NEXUSGUARD

White Paper

DDoS Mitigation

White Paper

V2.1.1

Content

Introduction 3

Common types of DDoS attacks 4

- Volumetric attacks 4
- Protocol attacks 4
- Application layer attacks 4

Impact of DDoS attacks 7

- Revenue loss 7
- Brand damage 7
- Foreseeable and unforeseeable costs 7

DDoS protection options 8

- Carrier-based DDoS mitigation 8
- CDN-based DDoS mitigation 9
- DNS-based DDoS mitigation 9
- Hosting provider DDoS mitigation 10
- In-house DDoS mitigation 11
- Outsourced specialist DDoS protection 11

Nexusguard's integrated solution in countering DDoS attacks 12

- Overview 12
- Nexusguard's multi-layered defense system 14
- Progressive challenge-response (C/R) algorithms 16
- SSL attack mitigation 19
- Patented crawler identification technology 22
- Origin Protection 25
- DNS Protection 26
- Content and networking optimization 29
- Visibility and control 30
- 24x7 SOCs staffed with DDoS experts 32

Conclusion 33

Introduction

In less than a decade ago, DDoS was a fairly straightforward cyber-weapon with which the attacker sends a large amount of requests to paralyze a web server and shut down online services. Over the years, as the Internet has evolved rapidly in terms of bandwidth and technology, DDoS attacks are no longer just a nuisance. Nowadays, apart from being gigantic in size and more advanced in complexity, they are increasingly being employed as part of a more complex hacking plot—very often targeting the sensitive or critical customer information, such as credit card data, and even real money.

As enterprises and organizations, no matter big or small, increasingly rely on uninterrupted online presence, the impact of DDoS attacks cannot be neglected. They can cause network, server and application downtime, and sometimes service degradation, which can cause direct revenue or customer loss, incur additional expenses and tarnish the victim's reputation overnight.

While legacy and on-premise appliances can handle simple, small-scale attacks to a certain extent, they are not designed to mitigate large-scale, volumetric attacks and more complex layer-7 attacks. Capacity wise, large-sized businesses and even the service provider upstream may also lack the bandwidth to absorb the traffic spike—not to mention small businesses on a limited budget. Moreover, the in-house security team may lack the expertise and skills to handle and respond to sophisticated DDoS attacks.

Nexusguard is the world's leader in countering DDoS attacks, leveraging its high-performance scrubbing centers located around the world—collectively loaded with 1.44Tbps of mitigation capacity. The scrubbing network is highly scalable and is ready at all times to buffer and mitigate attacks of all sizes and shapes. We employ multi-layered mitigation mechanism, including proprietary technologies, to identify, mitigate and analyze attacks effectively. As the solution is delivered over the cloud, there is no upfront cost and ongoing costs for appliance maintenances or software upgrades.

This paper provides an overview of the common types of DDoS attacks, negative impact of DDoS attacks on business and operations as well as basic mitigation techniques. The document further presents Nexusguard's integrated solution in countering DDoS attacks and other cyber threats, including our multi-layered mitigation technology, SSL attack mitigation methods, Web Application Firewall (WAF), content and networking optimization, as well as the Customer Portal.

Common types of DDoS attacks

DDoS attacks are a growing threat to enterprises and organizations relying on uninterrupted online presence. In its most native form, the attacker floods the target server or network by sending massive requests to saturate network capacity or exhaust network resources, so that legitimate users are denied access. Larger, more sophisticated DDoS attacks are often masterminded by the attacker behind botnets, or networks of compromised computers or devices. DDoS attacks can be broadly divided into three broad categories:

Volumetric attacks

Volumetric attacks are floods of junk traffic sent by the attacker through distributed botnets to overwhelm the target infrastructure. They include UDP floods, ICMP floods, and other spoofed-packet floods. The scale is measured in bits per second (bps).

Protocol attacks

Protocol attacks exploit a specific feature or implementation bug of some protocol installed at the victim in order to consume excess amounts of its resources. This type of attack consumes actual server resources, or those of intermediate communication equipment, such as firewalls and load balancers, and is measured in packets per second (pps). Common protocol attacks include smurf attack, SYN attack, UDP attack, ICMP attack, CGI request attack, authentication server attack, attack using DNS systems, and attack using spoofed address in ping.

Application layer attacks

A Network Time Protocol (NTP) amplification attack relies on the use of publicly accessible NTP servers to overwhelm a victim system with UDP traffic. The NTP service supports a monitoring service that allows administrators to query the server for traffic counts of connected clients. This information is provided via the "monlist" command. The basic attack technique consists of an attacker sending a "get monlist" request to a vulnerable NTP server, with the source address spoofed to be the victim's address.

The attack relies on the exploitation of the 'monlist' feature of NTP, as described in CVE-2013-5211, which is enabled by default on older NTP-capable devices. This command causes a list of the last 600 IP addresses, which connected to the NTP server to be sent to the victim. Due to the spoofed source address, when the NTP server sends the response, it is sent instead to the victim. Because the size of the response is typically considerably larger than the request, the attacker is able to amplify the volume of traffic directed at the victim. Additionally, because the responses are legitimate data coming from valid servers, it is especially difficult to block these types of attacks. The solution is to disable "monlist" within the NTP server or to upgrade to the latest version of NTP (4.2.7), which disables the "monlist" functionality.

