SteelShield – Performance / Scalability

Nitrado is the leading provider in game servers, also offering cloud servers, bare metal hosting and scalable, dynamic application hosting. We have more than 15 years of professional experience in systems engineering and software development.

SteelShield is our advanced, proprietary DDoS mitigation solution. Due to the nature of our business – hosting gaming applications for hundreds of thousands of players – we're massively targeted by DDoS attacks, and we're experts at dealing with it.

This document outlines SteelShield's performance and scalability, as well as different deployment scenarios. For a high-level overview on SteelShield's features and capabilities, please refer to our SteelShield whitepaper ("Reinventing DDoS mitigation").

STEELSHIELD

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

Deployment Scenarios

SteelShield is not tied to any particular deployment scenario and can be integrated in a number of different ways.

Scrubbing center

We act as a scrubbing center - we announce your IP networks at our PoPs, either permanently or temporarily during an ongoing attack, and forward filtered traffic to your origin network via GRE tunnels. This is recommended for infrequent attacks.

Direct IP interconnect / transit

We act as transit provider for your network(s) and maintain direct IP interconnects with you at our points of presence. This eliminates the operational complexity and packet overhead associated with GRE tunnels. Recommended if you need permanent mitigation.

Bare Metal Hosting

Similar to IP transit, but you buy hosting services from us.

SteelShield is tightly integrated with our bare metal hosting offering and management tools. Refer to our "Better Bare Metal" whitepaper for more details on our hosting solutions.

Custom Deployment

We deploy SteelShield inside your infrastructure. This is the most flexible option for existing setups. SteelShield can be deployed both in virtualized and physical environments.

It can operate in two modes:

Inline mode ("bump in the wire")

SteelShield forwards packets without routing them, like a switch would. Requires a layer 1 or 2 failover setup like a hardware bypass device or a highly available setup with LACP trunking.

BGP offramp setup

In a BGP setup, SteelShield acts as a route injector and peers with your edge or core routers, announcing either regular or Flowspec routes. You can control announcements through either a gRPC or REST API. This is our preferred deployment scenario, since it allows us fine-grained control over traffic redirection, and an easy way to achieve high availability by withdrawing routes in case of failures. The exact setup depends on the router vendor (our reference design uses Juniper MX routers).

We can combine any of these scenarios depending on your requirements. Since we own the SteelShield software and network stack, we can support complex and custom network topologies.

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

Our network backbone

Our global network is protected by **Link11's award-winning filtering technology**, which mitigates large volumetric attacks, while SteelShield takes care of smaller volumetric/spoofing and application layer attacks. Since any delay introduced by on-demand mitigation would be unacceptable, our traffic is permanently routed through Link11's network.

Our points of presence for direct IP interconnection:

- One Wilshire Building (Coresite), Los Angeles
- Quadranet, 530 West 6th St., Los Angeles
- Quadranet, 6171 W. Century Blvd., Los Angeles
- TelX/DigitalRealty, 36 NE Second Street, Miami
- TelX/DigitalRealty, 60 Hudson Street, New York City
- Telehouse, 7 Teleport Drive, Staten Island, NY
- Telehouse, Kleyerstraße, Frankfurt am Main
- Equinix FRA5, Frankfurt am Main
- Equinix, Singapore (available Q3 2018)
- Equinix, Sydney (available Q4 2018)

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

Filtering Performance

SteelShield has very high sensitivity and specificity - it causes minimal false positives, and is very efficient at identifying malicious traffic, at a much better ratio than any competing solution.

We successfully filter thousands of individual attacks every month. While DDoS attacks are a rare occurrence for most networks, we're seeing attacks 24/7. In June 2018 alone, we detected and mitigated more than 1,800 individual attacks.

