Product Portfolio

# Inter.link DDoS Protection

## Detection and Mitigation

DDoS attacks are the most relevant attacks (by number, price, and frequency) against IT infrastructure.

In the past years volumetric attacks have been growing from a few Gbps to multiple Tbps. Additionally, targeted and expensive application level attacks are harder than ever to mitigate.

DDoS detection and mitigation products protect customers from these threats by detecting, filtering, and blocking malicious traffic.

In general, one must differentiate between protection on Layer3-4 (IPs, Protocols & Ports) and Layer3-7 which includes application-specific payloads (ex. HTTP, DNS & more).

## What We Bring to the Table

We use our own backbone and multiple scrubbing centers to offer protection for everything from the smallest ISPs to large Enterprises.

**Clear Billing -** We don't base billing on "clean traffic" or "number of attacks" which you can't predict or influence. Instead you only pay for the features you need.

**Free Options -** Our basic DDoS protection is included in all IP-based products for free.

**Automated Services -** All of our tiers can be implemented easily and will be included in our automated service portal.

**Variety of Tiers -** We offer Plus, Premium, and Enterprise tiers to fit everyone's needs. This allows us to offer everything from a basic insurance-like option all the way to always-on Enterprise protection.

| Features | Basic | Plus | Premium | Enterprise |
|---|---|---|---|---|
| **Protection Layer** | Layer3-4 | Layer3-7 | Layer3-7 | Layer3-7 |
| **Max Attack Bandwidth** | 100 Gbps | 1 Tbps | 2 Tbps | 5 Tbps |
| **Scrubbing Capacity** | None | 100 Gbps | 250 Gbps | 500 Gbps |
| **Scubbing Center** | None | 1 Region | 2 Regions | 3 Regions |
| **Protected /24 & /48** | 2 | 4 | 10 | 20 |
| **Native IPv6 Support** | X | X | X | X |
| **Best Practice Filter** | X | X | X | X |
| **Attack Alerting** | | X | X | X |
| **Dashboard & Portal** | | X | X | X |
| **GRE Tunnel Support** | | | 1 | 2 |
| **Custom Filter** | | | X | X |
| **White- & Blacklists** | | | X | X |
| **FlowSpec Routes** | | | 10 | 25 |
| **Pulling of Traffic** | | | | X |
| **Enterprise Reports** | | | | X |
| **Adaptive Mitigation** | Charged per pricelist | 1h/month free | 3h/month free | Unlimited |

## Benefits Across All Tiers

- Pricing not based on number of attacks or clean traffic bandwidth.
- Integrated dashboard with automated reports.
- 24/7 proactive mitigation in German and English.
- An attractive reseller/partner model.
- By creating more specifics, we can pull the traffic via AS5405.
- Mitigated vectors:
  - TCP SYN, SYN-ACK, PUSH, RST, and FIN Flood.
  - UDP, DNS, HTTP, and ICMP Flood.
  - Session and Fragmentation Attacks
  - Protocol Violations, Faulty Applications

## Support

In general, the Inter.link SLA applies to all services. Contact addresses, reaction times, and our escalation matrix can be found on our website under https://inter.link/sla.

**Want to find out more?**

Reach out to us: sales@inter.link

Visit us on the web: Inter.link

Or call us: +49.30.577 03 74 - 10

# DDoS Reseller Program

The Inter.link DDoS Protection is also available for resellers. It offers a way to easily extend the products you offer your customers with the protection and trust of the Inter.link brand.

In order to participate each Reseller must sign the Inter.link's Reseller Agreement and must have at least one Inter.link IP Connectivity Port connected to their network. For the purpose of reselling DDoS Protection, discounted IP Ports are available to Resellers. The Reseller agrees to resell one tier per customer. Capacities from one tier which are dedicated to a specific customer can not be shared with a different customer or entity.

## Responsibilities

The Reseller is always responsible for actions taken and configuration changes for their customers. Inter.link bears no responsibility for those changes. The reseller is also responsible to provide support to their customers. Inter.link will only handle support requests directly with reseller personnel but not with the reseller's end customers.

## Fair-Use Policy & SLA

Inter.link is committed to keep every customer online and IP addresses protected. Therefore, all resellers and customers must adhere to our Fair-Use Policy. The users must not use Inter.link services with unsound intentions, carelessly, and/or putting Inter.link and its customers needlessly at risk.

However, every DDoS protection tier comes with its maximum capacity and limitation. If a customer or an attack exceeds the contractually agreed limits, Inter.link reserves the right to rate-limit or blackhole one or more IP address (es) as agreed in the SLA.

