corero | OpenCape

Corero Powers Next-Generation DDoS Defense

DDoS Attacks – Increasingly Frequent and Complex

Contents

As Certain as Death and Taxes

Along with email spam, phishing, malware and crazy cat videos, DDoS attacks remain a persistent blight on the Internet. The Verizon 2018 Data Breach Investigations Report notes "that the degree of certainty that they will occur is almost in the same class as death and taxes". Cyber criminals employing automated methods for launching attacks have escalated what was once an occasional, but often severe, nuisance into a widespread, ever-present and worsening threat. According to the Corero Full Year 2018 DDOS Trends Report, the average number of attacks per customer was up 16% over the previous 12 months.

Financially motivated criminal organizations and nation state actors bent on cyber warfare have aligned with malicious hackers to pool their collective knowledge and experience for generating increasingly complex, multi-vector, attacks that are more difficult to detect and mitigate. The vast majority of DDoS attacks are either volumetric in nature – consuming a high percentage of network bandwidth – or focused on exhausting protocol-processing resources in the host systems under attack. Both types are highly effective in knocking out Internet applications and services, sometimes for hours and with severe consequences for service providers, businesses and consumers.

A Clear and Present Danger

DDoS attacks are often launched from large-scale botnets created by hijacking poorly secured endpoints, including servers, PCs, laptops and consumer IoT devices, such as webcams. According to Verizon, over 75% of attacks also utilize amplification techniques to jack up attack intensity by exploiting vulnerabilities in Internet services and host systems that can generate a flood of data targeting a victim just by sending a series of simple requests.

Yet Corero reports that low-volume, sub-saturating attacks continue to dominate, with 95% of all events observed in 2018 below 5 Gbps in intensity. While this may seem counter-intuitive, it makes perfect sense given there are finite resources that an increasing number of bad actors can exploit to disrupt a growing number of targets. The fact that attacks are becoming shorter in duration (with 82%, observed during 2018, lasting less than 10 minutes) is further validation of this trend.

Application layer attacks are on the rise but still represent a miniscule percentage of all DDoS attacks. An Akamai 2018 State of the Internet blog post reported that "Layer 3 and 4 attacks continued to account for the vast majority (99.1%) of the DDoS attacks seen by Akamai", indicating that the vast majority of cyber criminals are content to exploit network and transport layer vulnerabilities, as opposed to denying service at the application layer. Furthermore, cyber criminals intent on exploiting application vulnerabilities typically concentrate on stealthily breaching security systems to exfiltrate or manipulate data, rather than simply disrupting service.

An Ever-Expanding Attack Surface

The vast complex of public networks spanning the globe is constantly growing and evolving, reaching into every corner of society. New devices, hosts and networks come online every hour of the day, offering bad actors a constantly growing number of targets and potentially vulnerable endpoints that can be exploited to launch attacks. Cyber criminals continuously crawl the Internet looking for new devices, which they can discover and hack within minutes of going online.

Internet trends are pushing the DDoS battle onto two separate fronts:

1. An increasing number of rapidly growing, high-capacity, access networks and edge data centers supporting a growing number of next-generation IoT use cases.
2. Massive, hyperscale data centers for delivering multi-cloud applications and services.

IoT adoption is driving a proliferation of intelligent devices that will ultimately exceed the number of user endpoints. Machine-to-machine connections at the edge will power a wide range of IoT use cases, spanning real-time industrial process control, remote healthcare, environmental sensing and "smart" physical infrastructure – power, water, transportation, roads and buildings. Mission-critical IoT networks will no doubt become targets for potentially catastrophic DDoS attacks.

Bandwidth at the Internet edge continues to scale up. Gigabit consumer broadband is already here. Multi-gigabit 5G wireless networks will soon be a reality. More bandwidth at the edge is driving the need for more backbone capacity. Fat pipe connections are moving from 10Gbps to 100Gbps and beyond. A faster edge also enables higher intensity attacks, with fewer endpoints needed to generate traffic for crippling volumetric attacks.