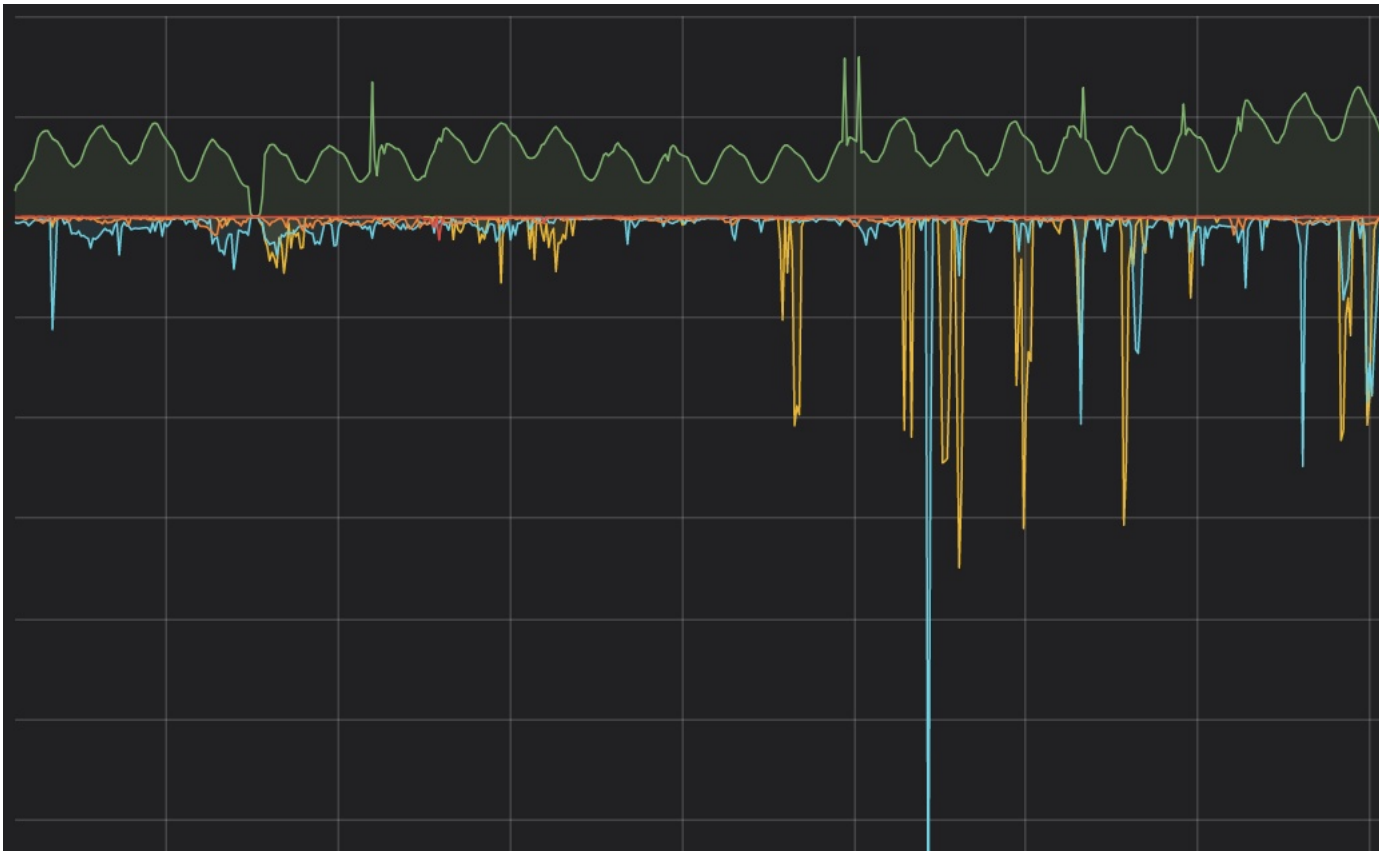

SteelShield statistics for June 2018 (after pre-filtering by ISP).

Clean traffic is at the top, and different kinds of mitigated attacks are at the bottom. The small irregularities seen in the clean traffic are caused by large-scale software deployments.

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

High-resolution mitigation overview (one-second interval). Clean traffic is unaffected by the attacks. Like in the previous graphs, one color represents one mitigation strategy.

Mitigation is instant - not a single packet made it through.