## Onboarding New Customers

Every new reseller customer needs to be reported to sales@inter.link. We'll track new customers automatically but for reporting and statistics, an e-mail is required.

## Offboarding Customers

End customer contracts must be cancelled by the reseller. The reseller will inform Inter.link in writing or will self-cancel the service through the Inter.link portal. The end customer cannot cancel their services directly with Inter.link since there is no direct contractual relationship.

## Billing

Inter.link will invoice the resellers for all consumed services for themselves and their customers. The resellers must invoice their end customers directly. The reseller must not create or include end customer IP address space under their own account. End customer IP space is defined as IP space not directly or contractually belonging to the reseller.

## Pricing

There are three different tiers of DDoS Protection which can be purchased with different associated costs.

The "Basic" Tier is included in all IP Services provided to customers.

The Reseller must sell one tier per end customer, aggregation/pooling of customers is not allowed. The reseller handles first-line support and all customer communication. The reseller is the contract holder.

## Additional Options

**Protected Networks:** Protected networks are included in all Tiers starting from Plus. Protected networks are those configured within the DDoS protection and Traffic is routed to the DDoS protection.

**GRE Tunnels:** Additional GRE Tunnels cost 250 Euros monthly per tunnel.

**Additional FlowSpec Routes:** See IP-Transit FlowSpec Route packages.

**Want to find out more?**

Reach out to us: **sales@inter.link**

Visit us on the web: Inter.link

Or call us: +49.30.577 03 74 - 10

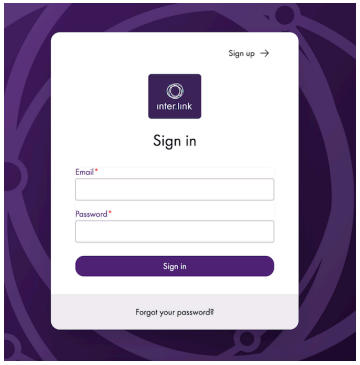


# Mitigation Portal

The Inter.link Portal provides a one-stop location to order and manage network services.

Mitigations, Dashboards, and Reports can be managed through our portal.

- Segment customers through groups
- Assign Zones (protected objects to groups)
- Create Mitigation-Policies and re-use them with templates to be assigned to one or many Individual Zones or Groups.
- Granular User-Management applicable on group-levels
- Configure notifications globally (or for each tenant/group)
- Configure reporting (group level)

## General Definitions

### Route Announcements

/24 or bigger prefix over transit. Onboarding of smaller subnets with BGP-community possible.

### Group

A group is a collection of zones. A group normally reflects a customer (tenant), a big project or a special part of the infrastructure (for example managed by a specific department[s]). Things like notifications and reporting can only be configured on a group level (and not zones).

### Zones

A zone is a collection of prefixes (aka protected objects). Usually, a zone represents a specific part of the infrastructure. A zone can contain a single IP (IPv4/32 and/or IPv6/128) but also a list of all (big) aggregates/prefixes within the network. Note that IPv4 and IPv6 can be mixed within the same zone. Most importantly the zone decides how the traffic is treated. A zone has a policy attached to it describing how to deal with traffic and attacks toward its protected objects.

# General Definitions

## Policy

A set of predefined rules, how traffic is classified and identified as attack-traffic. This also included the action being applied to suspected attack traffic from dropping single packets to blocking malicious IP for predefined time-ranges. Easily applicable to multiple Zones at once without further configuration involved if suited. Each policy and vector/pattern can be set to either "Monitoring" or "Mitigate" mode.

## Policy Modes

### "Monitoring"

In "Monitoring" mode the scrubbing centers will preform all actions (Analyze, Detect, Graph, and Report) except for actually dropping any traffic. This mode is best suited to perform pre-analyze the zone traffic and do a "soft onboard" to specific protected objects/prefixes. This mode can quickly be switches to "Mitigate" (by Customers/Resellers of Inter.link).

### "Mitigate"

If a policy is in "Mitigate" mode it will be fully automatic in terms of traffic- and attack response. The platform will start dropping attack traffic as soon as the attack is detected (depending on the vector, pattern and distribution 1 to 60 seconds is the normal response time).

## Reporting

Reports can be automatically generated and delivered via email on a regular interval, which represent the past history of your zone traffic and attacks including:

- Inbound Traffic - Packet Rate per hour (pps)
- Inbound Traffic - Bit Rate per hour (Bandwidth)
- Protocol Distribution in % (TCP, UDP, ICMP, Other)
- Attack Breakdown by Attack Type (Attack Count, Attack Packets, Attack Bits)
- Top 10 Attacks (Sort by bytes)
  - Attacked IP
  - Total Attack Bits / Dropped Attack Bits
  - Start and End time

The intervals/time ranges are daily, weekly, and monthly.