Consumers already live in a multi-cloud world of mobile apps, streaming media, e-commerce, online gaming and social media. Business users have been enthusiastic adopters of SaaS applications and enterprise IT managers are migrating business critical applications to hybrid public/private multi-cloud environments. Hyperscale data centers, consisting of thousands of servers, are both targets for attack and potential platforms for launching attacks. Content distribution networks, for scaling cloud service delivery, are also targets for attacks that could cripple services affecting millions of users.

Enterprise multi-cloud adoption is occurring in parallel with a move away from private MPLS WANs to SD-WANs based on inexpensive and readily available Internet broadband connections. SD-WAN overlay networks are vulnerable to DDoS attacks and Internet breakout connections for SaaS applications are also potential points of attack.

A Never-Ending Battle

The Internet is so vast, complex and constantly changing that there is no known method for eradicating the sources of DDoS attacks. Network operators need to be ever-vigilant and prepared to deal with attacks when – not if – then occur.

DDoS defense is a never-ending series of battles that swing back and forth with the black hats and white hats alternately enjoying the advantage. Cyber criminals discover and exploit vulnerable hosts to launch attacks. Defenders monitor network activity to compile a catalog of attack profiles that are used to take the necessary mitigation actions. New attack vectors are added to the catalog and any associated host vulnerabilities are circulated within the security community so that network operators can take preventive measures.

While massive-scale, high-intensity, DDoS attacks measured in hundreds of gigabits get the headlines in the press, the war is waged in an endless series of smaller-scale skirmishes. High-intensity attacks may rise for a short period of time but then cyber criminals are forced to regroup as network operators mount defenses to reconfigure and protect vulnerable hosts, and so attack intensity subsides. Yet the cycle of DDoS attack and defense continues on, with no end in sight.

FIGURE 1

Never-Ending DDoS Attack/Defense Cycle

TRADITIONAL DDOS DEFENSE IS BECOMING PROBLEMATIC

Complex, Costly and Time-Consuming

While traditional DDoS defense solutions have proven a certain level of effectiveness in preventing widespread carnage on the Internet, these defenses remain complex, costly and time-consuming for network operators. With attacks increasing in frequency, intensity and technical sophistication, DDoS defense has grown more resource-intensive, requiring additional investment in the supporting infrastructure needed to detect and mitigate against attacks, increasing both CAPEX and OPEX.

A major shortcoming of traditional DDOS defense is the significant lag between the time an attack is first detected, to when the operator initiates the appropriate mitigation actions. The lower bound on this time is determined by the manual network operator and security analyst workflows that typically consume precious minutes, or even hours in the event of a highly complex attack. Meanwhile, during this interval, service is likely to be disrupted and users impacted by the attack.

Scrubbing Centers Dominate

Scrubbing centers dominate traditional approaches to DDoS defense. Once an attack has been detected and identified by its profile, operators typically reroute traffic flows through special-purpose data centers equipped with DPI-based network devices designed to filter and block malicious traffic flows. Mon-malicious traffic is passed through the scrubbing center and routed on to its final destination. However, the shortcomings of this approach are becoming evident as attacks increase in intensity and variability, with sudden spikes in bandwidth and constantly changing attack vectors that cycle through multiple modes and targets over time.

Service Provider Scrubbing Centers

Most large-scale network service providers operate their own scrubbing centers for detecting and mitigating attacks that originate, transit or terminate in their networks. When an attack is detected, the operator reconfigures its network routers to redirect specific traffic flows through the scrubbing center.

Two significant drawbacks of this approach are the additional bandwidth consumed by backhauling the attack traffic across the core network, to the scrubbing center and the risk of congestion or overload at the scrubbing center. This has the negative effect of introducing additional latency and risk of loss for the legitimate traffic that is also passed through the scrubbing center.

In addition, as the intensity and frequency of attacks increases, the service provider needs to keep up by provisioning additional scrubbing center capacity, increasing CAPEX. Plus, there is always the risk that a very large-scale, high-intensity, attack either exceeds scrubbing center capacity or the bandwidth available on connections into the scrubbing center.