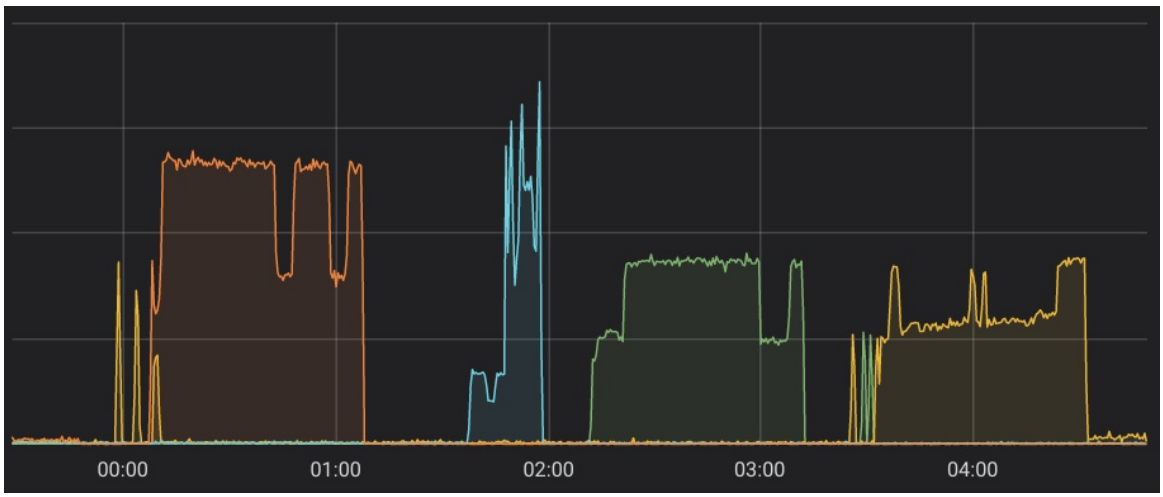

A different view, showing the largest attacks at the time. Each color represents one target host. In this case, an attacker probed four different hosts before finally giving up..

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

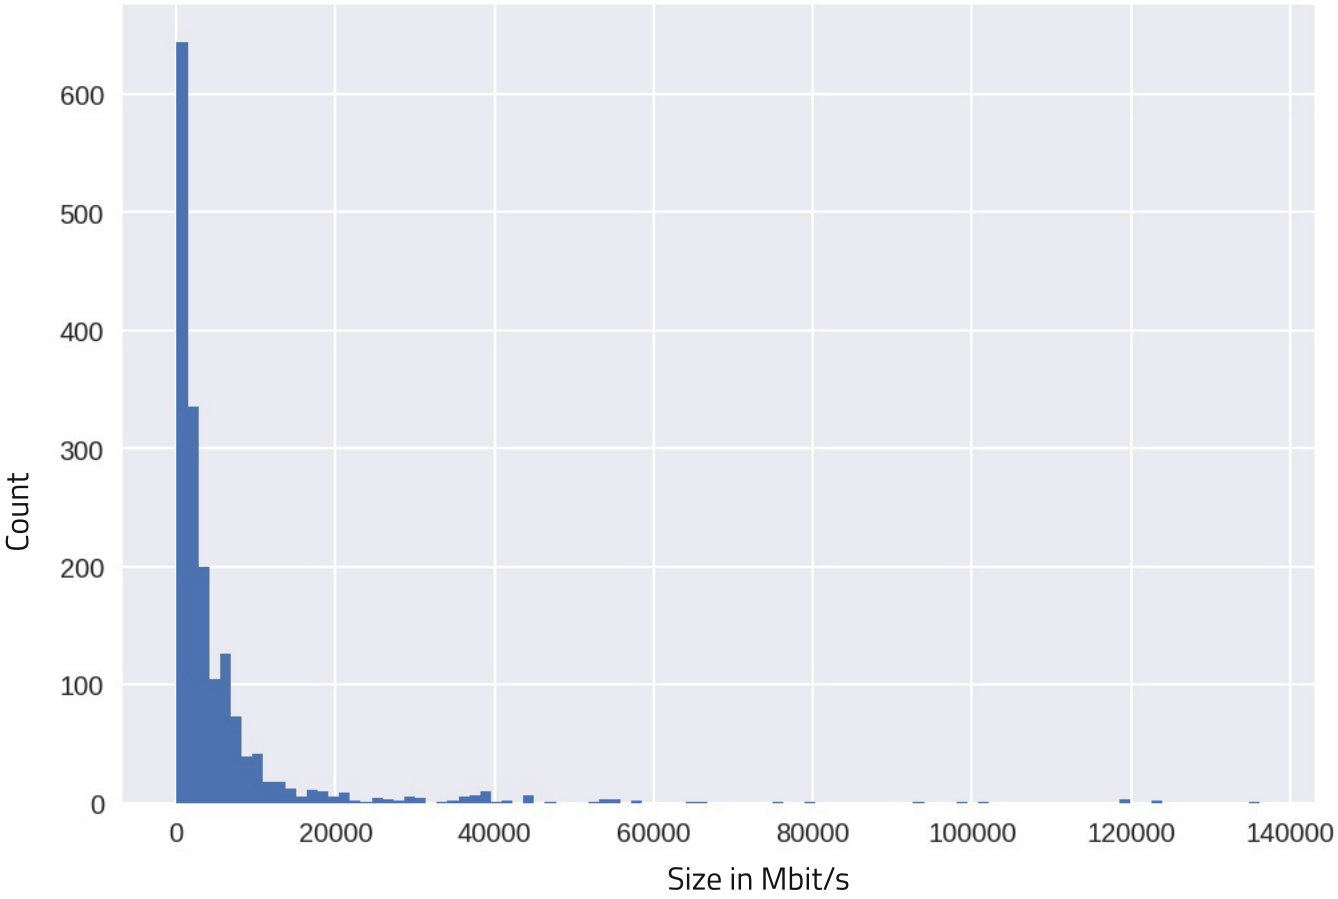

Histogram of DDoS attack bandwidths in June 2018, ranging from 0 to 140 Gbit/s

An interesting, 'spikey' application-layer attack which attempts to exploit mitigation delays common in other solutions. After realizing the attack's futility, the attacker switched to another, equally ineffective method.