Hypertext Transfer Protocol (HTTP) flood

In HTTP flood attack, the attacker exploits seemingly legitimate HTTP GET or POST requests to attack a web server or an application. HTTP floods do not use malformed packets, spoofing or reflection techniques. On the contrary, it requires less bandwidth than other attacks to bring down the targeted site or web server. The attack becomes most effective if it can force the server or application to allocate the maximum resources possible to handle each request.

Zero-day attacks

"Zero-day" attacks are literally previously unknown attacks, exploiting inefficiencies and vulnerabilities for which no patch has yet been released.

UDP flood

User Datagram Protocol (UDP) is a sessionless networking protocol. In a UDP flood, the attacker sends large volumes of UDP traffic to flood the random ports on the target machine with packets that force it to listen for applications on those ports and report back with a ICMP packet. The goal is to disable the machine or fill the bandwidth up with attack traffic.

SYN flood

In a SYN flood, the attacker sends a succession of SYN requests to the target system in an attempt to consume enough server resources to make the system unresponsive to legitimate traffic. The TCP three-way handshake is the foundation of this form of attack. A SYN flood attack works by not responding to the server with the expected ACK code. The malicious client can either simply not send the expected ACK, or by spoofing the source IP address in the SYN, causing the server to send the SYN-ACK to a falsified IP address—which will not send an ACK because it "knows" that it never sent a SYN.

The server will wait for the acknowledgement for some time, as simple network congestion could also be the cause of the missing ACK. However, in an attack, the half-open connections created by the malicious client bind resources on the server and may eventually exceed the resources available on the server. At that point, the server cannot connect to any clients, whether legitimate or otherwise. This effectively denies service to legitimate clients. Some systems may also malfunction or crash when other operating system functions are starved of resources in this way.

Nuke

Corrupt and fragmented ICMP packets are sent via a modified ping utility to keep the malicious packets to be delivered to the target. Eventually, the target machine goes offline. This attack focuses on comprising computer networks and is an old distributed denial of service attack.

Multi-vector attacks

Multi-vector attacks are the most complex forms of DDoS attack. Instead of utilizing a single method, a combination of tools and strategies are used to overwhelm the target and take it offline. Often times, multi-vector attacks will target specific applications on the target server, as well as, flood the target with a large volume of malicious traffic. These types of DDoS attacks are the most difficult to mitigate because the attack come in different forms and target different resources simultaneously.

DNS attacks

DNS is the network address behind every IP address. DNS transforms a domain name and converts it into the actual IP address. It is possible for many domain names to have the same IP address because one server can support a huge number of domain names. One DNS name can also be configured to map to several IP addresses. For example, if a URL maps to five different addresses, a web browser will go to any one of them to access the site. DNS attacks are used for attacking both the infrastructure and a DNS application. This type of attack allows the attacker to use both reflection and spoofed direct attacks that overwhelm the target's infrastructure by sucking up all available network bandwidth.

Advanced Persistent Threat (APT)

An APT is a sustained cyber espionage led by a powerful entity, such as a government or organized criminals, with the intent to gain access to a specific target, such as a political resistance group or another government. APTs often use DDoS attacks, combined with other hacking techniques. APT-style DDoS attacks can be characterized by long durations, repetition and changing attack vectors aimed at evading simple, signature-based defense systems.

A DDoS attack may also be used as a smokescreen to distract IT and security teams from a more serious cybercrime, such as theft of critical data, intellectual property, customer data, etc. The damage caused by severe data breach to the organization and its customers could be disastrous.S

Impact of DDoS attacks

The impact of DDoS attacks can be devastating especially if the organization or business relies heavily on uninterrupted online presence to maintain e-commerce business, online operations and/or brand equity. Unprepared for such threats, a sudden attack could cause direct revenue and customer loss, damage brand and reputation, and incur extra costs in taking remedial actions.

Revenue loss

Your customers will be upset about your website and service if they need to wait a long time to load the page or access to its applications, which in turn will discourage them from paying for your product or service again. If e-commerce business—anything from online toy store to security trading platform—is your major revenue source, downtime causes direct revenue loss and it could be huge, taking into account the revenue loss during the downtime and the reputational damage, combining negative word of mouth and customer dissatisfaction, going forward.

Brand damage

While revenue losses caused by downtime can be calculated, the tarnished reputation is more difficult to measure. Beyond that, don't forget that first-time visitors, who were denied access to your website, were unimpressed and chance is they will never return. Existing customers are also likely to switch to another vendor or service provider, especially if the industry is highly competitive, where products or services are largely homogeneous (consider stock trading). If you a service provider that has service level agreements for uptime, you may be subject to penalties.

Foreseeable and unforeseeable costs

Even in worst-case scenarios, you could still seek help from a reputable DDoS protection service provider, such as Nexusguard, to get emergency protection when you are caught off guard by a sudden attack. But the costs in taking remedial actions could be far beyond previously thought, such as the extra costs to the hosting service provider for extra bandwidth consumed, overage charges, staff overtime, extra manpower, etc; and even worse, the costs needed to clean up the mess when a data breach occurs. Other financial costs and resources may include investor communications, customer service, rebuilding of reputation, lawsuits, recovery operations, etc.

DDoS protection options

As DDoS attacks present a growing security risk, companies that offer to protect your business and mitigate attacks make up a growing industry. There are many options for DDoS mitigation services, including:

- Carrier-based DDoS mitigation
- CDN-based DDoS mitigation
- DNS-based DDoS mitigation
- Host-based DDoS mitigation
- In-house DDoS mitigation
- Outsourced specialist DDoS mitigation

Each solution has its own benefits and limitations.