Finally, as Internet bandwidth continues to increase at the end-user edge, the prospect of malicious traffic originating on gigabit access connections could well obsolete scrubbing centers as a viable way to mitigate attacks.

Cloud-Based Scrubbing is Complicated by Multi-Vector Attacks

Many large enterprises rely on traffic scrubbing managed by a third-party DDoS protection service provider. In the event of an attack on their customers, the cloud-based DDoS service redirects and scrubs all traffic to ensure that only non-malicious traffic reaches the customer.

Cloud-based DDoS protection appeals to organisations who would prefer not to manage their own DDoS infrastructure. Cloud scrubbing centers are shared resources for mitigating attacks across a large pool of customers. The model works because at any given time the vast majority of customers are not under attack. However, there are complications that arise from managing access to shared scrubbing center capacity in the event of complex attacks with widely varying intensity.

First, a cloud scrubbing center needs to be provisioned with sufficient capacity to process the aggregate bandwidth of all traffic flows for customers under attack. Individual connections into the data center must also have sufficient bandwidth to transport the traffic to be scrubbed, independent of total capacity. If there is insufficient bandwidth, or scrubbing capacity, then malicious traffic won't be blocked and non-malicious flows may be impacted due by severe packet loss.

The second problem is allocating scrubbing center capacity. A troubling trend is attack vectors that vary over time, sending relatively short bursts of traffic at intensity levels which may jump rapidly from under 10 Gbps to over 100 Gbps. The DDoS provider has to manage the allocation of shared scrubbing capacity and the bandwidth available to transport traffic into the scrubbing center. This is operationally complex during highly dynamic attacks with intensity levels and vectors that constantly vary over a short period of time, possibly targeting multiple customers.

The necessary bandwidth or scrubbing center capacity can instantly become inadequate if there is a sudden spike in intensity or too many attacks occur at the same time. Provided there is sufficient bandwidth and capacity available, scrubbing center operators might eventually be able to successfully mitigate these types of attacks. However, rerouting traffic and dynamically allocating capacity consumes precious time and is far from an exact science.

Security Analyst and Network Operator Workflows Gate Time-To-Mitigation

DDoS detection involves monitoring network traffic and extracting flow metadata to derive attack vector signatures that are matched against a catalog of known profiles. Monitoring can be performed in-line or out-of-band to extract metadata from flows, which is correlated with other network monitoring data to fully characterize attacks so they can be mitigated.

The downside of traditional DDoS defense solutions is that security analysts are integrally involved in the analysis of monitored traffic. Network operators, and cloud-based DDoS providers, staff teams of security analysts who swing into action when an attack is detected to analyze data, identify the type of attack and then initiate workflows to take the appropriate mitigation actions. Security analysts then coordinate with network operators to initiate mitigation, which typically involves rerouting specific traffic flows through a scrubbing center.

While workflow automation tools can help speed the process of detecting and mitigating attacks, operational workflows that put humans in the critical path result in time-to-mitigation measured in minutes, not seconds. A straightforward attack might take 5 minutes to mitigate, while a more complex attack could take 20-30 minutes. A tricky attack could take hours to sort out, during which time service may be severely impacted, affecting many thousands or even millions of users.

BGP Router-Based DDoS Mitigation

Another aspect of traditional DDoS defense involves service providers reconfiguring routers at network ingress points to mitigate attacks. Network operators can use BGP to "black hole" or block traffic based on destination or source IP address. While effective for blocking attacks, black holing, based on destination IP, sacrifices that customer for the benefit of everyone else, because that endpoint will no longer receive any traffic, whether or not it is legitimate.

Operators are also starting to utilize BGP FlowSpec, which is a protocol for reconfiguring Internet routers so that flows can be filtered and blocked based on the content of the classic 5-tuple in IP packet headers. However, BGP FlowSpec is currently a fairly coarse-grained mechanism for blocking flows because filters can't be applied based on visibility into packet payload, which excludes filters based on encapsulated protocol headers and application layer content.

