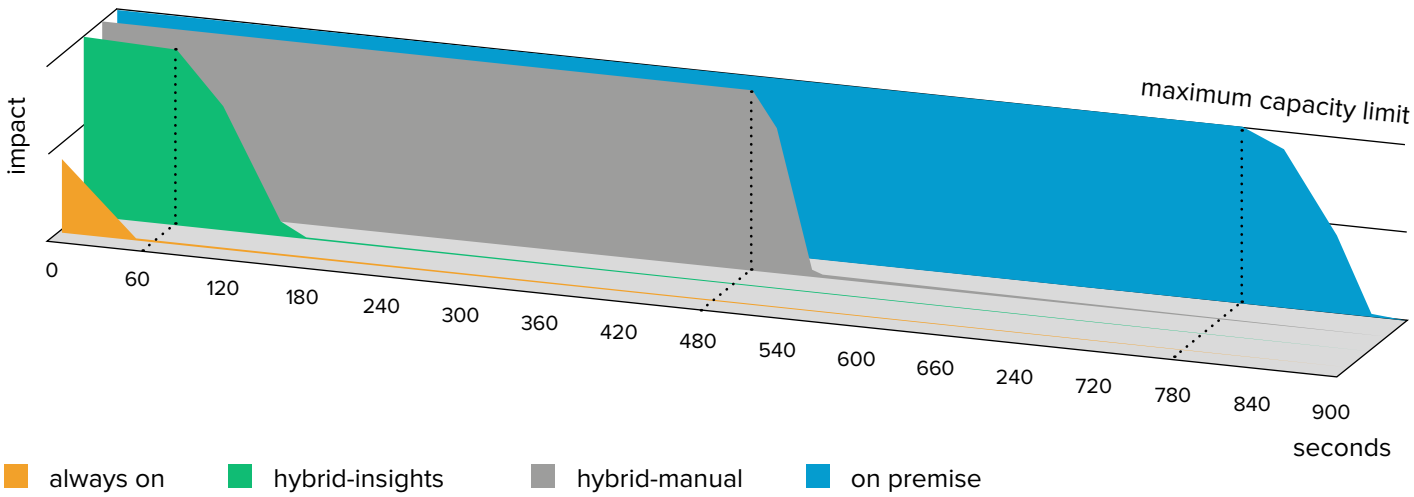

The top line on the y-axis represents the maximum capacity limit of the data traffic or bandwidths. If this is reached, there is an outage and potential collateral damage in the corporate network. The TTM varies greatly between the different modes:

a) Always on: The maximum capacity limit is never reached, and mitigation occurs within seconds.
b) Hybrid Insights: Automatic BGP rerouting can still take 45-60 seconds for the global route tables to converge. A temporary block occurs until traffic is fully redirected to the Link11 platform.
c) Hybrid Manual: It may take another 7 minutes or so to identify the destination IP and make the necessary BGP changes for the enterprise. The duration may be even longer if the enterprise network does not operate a 24/7 IT network operations center and must rely on on-call personnel.
d) On-site solutions: Local devices would be completely at the mercy of an attack for the entire duration. Until the carrier has been notified of the attack and has implemented what is known as „blackholing“ at the various upstream providers, all services and applications are unavailable.

The costs of an outage caused by a DDoS attack include lost revenue, service restoration and additional operational overhead costs. These losses have been estimated to average $5,600 per minute of downtime[4]. Based on these estimates, an on-site-only strategy would cost $84,000 for the assumed 15-minute DDoS attack. DDoS attacks are becoming larger, more sustained, and more complex. These costs will continue to rise in the event of an attack and the associated downtime without an adequate and efficient solution.

To learn more, talk to the cyber resilience experts at Link11. Link11 relies on artificial intelligence, machine learning, automation, and real-time defence with its Insights feature for holistic protection of IT infrastructures and critical applications and to strengthen cyber resilience. We look forward to advising you on the topic and highlighting alternatives.

Your contact for hybrid protection solutions and their benefits:

Jag Bains
Link11 GmbH,
Vice President Solution Engineering
j.bains@link11.com

Sources

1 https://www.link11.com/en/downloads/ddos-report-1st-half-2022

2 https://www.darkreading.com/vulnerabilities-threats/link11-sets-new-standards-in-ddos-protection-as-test-winner

3 https://www2.cs.arizona.edu/~bzhang/paper/04-globecom-destreach.pdf

4 https://bigdata-madesimple.com/ddos-attacks-on-the-rise-a-closer-look-at-the-data

LINK11

www.link11.com

www.linkedin.com/company/link11/

twitter.com/link11gmbh

info@link11.com

Contact

Link11 GmbH
Lindleystr. 12
60314 Frankfurt