Carrier-based DDoS mitigation

Ordered directly from one of the major carriers that is responsible for handling online traffic, carrier-based DDoS mitigation is available from providers including:

- AT&T
- Verizon
- Telstra
- Tata
- PCCW

These services monitor traffic and identify the signs of an attack, refusing access from malicious traffic. However, a major concern with carrier-based mitigation is that traffic can only be scrubbed by the networks your chosen carrier operates. As attacks tend to be more widely distributed across multiple providers, this limited approach is rarely effective.

In addition, while the size of attacks has grown considerably in recent years, carrier-based DDoS mitigation often provides a restrictive amount of DDoS bandwidth. If attacks peak under 8-10 Gbps, this type of solution could cope. If not, your account could be taken offline until the attack ends.

Most carrier-based solutions are implemented using commercial off-the-shelf equipment, providing a basic level of protection with little to no customization. These devices are usually left as the only line of defense, without the support of a dedicated security operations center (SOC). As a result, enterprises have little protection against application layer attacks.

Finally, carriers tend to be significant targets for their own DDoS attacks. In 2012, AT&T, a worldwide leader in IP communications, suffered a DDoS attack that took DNS servers offline for more than eight hours.6 As high profile targets, a carrier-based anti-DDoS solution could attract attackers while defending against them.

CDN-based DDoS mitigation

A content delivery network, or CDN, brings together servers around the world that host and serve website content.. A major benefit of using a CDN is locally served static content, improved uptime, and a lighter load on your servers that provide dynamic content including resource intensive web applications.

Since a CDN offers large amounts of bandwidth across the entire network, many CDN providers offer a DDoS mitigation service. These services absorb simple volumetric attacks into the CDN, with enough remaining bandwidth to serve legitimate traffic.

For a CDN, scale of attack is not a significant issue. With enough hardware resources and bandwidth across every server, even the largest of volumetric attacks can be successfully mitigated. Legitimate requests for data are still served effectively, and web applications remain online.

However, this approach does not account for the more intelligent, application-based attacks that are becoming more and more prevalent. CDN-based mitigation is a simple solution to the simplest of attacks, but incapable of helping you combat more sophisticated threats.

In addition, as a CDN absorbs attack traffic, businesses are usually expected to pay for bandwidth usage. As attacks grow larger and last longer, these uncontrollable costs can have a serious financial impact. In this sense, an attack causes business disruption, even with no service loss.

Finally, while a content delivery network can continue to serve static content even if individual servers go offline, important dynamic content is not stored in the CDN. This creates a significant target for attackers; if a dynamic content server is taken offline, crucial functionality is lost.

DNS-based DDoS mitigation

When legitimate users enter a domain name, a Domain Name Server (DNS) is responsible for translating this name into an IP address. Domain name servers are a routine target for DDoS attacks themselves, and act as a channel for traffic in all types of attack.

Similarly to CDNs, DNS providers have access to large quantities of bandwidth. This bandwidth can be used to absorb attack traffic, ensuring that Layer 3 and 4 attacks do not interrupt the service of legitimate users.

However, DNS-based DDoS mitigation is limited in the same way as a CDN-based solution. The reality is that, while solutions from DNS providers are convenient, they fail to offer blanket protection across every threat. Organizations that use DNS-based mitigation remain open to Layer 7 attack traffic, which is difficult to detect without careful analysis.

Domain name servers can be taken offline by intelligently designed, low bandwidth attacks. DNS-based DDoS mitigation does nothing to eliminate this weakness.

Hosting provider DDoS mitigation

As the size of the DDoS mitigation market continues to grow and companies become more aware of the threat of attack, it is little surprise that web hosting companies have seized the opportunity to add value to their services. Anti-DDoS services can usually be added to existing web hosting plans at a very affordable price.

To add DDoS mitigation as a selling point of hosting services, companies can use a range of different methods. A host could deploy and manage its own hardware within its network, or leverage other services including carrier-based, DNS-based, and CDN-based mitigation. As a result, solutions have different benefits or limitations depending on the type of solution that has been used.

However, DDoS mitigation services from hosting companies tend to have certain things in common. First, these solutions can only protect data that is hosted within their networks. If companies use a combination of hosted services and in-house platforms to serve content and applications, they need to launch and maintain multiple solutions at significant cost.

In addition, DDoS protection from a hosting company usually works on shared bandwidth. In the same way as hosting services can drop in performance and reliability if a few users monopolize hardware resources, the quality of host-based DDoS mitigation can be unpredictable. In this way, some anti-DDoS services from hosting companies fail to mitigate even basic volumetric attacks.

Finally, host-based DDoS mitigation reflects the increasing diversification of hosting services. Although a web hosting company may have extensive expertise in storing and serving content, it is unlikely that the same company will have the specialist expertise that is required to effectively protect against DDoS attacks.

In-house DDoS mitigation

Faced with the limitations of carrier-based, CDN-based, DNS-based and host-based anti-DDoS services, many organizations understand that a more specialized approach is required. One way to achieve this is to bring mitigation in-house. There are many vendors that offer dedicated hardware for DDoS mitigation, including threat tracking and powerful behavioral analysis.