# General Definitions

## Notification

In case of a detected attack, a notification email will be sent immediately out to predefined recipients. This Notification contains a brief overview of the nature of the attack including, but not limited to Unique-ID, IP-/ Zone-/Groups-, Attack Vector, Attack rate (pps/bps).

## Onboarding

The onboarding consists of two phases/stages. Self On/Off Boarding of protected objects (Zones) via BGP-Community.

1. Traffic Redirect through the Scrubbing Centers

To be able to detect and mitigate traffic we need to send the traffic through our scrubbing centers. Customers with an ASN can use a BGP community to redirect traffic through the centers. For “static/access” customers this prefix onboarding must be performed via an Inter.link ticket. In both cases, it is also planned to offer onboarding through the mitigation portal.

2. Configure a Zone (Prefix + Policy[Action])

A zone needs to be configured with the protected object (prefix) and its attached policy. This will tell our platform how to deal with traffic towards the protected object(s).

**Want to find out more?**

Reach out to us: **sales@inter.link**

Visit us on the web: Inter.link

Or call us: +49.30.577 03 74 - 10

# Appendix 1: Capacities

With our DDoS Protection customers and resellers are being enabled to defend themselves against the following common attack vectors. In addition, multiple built in Layer3/4 filters are constantly active and monitored in our network (ex. Udp/0, memcached, uvm).

Below is a list of common attack vectors that our product is able to mitigate.

### SYN Flood

In a SYN flood, a victim server receives spoofed SYN requests at a high packet rate that contain fake source IP addresses. The SYN flood overwhelms the victim server by depleting its system resources (connection table memory) normally used to store and process these incoming packets, resulting in performance degradation or a complete server shutdown. A well-crafted SYN flood often fools deep-packet inspection filtering techniques.

### SYN-ACK Flood

Host servers generate SYN-ACK packets in response to incoming SYN requests from clients. During a SYN-ACK flood, the victim server receives spoofed SYN-ACK packets at a high packet rate. This flood exhausts a victim's server by depleting its system resources (memory, CPU, etc.) used to compute this irregularity, resulting in performance degradation or a complete server shutdown.

### ACK & PUSH ACK Flood

After a TCP-SYN session is established between a host and a client, ACK or PUSH ACK packets are used to communicate information back and forth between the two until the session is closed. During the ACK flood, a victim receives spoofed ACK packets at a high packet rate that fail to belong to any session within the server's connection list. The ACK flood exhausts a victim's server by depleting its system resources (memory, CPU, etc.) used to match these incoming packets, resulting in performance degradation or a complete server shutdown.

### Fragmented ACK

A variation of the ACK & PUSH ACK Flood. The attack uses 1500 byte size packets to consume large amounts of bandwidth, while generating a relatively moderate packet rate. Because routers do not reassemble fragmented packets at the IP level, these packets usually pass-through routers, ACL, firewalls, and IDS/IPS unimpeded. The packet content is usually randomized, irrelevant data. The attacker's goal is to consume all bandwidth of the victim's network.

## RST and FIN Flood

In order to close a TCP-SYN session between a client and a host, the servers exchange RST or FIN packets to close the session using a three-way or four-way TCP communication handshake. During a RST or FIN flood, a victim server receives spoofed RST or FIN packets at a high rate that do not belong to any session within the server's databases. The RST or FIN flood exhausts a victim's server by depleting its system resources (memory, CPU, etc.) used to match these incoming packets, resulting in performance degradation or a complete server shutdown.

## Synonymous IP

A victim receives spoofed TCP-SYN packets at a high rate that have the victim's information specified as both the Source IP and Destination IP. This attack exhausts a victim's server by depleting its system resources (memory, CPU, etc.) used to compute this irregularity, resulting in performance degradation or a complete server shutdown. Although the packet's Source and Destination IP are identically defined within the Synonymous IP attack, the content is irrelevant because the attacker is simply depleting the victim's system resources.

## Fake Session

This attack fakes a complete TCP communication and is designed to fool new defense tools that only monitor incoming traffic to the network. There are two variations of this attack: the first variation generates multiple forged SYNs, then multiple ACKs, followed by one or more FIN/RST packets, and the second variation skips the initial SYN, and starts by generating multiple ACKs, followed by one or more FIN/RST packets. The low TCP-SYN rate makes the attack harder to detect than a typical SYN flood while achieving the same result: the depletion of the victim's system resources.