TECHNOLOGY ADVANCES ENABLING NETWORK-BASED DDOS DEFENSE

Next-Generation Networking Technologies

BGP router-based DDoS mitigation is only the first step on the path to network-based DDoS defense. Next-generation networking technologies, including programmable silicon, software-defined networking (SDN) and Big Data analytics will power new network-based DDoS defense solutions that are dramatically more effective and less costly to implement. And, it's important to note that, the first iterations of these enabling technologies have already been deployed by many ISPs, as part of their existing Internet routing infrastructure.

DDoS Mitigation Leverages Powerful Silicon and SDN

Powerful ASICs from equipment manufacturers and merchant silicon vendors already support flexible methods for filtering packet flows in-line at the silicon level. Routers based on these ASICs can be reconfigured on-the-fly with flexible rules to filter and drop packets based on data examined in both the packet header and payload. This allows malicious flows to be blocked using criteria for matching complex attack vector signatures.

Network operators can use standard SDN protocols such as NETCONF to automatically distribute these complex DDoS mitigation filtering rules to Internet routers, from a centrally located controller. When a new attack is detected, and its profile is identified, operators can immediately push out new filtering rules to mitigate the attack and reconfigure router interfaces at the relevant ingress points to start dropping the offending packets.

Network operators also require real-time confirmation that router-based filtering rules are actually blocking the intended traffic flows. The new generation of routing ASICs can track per-filter metrics, that the router's operating system can export to a centralized monitoring platform, using a streaming telemetry protocol such as gRPC. This provides operators with network-wide visibility into which router-based filters are proving effective in mitigating a DDoS attack.

Sampled DPI for DDoS Detection

Another key capability in the latest routing ASICs, is the ability to extract packet header and payload data based on a specified offset and number of bytes. Because these ASICs operate at terabit rates, packet flows can be continuously sampled, with raw "sampled DPI" data exported via streaming telemetry. Analysis of sampled metadata has proven to be 99.9% accurate for identifying malicious DDoS traffic and sampling overcomes the challenge of maintaining visibility as Internet access, backbone and cloud connections continue to scale up.

Real-Time Big Data Analytics

Real-time, Big Data, analytics is critical for network-based DDoS detection. A key component of the latest generation of network-based DDoS defense is a Big Data analytics engine that continuously ingests sampled DPI metadata, performance metrics and event data from routers and other network elements. Operators can leverage the analytics engine to perform real-time, multi-dimensional analysis of data collected to detect attacks, identify sources and characterize attack profiles. The analytics engine catalogs these profiles and generates packet filtering rules corresponding to each attack's specific signature. An SDN controller can then automatically distribute these filtering rules to reconfigure the relevant router ingress interfaces on-the-fly and block offending flows.

Real-Time, Automated DDoS Defense

Software can now automate network-based DDoS defense, so that the process of attack detection and mitigation takes place within seconds – not minutes or hours. An automated system can immediately detect malicious packets flows, formulate mitigation filtering rules and reconfigure routers with these rules, on the fly, without any intervention from security analysts or network operators.

Automated network-based DDoS defense enables DDoS attacks to be mitigated in real-time, as soon as the first flood of malicious packets start traversing the network. Automation removes security analysts and network operators from DDoS defense workflows, for all but the tiny proportion of attacks that are too complex to fully characterize.

CORERO SMARTWALL TDS DDOS SOLUTION

Real-Time, Line-Rate DDoS Defense

Corero Network Security has combined its software innovation with advances in Intel x86 multicore CPU technology, DPDK packet processing acceleration and high-performance NICs to develop a new generation of appliances providing breakthrough price/performance for DDoS defense. The Corero SmartWall TDS family of products provide line-rate protection at the network edge at connection speeds up to 100 Gbps.