Using this hardware within the enterprise, a company can retain complete control and ownership of their solution. However, bringing mitigation in-house brings with it a number of potential problems. The high costs of purchasing hardware, installation and maintainance, as well as the extra manpower dedicated to keeping the equipment up and running, are some of the disadvantages. Expenses can quickly grow out of control, especially for small businesses.

Considering these capital and operating expenditures, many companies that take an in-house approach to DDoS protection will not break even for between three and five years. In a threat landscape that is constantly changing and evolving, new equipment is usually required before the three to five year period. The prohibitive costs of in-house DDoS mitigation coupled with regular upgrades and innovations make it very difficult – almost impossible – for companies to break even.

And very often, these hardware appliances could become the bottlenecks themselves when a large-scale, volumetric attack exceeds the threshold they could handle. In addition, implementing in-house DDoS mitigation requires a significant investment of time and expertise. Often, employees have a good understanding of the general issues surrounding DDoS attacks, but are unable to commit themselves to the topic on a full-time basis.

Outsourced Specialist DDoS Protection

Nexusguard currently partners with a number of service providers to offer turnkey DDoS protection to their end-clients. These service providers are well backed up by Nexusguard's globally located scrubbing centers and mitigation technologies.

Alternatively, to get the right, sufficient and cost-effective DDoS protection, organizations are strongly advised to seek out dedicated DDoS protection services, such as those provided by Nexusguard, that are highly scalable, resilient and can withstand DDoS attacks of all sizes and types.

Nexusguard's integrated solution in countering DDoS attacks

Overview

DDoS mitigation services deployed over the cloud is the ultimate solution for organizations looking for comprehensive protection. In other words, Nexusguard as your DDoS protection service provider stands in front of the website and identifies, mitigates and keeps track of attack traffic by checking incoming traffic thoroughly against a database of constantly updated attack signatures and IP reputation.

Nexusguard's rigorous, multi-layered scrubbing process is spread out over a vast scrubbing network geographically located around the world. It is equipped with the highest levels of scalability and resilience required to handle the largest network DDoS attacks.

We provide comprehensive, professional and cost-effective protection against the most potent DDoS attacks. Enterprises of all sizes across all industries can benefit from the protection with Nexusguard's massive, globally distributed cloud infrastructure. We deliver 100% uptime guarantee, and faster, more stable performance and offer immediate response and support from our 24/7 Security Operations Centers.

1.44Tbps of mitigation capacity

Globally distributed scrubbing centers located in San Jose, CA; Los Angeles, CA; Miami, FL; Ashburn, VA; London; Amesterdam; Singapore; Taiwan; and Hong Kong.

Types of DDoS attacks mitigated

To this day, DDoS protection, SSL attack mitigation and WAF come in different solutions or packages and need dedicated internal resources and manpower to keep them up and running. Our integrated solution combines the essential modules to make businesses resilient to DDoS and other forms of cyber-attacks. Delivered through cloud, our solution is highly scalable and is stand by at all times to detect, mitigate and eliminate them all, without compromising on user experience.

We are capable of accurately detecting and mitigating a large variety of attacks, including but not limited to:

Category	Attack Type	
Bandwidth/ Network Depletion Attacks	Protocol Flood / Exploitation Attacks	TCP Flood
		UDP Flood
		ICMP Flood (Smurf, Ping Flood, Ping of Death, ICMP Echo)
		TCP SYN, SYN/ACK, RST, FIN Flood (Spoofed and Non-spoofed)
		IP Null
		Fragmentation (IP/UDP, IP/ICMP, IP/TCP, Teardrop)
		DNS Amplification
		Fraggle
		Nuke
		TCP Flag Abuse
		Zombie / Bots Attack
Application-based Attacks	HTTP Attacks	HTTP GET/POST Flood
		HTTP Page Flood
		HTTP Connection Flood
		HTTP Malformed Request
		HTTP 404
		Slowloris
		Socketstress
		Slow HTTP
	DNS Attacks	Reflected DNS
		DNS Query
		DNS UDP flood
		DNS TCP flood
		DNS Malformed Query
		Protocol and Vulnerability Exploitation DoS/DDoS
	Hacks	SQL Injection
		Cross Site Scripting (XSS)
		Cross Site Request Forgery (XSRF)
		Session Hijack
Others		Malicious Headers
		Malicious Payloads
		Pucodex
		Zero Day Exploits

Nexusguard's multi-layered defense system

Preventing dynamic DDoS attacks require real-time, comprehensive, and meticulous detection and action. Nexusguard guards against DDoS attack traffic through multiple layers of inspection to deliver fast, clean traffic.