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

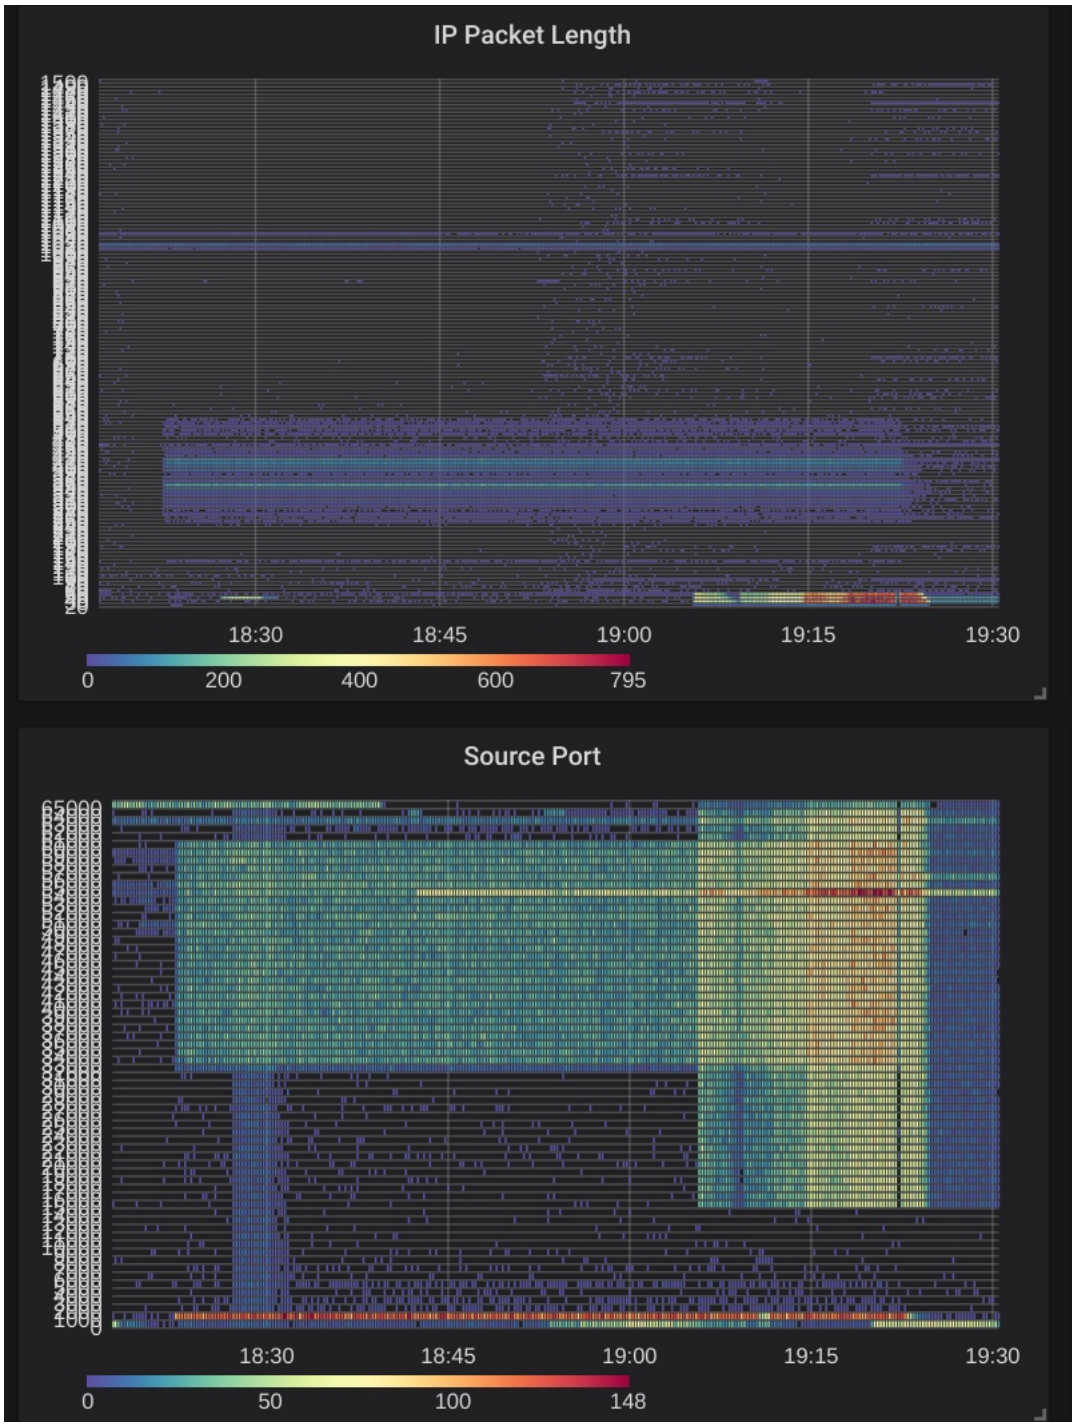

Our telemetry backend visualizing an application-layer attack where the attacker employs different strategies and source port randomization mechanisms. This is one of many highly detailed analytics tools which can be used to analyze attacks even days after they happened, with full details and filtering capabilities (no aggregation).

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net

Scalability

SteelShield is able to filter both "dumb" volumetric attacks and sophisticated (but smaller) application-level attacks. In our environment, it's more economical to offload easily-filtered volumetric attacks to our ISPs and handle only smaller application-layer attacks which our ISPs cannot filter.

However, SteelShield can be scaled to **handle any amount of traffic**, only limited by external factors (ingress bandwidth, network devices on the ingress path, physical space).

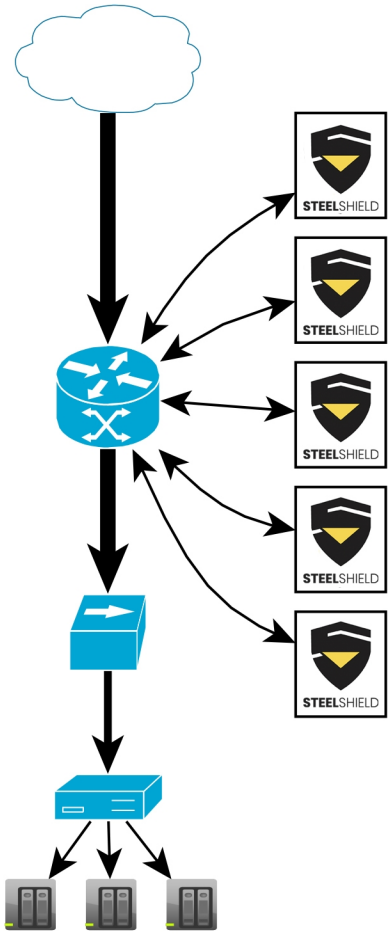

Most of our filter algorithms are completely **stateless**. For inherently stateful features like flow tracking, we use a highly available state synchronization mechanism.

This makes SteelShield exceptionally scalable – traffic can be distributed across a cluster of hosts using standard networking technologies like port trunking or equal-cost multi-path (ECMP) routing.