SmartWall TDS appliances perform line-rate DPI to generate security metadata from traffic flows. The internal rules-engine examines this metadata to flag offending packet flows in real-time and instantly block attack packets. At the same time, the security metadata is streamed to the Corero SecureWatch Analytics platform, where further analysis involving correlation with other performance metrics and event data enables rapid identification of new attack vectors.

SecureWatch Analytics also formulates new mitigation rules for these vectors that are automatically distributed out to each SmartWall TDS appliance.

Corero SecureWatch Analytics is based on Splunk's Big Data analytics engine and is a critical component of Corero's front-line defense against DDoS attacks. SecureWatch Analytics features a web portal providing easy-to-read dashboards for monitoring routine operations and incident response. Operators can also perform complex queries to conduct sophisticated security forensic analysis.

Corero's DDoS defense solution has proven more than 99% effective in automatically detecting and mitigating attacks within seconds. This degree of effectiveness, speed and accuracy would not be possible without Corero's innovation and incorporating a Big Data analytics engine that can perform rapid analysis of high velocity security metadata.

Corero's SmartWall solution is fully automatic, detecting and mitigating attacks without the intervention of security analysts or network operators. Customers are usually unaware they have been under attack, until they check the SecureWatch Analytics dashboard for alerts.

FIGURE 2

Corero SmartWall TDS DDoS Protection

CORERO & JUNIPER NETWORKS — IN THE VANGUARD OF NETWORK-BASED DDOS DEFENSE

Corero and Juniper Networks Global Partnership

In 2018, Juniper Networks® signed a global partnership to sell Corero's SmartWall® DDoS products in conjunction with Juniper's MX Series routers. Juniper and Corero have developed an integrated solution for network-based DDoS defense that leverages powerful capabilities in the latest generation of MX router ASICs.

Silicon-Based Packet Filtering

MX Series router ASICs can be configured with flexible filtering rules that match on fields within a packet's header and payload. These filters are applied to every packet processed by the ASIC, to block malicious traffic flows at the silicon layer. Juniper utilizes software-defined networking protocols to distribute filter rules defined by Corero's SmartWall software to MX routers. When SmartWall identifies new attack vectors, and the corresponding filtering rules are distributed to MX routers, Juniper's Junos OS dynamically configures the new mitigation rules on-the-fly.

Silicon-Based Packet Data Export and Flow Telemetry

In the Corero-powered Juniper Networks solution, MX router ASICs are configured to sample packet header and payload data, which is exported using the Junos OS streaming telemetry protocol. Corero's SmartWall Virtual Detection Engine converts this sampled DPI data into security metadata that is examined to determine if packet flows match attack vector signatures. When a malicious flow is detected, the detection engine relays the mitigation rules to Juniper's software which pushes new filters out to MX routers using the NETCONF protocol.

Security metadata is also ingested by SecureWatch Analytics and correlated with performance metrics and event data to identify new attack signatures not yet known by the detection engine. When a new attack vector is discovered, the analytics engine formulates the corresponding mitigation rules, which are then relayed to Juniper's software to reconfigure the MX routers accordingly.

Security analysts and network operators need confirmation that MX routers are configured properly to block flows and mitigate attacks. To ensure this, MX router ASICs track per-flow metrics that are exported via streaming telemetry to monitoring systems, providing fine-grained visibility into the specific flows that are being filtered and dropped at MX router interfaces across the operator's entire network.