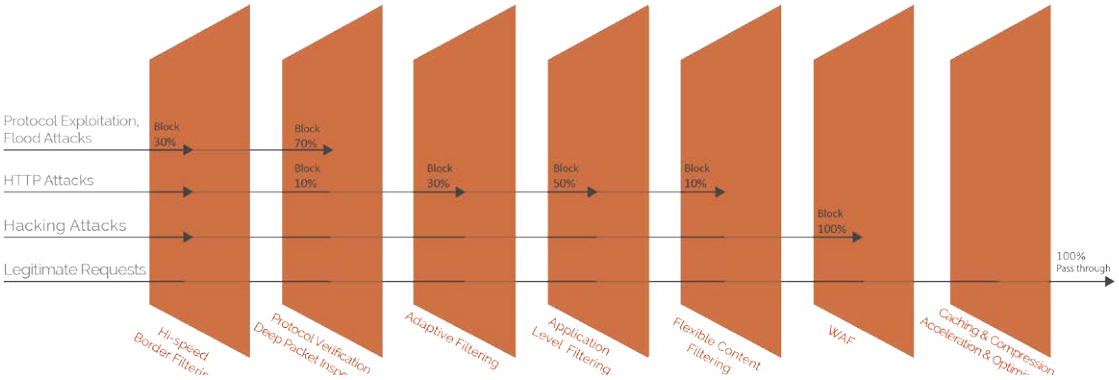

Hi-speed border filtering

We have established peering connections with multiple core ISPs to provide multi-gigabit attack protection. Each peer is closely monitored and continuously evaluated in order to deliver the fastest response time to customer's critical and latency-sensitive applications. At Nexusguard's border, traffic is filtered for bandwidth flood using wire-speed Access Control Lists. Nexusguard also keeps tracking lists of bogus IPs and infected hosts, which are also filtered at this layer.

Protocol verification and Deep Packet Inspection (DPI)

At this level, protocol state such as TCP three-way handshake is verified. SYN flood and other similar attack attempts that do not conform to protocol standard are also filtered out. To mitigate spoofed attacks, Nexusguard uses progressive challenge-response algorithms (see the next section for details) to distinguish between spoofed and legitimate traffic with surgical precision.

Adaptive filtering

Nexusguard enforces both Statistical Analysis and Anomaly Recognition filtering for zero day attacks. Using Statistical Analysis, unusual number of packets or high traffic rate from zombie clients can be identified and filtered. Using Anomaly recognition, auto-learning of normal baselines for protocol and source networks flows can be used to identify and filter malicious activities.

Application-level filtering

Nexusguard's DPI engine provides comprehensive application-layer intelligence, allowing Nexusguard to understand what applications are running on the client's network to efficiently detect and deter application traffic violations.

When faced with an increasing number of attacks from larger-sized clients (or zombies) using valid established connections to overwhelm the system resources, Nexusguard's anti-zombie system mitigates such HTTP attacks by using a challenge response authentication process to differentiate between legitimate browsers and zombie programs that access the attacked site.

To further mitigate application specific level attacks - HTTP attacks, Nexusguard can enforce intelligent HTTP Malformed filtering to ensure the validity of HTTP transactions, and limit the number of connections or request to specific objects.

Flexible content filtering

Nexusguard DDoS Mitigation system continuously monitors application traffic for unusual pattern and behavior. Using its proprietary pattern recognition and analysis system, Nexusguard deters morphing HTTP Flood attacks by adapting flexible-content filters to counter evasive intents rapidly.

Rate limiting

Rate limiting will be applied to further limit exploitation of system and bandwidth resources against baseline statistics.

Web Application Firewall (WAF)

As part of Nexusguard's DDoS protection service offerings, our Web Application Firewall (WAF) aims to protect web applications, mobile apps and application program interface (API) apps against unsolicited intrusions with layered protection.

It stops attacks at the network edge, protecting your website from common web threats and specialized attacks before they reach your servers. It covers both desktop and mobile websites as well as applications by examining HTTP requests. It examines both GET and POST requests and applies rules to help filter out illegitimate traffic from legitimate users.

You can decide whether to monitor or block malicious and attack traffic, including an option to challenge suspicious users. With blocking and challenging, our WAF will block any traffic identified as illegitimate before it reaches your origin web server.

Our WAF employs blacklist rules to address known security risks, in which the rule set is primarily based on the OWASP ModSecurity Core Rule Set as well as those we have developed based on the previous and current attacks on our other customers.

Our WAF examines all requests and you can decide whether to monitor or block an attack, with an option to challenge suspicious users. More specifically, our WAF checks HTTP request headers, where things like cookies, authorization tokens and miscellaneous request information is stored. It then checks the HTTP request body, including the parameters and data values being sent to the application.

During the process, the WAF detects the OS, web framework/component, database, and even the development language running on the back-end server. This enables the WAF to apply OS-specified rules to make the detection more accurate and efficient. Besides, it makes virtual patching more accurately--some applications tend to be more vulnerable when running on a specific OS.

We work with our clients to define site-specific rules and constantly update our rule set to ensure that our WAF is well equipped against the latest threats. On the Customer Portal, customers may also define and update their own rules for their sites as well.

Progressive challenge-response (C/R) algorithms

Overview

Our DDoS mitigation engine consists of the challenge-response (C/R) algorithms drawing on years of experience in cyber-security and intelligence collected from a large pool of clients. Our proprietary C/R algorithms can effortlessly defend the application layer of the network against all sorts of abuses and attacks, while keeping the user experience as smooth and seamless as possible. It employs non-intrusive authentication challenges depending on user behavior. The service delivers a non-disruptive browsing experience and zero false-positive mitigation errors.

The challenge authentications are based on continuous learning of user behaviors followed by dynamically tuned progressive challenge thresholds. Requests that do not comply with the unique identifiers in a browser are considered suspicious and will be directed to go through a set of progressive challenges.

Nexusguard's progressive C/R protocol comprises of the following validation methods. To minimize the impact on user experience, Captcha authentication is only implemented as the last resort.

Level 1

HTTP protocol behavior validation – It first looks at the HTTP protocol's behavior for clues of malicious activity, such as HTTP header field ordering, forbidden operation response, improper HTTP version response, and improper protocol responses.