Each SteelShield cluster node is designed to **handle 20 Gbps of small-packet line-rate** traffic (28 million 64 byte packets per second). According to our current benchmarks, each node can actually handle up to 36 Gbps of traffic - but our design includes a large safety margin, which increases reliability and allows for future software upgrades that introduce algorithms which are more computationally intensive than our current ones.

64 byte packets are the worst-case scenario which we optimize for, but all large-volume DDoS attacks we've seen used the largest-possible packet size - 1,500 bytes - saturating network interfaces at much lower packet rates, which are easier to handle (there's a fixed per-packet overhead rather than per-byte). At packet sizes larger than 256 bytes, we easily achieve line rate at 40 Gbps.

Our current-generation hardware can be upgraded to support ~70 Gbps, but we aren't currently pursuing this since we're focusing on "scaling out" by increasing the number of filtering nodes, rather than operating each node at its theoretical maximum saturation.

For more in-depth information, please feel free to contact us anytime. We're happy to work directly with your engineers to discuss the technical requirements.

marbis GmbH
Griesbachstr. 10
D-76185 Karlsruhe
Germany
m.boessendoerfer@marbis.net

marbis / Nitrado USA, Inc.
4600 Kietzke Lane, H-188
Reno, Nevada 89502
USA
volker.huber@marbis.net