FIGURE 3
Corero & Juniper Networks DDoS Defense Solution

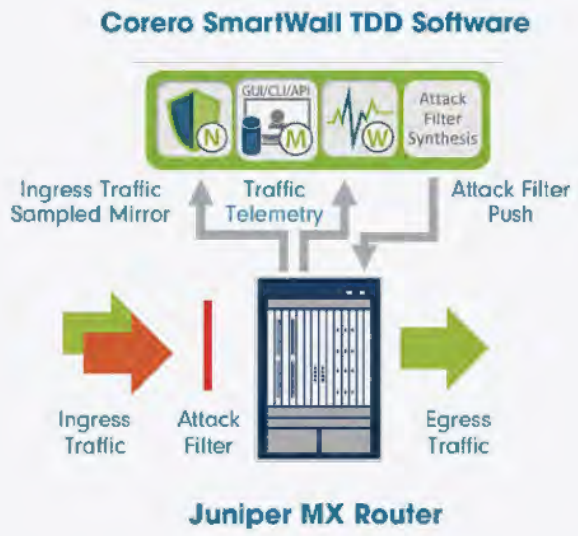

FUTURE DIRECTIONS IN NETWORK-BASED DDOS DEFENSE

Shift from Bolt-On to Built-In DDoS Defense

Internet evolution trends and advanced networking technologies are shifting DDoS defense from the traditional "bolt-on" scrubbing center approach to "built-in" network-based solutions that will further leverage technology breakthroughs in programmable silicon, SDN and machine intelligence powered by Big Data analytics

Software Control and Programmability at Every Layer

SDN enables disaggregation in network infrastructure, decomposing monolithic switches and routers into separate layers for the control and data planes. P4 programming is the extension of SDN down to the silicon layer, disaggregating firmware from the chip itself, enabling silicon to be easily adapted for specialized packet processing functions with a new firmware load.

P4 Runtime is an open, extensible API for dynamically reconfiguring network elements with new "match" criteria for the P4 pipeline, so that network elements to be updated with specific rules for processing packets, such as mitigation filtering rules for DDoS defense.

SDN protocols such as NETCONF provide mechanisms for automated, on-the-fly network reconfiguration for distributing DDoS mitigation rules. Streaming telemetry protocols such as gRPC can be used to export sampled DPI data and per-flow metrics for DDoS packet filtering.

At the routing layer, there are two IETF drafts that propose extending BGP FlowSpec to support payload matching and for applying rules on a specific set of router interfaces. BGP FlowSpec with payload matching could be combined with P4 programmable silicon to streamline the automated distribution of mitigation filtering rules to Internet routers.

Breakthroughs in Programmable Silicon

Router manufacturers, Juniper Networks and Nokia, have developed powerful custom ASICs that can be configured on-the-fly with flexible packet filtering rules to mitigate DDoS attacks. These ASICs can also export sampled DPI data and per-flow performance metrics. Broadcom's Jericho II merchant silicon ASIC supports comparable functions in a new generation chip for hyperscale data center spine switches.

Barefoot Networks' Tofino ASIC is fully programmable at the silicon layer and was specifically designed to implement packet flow processing functions using the P4 language (include link to P4 Language Consortium: p4.org). P4 was conceived to enable field-reconfigurable packet forwarding that is protocol and target independent, allowing switches to be configured at boot time with specific capabilities suited to a particular application or use case.

P4 programming specifies a "match-action" pipeline that defines the packet forwarding process for a network element, including "match" rules based on flexible filter definitions and "actions" that include extracting DPI data, dropping packets and tracking per-flow metrics.

Machine Intelligence Powered by Big Data Analytics

Real-time Big Data analytics has proven to be a key enabler for DDoS defense. Big Data engines provide the massively scalable compute and storage capacity required to ingest and process a huge volume of critical network telemetry such as flow metadata and sampled DPI to generate actionable security intelligence.

Big Data analytics also provides the foundation for applying machine learning and AI techniques that will further speed up and improve the effectiveness of DDoS defense solutions. These techniques will likely prove helpful in thwarting future complex, multi-vector, attacks that still require the intervention of security analysts.

NETWORK-BASED DEFENSE SPANNING THE EXPANDING INTERNET

Access, Backbone and Multi-Cloud Protection

Corero provides real-time, line-rate DDoS protection for customers including shared hosting facilities, large enterprises, government agencies, online gaming services, critical Internet infrastructure and cloud-native digital enterprises such as SaaS providers.

ISPs can deploy SmartWall TDS appliances always-on at the edge, or in scrubbing centers for on-demand mitigation. Service providers can also deploy Corero as part of a managed security service, benefiting from Corero's ability to automatically detect and mitigate DDoS attacks while leveraging a platform that leads the industry in price/performance.