HTTP redirect authentication – This technique artificially redirects HTTP 302, which distinguishes legitimate browsers from automated tools.

HTTP secure-cookie authentication – This method works like, and is usually used together with, HTTP redirect authentication. During the examination, the way the browser handles cookie is tested. Clients that do not carry cookies in subsequent HTTP requests are clear suspects and can safely be blocked.

Level 2

JavaScript compute engine verification – A piece of JavaScript code embedded in the HTML is sent to clients as a challenge. Only clients equipped with a full-fledged JavaScript engine can perform the computation. It would not be economical for DDoS attack tools to hijack or otherwise make use of a real heavyweight browser to carry out attacks. The purpose of this authentication is to identify if the HTTP request is sent from a real browser.

Level 3

Captcha authentication – As the final gatekeeper, Captcha challenges will be inserted into suspicious traffic. If the client is successful in solving the Captcha, it will be whitelisted for a certain period of time or for certain amount of subsequent traffic, after which it will need to authenticate itself again.

Captcha authentication is intended to distinguish whether the request is initiated by a real human or a bot. This method proves to be the most effective way to authenticate legitimate users, but inevitably diverts the user's attention, meaning user experience is compromised. To minimize disruption to users, Nexusguard uses Captcha as the last resort only.

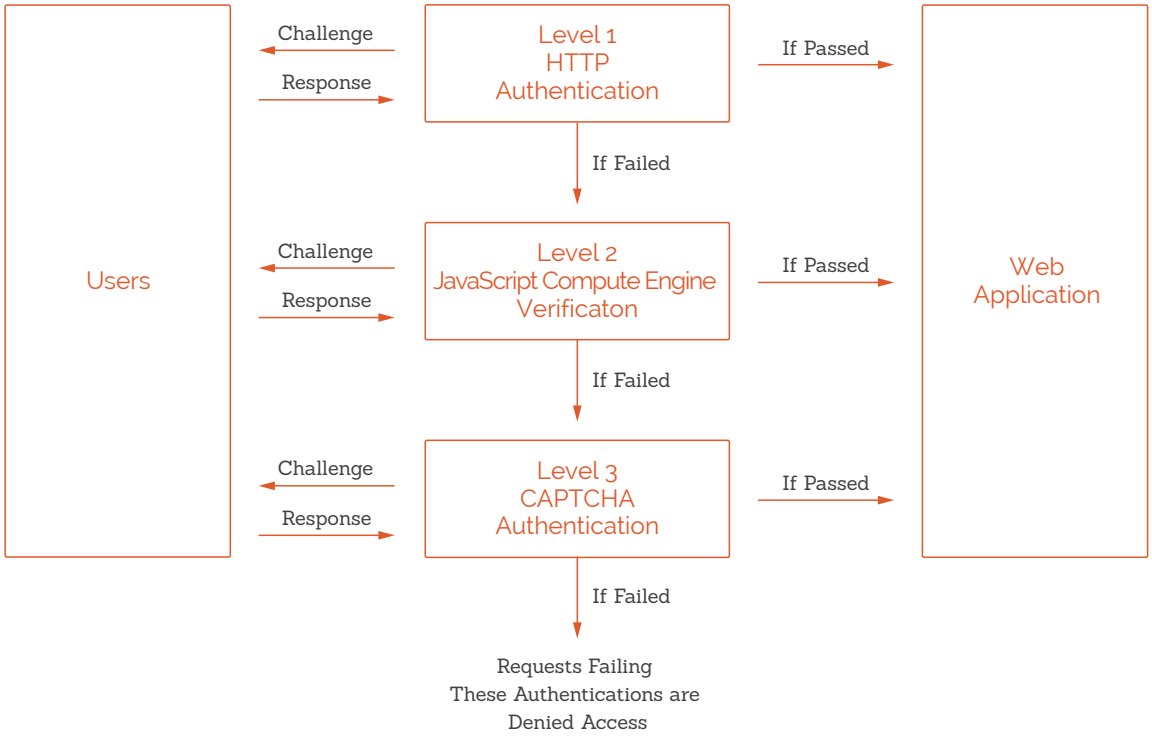

Source- and destination-based mitigation techniques allow for surgical precision in countering DDoS attacks

Nexusguard's mitigation technology employs both source- and destination-based techniques. In other words, malicious traffic is assessed and blocked based on the source IP, followed by the destination IP, well before they can reach the client's server.

Nexusguard's progressive C/R is the technique employed to block the attack right at the very source while allowing the remaining legitimate traffic to be routed normally. The advantage is that malicious traffic is traced down to the browser by inspecting each browser session. Such source-based blocking technique can significantly reduce false-positive errors by not blocking IP addresses shared by compromised and innocent computers but pinpointing only the compromised machines.

If a large botnet is used or if the attacker uses spoofed IPs, source-based blocking may not be enough. Therefore, Nexusguard's ISP-grade edge routers in the scrubbing center located closest to the source takes charge to filter out and drop immediately identifiable malicious packets. The rest of the traffic then goes through the multiple layers of filtering at Nexusguard's scrubbing centers around the world.

The combination of source-based and destination-based mitigation techniques guarantees a zero false-positive mitigation result. Legitimate users do not even feel the site they are visiting is under a prolonged DDoS attack.