## Session Attack

A valid TCP-SYN session is generated between a BOT and a victim. Once the session is established, the attacker delays responding with an ACK packet to keep the session open until a Session Time Out is triggered. The empty session exhausts the victim's server by depleting its system resources (memory, CPU, etc.) used to compute this irregularity, resulting in performance degradation or a complete server shutdown. Session Attacks are non-spoofed: the source IP is the actual public IP of the attacker BOT, and the source IP range is equal to the number of BOTs used in the attack.

## Misused Application Attack

The attacker does not use BOTs to consume the system resources of a victim's server. Rather, an attacker redirects valid clients belonging to a high traffic application, such as peer-to-peer services, to a victim's server. The target victim is then overwhelmed with traffic from a group of misdirected computers trying to form a legitimate connection with its server. The overwhelming connection requests received by the victim's server depletes its system resources, resulting in performance degradation of a complete server shutdown.

## HTTP Fragmentation

In this attack, the BOT (non-spoofed) establishes a valid HTTP connection with a web server. The BOT proceeds to fragment legitimate HTTP packets into tiny fragments, sending each fragment as slow as the server time out allows, holding up the HTTP connection for a long time without raising any alarms. For Apache and many other web servers designed with improper time-out mechanisms, this HTTP session time can be extended to a very long time period. By opening multiple extended sessions per BOT, the attacker can silently stop a web service with just a handful of BOTs.

## UDP Flood

During a UDP flood, a victim server receives spoofed UDP packets at a very high packet rate and with a large source IP range. The victim server is overwhelmed by the large number of incoming UDP packets. The attack consumes network resources and available bandwidth, exhausting the network until it shuts down. A full communication handshake is not used in the UDP software to exchange data, making UDP attacks difficult to detect and extremely effective in flooding the network bandwidth. UDP floods can overwhelm a network with packets containing randomized or fixed Source IP addresses and can be designed to target a specific server by using the victim's information as the Destination port and IP within the packets.

## UDP Fragmentation

A variation of the UDP flood. The attacker uses large packets (1500 bytes) to consume more bandwidth with fewer packets. Since these fragmented packets are forged and have no real relationship for reassembly, the victim server receiving these packets will spend CPU resources to "reassemble" useless packets. This often causes the processors to overload and sometimes reboot the entire system. This attack is harder to identify because it resembles good traffic.

## DNS Flood

An application-specific variation of the UDP flood. During a DNS flood, a victim DNS server receives valid but spoofed DNS request packets at a very high packet rate and from a very large pool of source IP. The victim server cannot determine which packet is from a real server and therefore proceeds to respond to all requests. The server is overwhelmed by the requests. This attack consumes network resources and available bandwidth that exhausts the network until it shuts down. Spoofed DNS attacks are well-crafted flood attacks - the content of spoofed DNS packets are designed to mimic actual DNS requests. Since they are 100% normal looking packets, this attack is not detectable by deep packet inspection. With a wide range of available attacking IP, the attacker can easily evade most traffic anomaly detection techniques.

## VoIP Flood

A variation of an application specific UDP flood. A victim VoIP server receives spoofed VoIP packets at a very high packet rate and with a very large source IP range. The victim server has to sort out the proper VoIP connections from the forged ones, consuming a detrimental amount of resources. VoIP floods can overwhelm a network with packets containing randomized or fixed Source IP addresses. A fixed Source VoIP attack mimics traffic from large VoIP servers, and can be very difficult to identify because it resembles good traffic.

## ICMP Flood

A victim server receives spoofed ICMP packets at a very high packet rate and with a very large source IP range. The victim server is overwhelmed by the large number of incoming ICMP packets. The attack consumes network resources and available bandwidth, exhausting the network until it shuts down. A full communication handshake is not used in the ICMP software stack to exchange data, making ICMP-based attacked difficult to detect. ICMP floods can overwhelm a network with packets containing randomized or fixed Source IP addresses. ICMP floods can target a specific server by using the victim's information as the Destination port and IP within the packets.

## ICMP Fragmentation

A victim server receives spoofed, large fragmented ICMP packets (1500 bytes) at a high incoming packet rate and these packets cannot be reassembled. The large packet size expands the bandwidth of an ICMP attack. In addition, it causes the victim CPU to waste resources when it attempts to reassemble useless packets. This attack will often cause victim servers to overload and reboot.

## Ping Flood

An application specific adaptation of ICMP flood. During a Ping flood, a victim server receives spoofed ping (IMCP echo requests) at a very high packet rate and from a very large source IP range. The victim server is overwhelmed by the large number of incoming Ping packets. The attack consumes network resources and available bandwidth, exhausting the network until it shuts down. The spoofed Source IP can be random or set as the address of the victim. Since the Ping requests are usually well formed and from a large number of source IP addresses, the Ping flood cannot be easily detected by either deep packet inspection or anomaly detection techniques.