The Corero – Juniper Networks Threat Defense Director network-based DDoS defense solution is particularly well-suited for mitigating high intensity attacks in large-scale ISP networks featuring high speed links (10 Gbps to 100 Gbps) and many routers distributed across the network edge.

Continued Internet expansion at the access and multi-cloud edges will ultimately drive a complete shift away from bolt-on DDoS defense to network-based solutions enabled by advanced technologies embedded in the Internet infrastructure.

FIGURE 4

Network-Based DDoS Defense Across the Expanding Internet

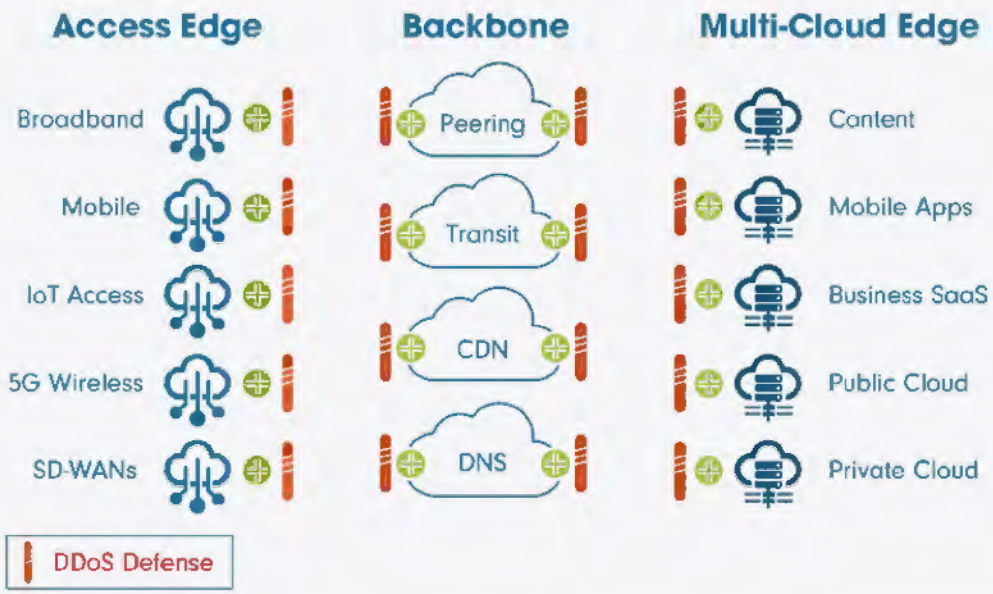

corero | OpenCape

Corero Network Security is a leader in real-time, high-performance DDoS defense solutions. Service providers, hosting providers and digital enterprises rely on Corero's award winning technology to eliminate the DDoS threat to their environment through automatic attack detection and mitigation, with comprehensive visibility, analytics and reporting. This industry leading technology delivers flexible protection that scales to tens of terabits, with a dramatically lower cost of ownership than previously possible.

OpenCape Corporation is a 501c3 nonprofit technology company headquartered in Barnstable Village at the Barnstable County Complex. OpenCape owns and operates a state-of-the-art fiber optic network built to serve local governments, businesses, and residents of Southeastern Massachusetts, the Cape & Islands. Our fiber network is on par with the most sophisticated and technically proficient fiber networks in the world. OpenCape Corporation also sells a variety of Internet and Voice-Over-Internet services as part of our continued focus on advancing the needs and interests of the communities we serve.

US HEADQUARTERS

Corero Network Security Inc.
293 Boston Post Road West, Suite 310
Marlborough, MA 01752
Tel: +1 978 212 1500
Email: info@corero.com

EMEA HEADQUARTERS

Corero Network Security (UK) Ltd.
St Mary's Court, The Broadway,
Amersham, Buckinghamshire, HP7 0UT, UK
Tel: +44 (0) 1494 590404
Email: info_uk@corero.com