SSL attack mitigation

SSL is the most widely used solution that uses "private" and "public" keys for encrypting traffic. Despite its trusted security, this solution uses much more processing power than most other server requests. As such, DDoS attacks on Source- and destination-based mitigation techniques allow for surgical precision in countering DDoS attacks.

Nexusguard's mitigation technology employs both source- and destination-based techniques. In other words, malicious traffic is assessed and blocked based on the source IP, followed by the destination IP, well before they can reach the client's server.

Nexusguard's progressive C/R is the technique employed to block the attack right at the very source while allowing the remaining legitimate traffic to be routed normally. The advantage is that malicious traffic is traced down to the browser by inspecting each browser session. Such source-based blocking technique can significantly reduce false-positive errors by not blocking IP addresses shared by compromised and innocent computers but pinpointing only the compromised machines.

If a large botnet is used or if the attacker uses spoofed IPs, source-based blocking may not be enough. Therefore, Nexusguard's ISP-grade edge routers in the scrubbing center located closest to the source takes charge to filter out and drop immediately identifiable malicious packets. The rest of the traffic then goes through the multiple layers of filtering at Nexusguard's scrubbing centers around the world.

The combination of source-based and destination-based mitigation techniques guarantees a zero false-positive mitigation result. Legitimate users do not even feel the site they are visiting is under a prolonged DDoS attack.

SSL-encrypted attacks are very effective, as a small traffic volume generated by the attacker is enough to overwhelm the target. Additionally, encrypted attacks can easily bypass detection if SSL-based protection is not supported.

As part of our total anti-DDoS solution, we support SSL-encrypted attack mitigation. Our SSL certification management follows the PCI Data Security Standard and ISO27001. In doing so, our scrubbing centres become the intermediate for all incoming traffic to your website, including SSL traffic. We offer three SSL traffic-handling options, namely Offloading, Bridging and Forwarding, to maximise the DDoS mitigation results and reduce false-negatives.

Offloading – SSL traffic is decrypted at our scrubbing centers and then returned to your web servers in clear-text format. This method relieves your web servers of processing heavy encrypting/ decrypting traffic via SSL, therefore improving server performance.

Bridging – SSL traffic is decrypted at our scrubbing centres and then re-encrypted when sent back to your web servers. As data is SSL-encrypted en route, this method offers the highest protection.

Forwarding – SSL traffic is forwarded to your web servers directly without decryption in between.

While most clients trust Nexusguard to look after their private key, for some reason, such as regulatory requirements or internal security policies, some organisations and companies do not want to turn over their private key to third party vendors.

Separate key pair signed by a trusted CA

Currently employed by Nexusguard, this easy yet highly effective approach uses two key pairs for the Visitor-to-Nexusguard and Nexusguard-to-Server sessions, respectively. To be specific, the original key pair, including the site owner's private key, is only used in the Nexusguard-to-Server session, so that the site owner remains in control of its private key.

How it works?

To do so, we require a valid SSL certificate for the site owner's domain to be installed in our cloud servers located worldwide. To provision this certificate, we partner with some of the world's leading certificate authorities (CAs) to enable the site owner to issue Nexusguard with a certificate for its domain, in a very simple way and at no additional cost.

The original web server's certificate will still be used for Nexusguard-to-Server requests, while the new Nexusguard-issued certificates will be used to handle all Nexusguard-to-Visitor traffic.

To ensure maximum security, these Nexusguard issued certificates are signed by the world's leading SSL authorities, supported by all major browsers and devices and are trusted by 99.9% of the Internet population. Moreover, they are using 2048 bit key encryption, thus boosting security for all 1024 bit certificate holders.

Due to the limitation that SSL requires its own IP address and that each of our data centres around the world needs to have a valid version of each certificate, we resort to using multi-domain SSL certificates (MDC), allowing us to include multiple domain names within the same certificate and solve the IP shortage problem. Simply put, we can use one single certificate to serve all domains of the site owner.

By this mean, the authorisation of the use of our custom SSL certificate for the site owner's domain is carried out via an API provided by our partner Certificate Authority (CA), which will be embedded into an email for the site owner's approval. This gives the site owner the right to revoke the certificate, at anytime, if necessary.

Patented crawler identification technology

Overview

Search engines and social networking sites send out bots, i.e. crawlers, to understand what your site content is about and help you spread the message across. Bots like crawlers are known as good bots and they are most welcome by site owners. But with good always comes bad: Apart from being used to perform productive tasks, bots are also increasingly used to do evil things, such as spam email harvesting, site cloning without permission, implanting viruses and worms, manipulating page views (to cheat advertisers), and many more. A large volume of bad bots could also paralyze the network, depriving legitimate users from accessing the resources.

It is apparent that websites need to let in the right amount of good bots in order to be seen and heard, and deny access from bad bots in order to protect the network from abuses and thefts. However, the proliferation of bad bots is already posing substantial threats to site owners -- not only could they steal your valuable digital assets, but they could also destroy your SEO efforts.

The old-fashioned way of blocking bad bots and rogue crawlers is to customize the robots.txt and .htaccess files after examining the web log for unsolicited bots. Most anti-DDoS service providers offer a simple whitelist filter to block some (not all) bad bots. But if the site owner wants customized protection, he is required to blacklist the suspicious IPs on his own. The drawback is that, without an authoritative crawler whitelist and blacklist validation, it is extremely difficult for average site owners to distinguish between good and bad bots.

Benefits

Effectively identifying bad bot traffic, good bot traffic and human traffic needs a great deal of expertise, experience and technology that are far beyond the capability of most site owners. Therefore, Nexusguard has developed a proprietary, Patented search engine crawler identification technology that accurately segregates legitimate crawlers from spoofed or illicit ones, delivering:

- 100% search engine access
- Rejection of all forged search engine requests
- Enhanced SEO optimization and improved search engine rankings

How it works?

As the proxy on behalf of the client's server, Nexusguard examines the incoming traffic during each HTTP session before passing it to the client's server. The bad bot filter, which consists of Nexusguard's proprietary signature database, effectively and accurately detects and blocks bad bots, and every action taken -- passing or dropping -- will be logged.

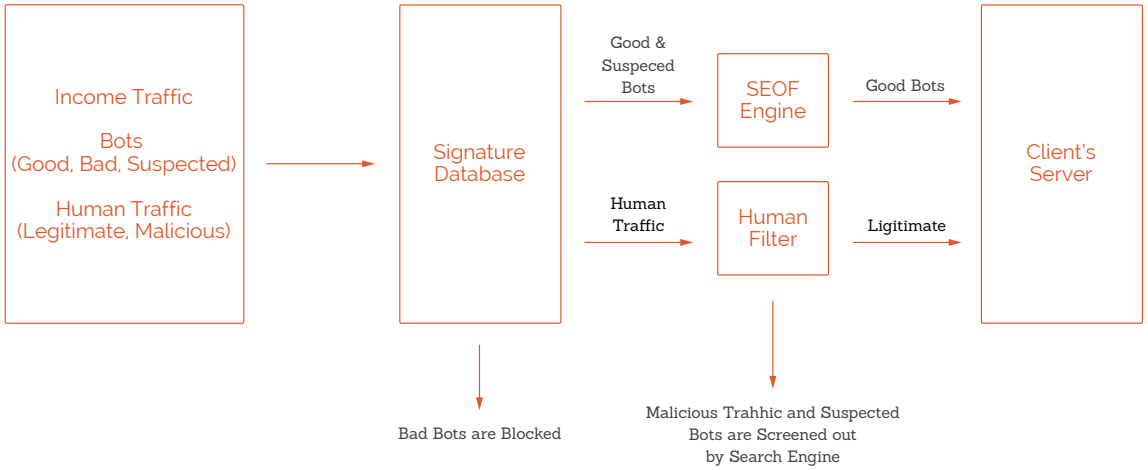

Two-tiered filtering of bad bots and malicious traffic

The signature database consists of a crawler whitelist authorised and maintained by leading search engines and a verification mechanism to cross-check the IPs claiming they are good bots. Under this stringent scrubbing process, spoofed sources stand no chance of distinguishing as good bots. At this point, good bots are let in, and bad bots are blocked.

The rest of the HTTP requests that have crossed the above threshold are considered human activity. If malicious activity, including suspected bots, is detected, Nexusguard's proprietary challenge/response (C/R) protocol kicks in to verify if a DDoS attack or other cyber threats take place.

Depending on activity levels, malicious traffic will have to go through a non-intrusive, progressive C/R identification process, from HTTP protocol behavior checking, secure-cookie authentication and JavaScript compute engine verification, to captcha as the last resort. All malicious traffic, including suspected bots, are screened out here and only legitimate traffic can reach the client's server.

Origin Protection

Overview

Nexusguard's Origin Protection covers all elements of your network, e.g. internal websites, email servers, FTP servers, and other applications, against all volumetric and protocol-based DDoS attacks, such as UDP, SMTP or SYN floods. All incoming traffic is routed through Nexusguard's scrubbing centers using BGP announcements and only clean traffic will be routed through a secure GRE tunnel back to your servers. In this sense, Nexusguard advertises all protected IP range announcements on your behalf.

Origin Protection protects multiple service types and protocols, even executed via a single IP address, without using BGP routing. Clients relying on a small number of IP addresses or even a single IP address to deliver high-traffic, non-HTTP services, as well as clients using cloud services over the public Internet via a dedicated IP address, will find Origin Protection exceptionally cost-effective and valuable.

Protection for individual IP addresses

Using this unique deployment model, clients who do not have an entire Class C subnet can benefit from infrastructure protection in its entirety. This feature enables smaller organizations to protect multiple service types and protocols, even for a single IP address, without using BGP routing.

Customers receive a "protected IP address" from Nexusguard, which inspects and filters all incoming traffic. A redundant, secure, two-way GRE tunnel is used to forward clean traffic to the origin IP and to return outbound traffic from the application to the users. Once in place, this tunnel is used to route clean traffic from our network to your origin, and vice versa.

You then broadcast the assigned IP addresses to your users via DNS, making these your nominal "origin" addresses. In the future, all incoming traffic flows through the Nexusguard's network, is inspected by our global scrubbing centers, and then forwarded to the origin via your GRE tunnel. Outgoing traffic also passes through Nexusguard's network.

Individual IP address protection is ideal for mission-critical online services, which have high-traffic, critical non-HTTP assets with low IP counts, as well as cloud deployments looking for direct-to-IP attack prevention.

DNS Protection

Overview

DNS plays a critical role in how the Internet and TCP/IP work. Especially in the case of Internet-facing DNS servers, it is important to ensure your DNS infrastructure is protected from attack from outside or even from within your organization.

Nexusguard's DNS Protection protects mission-critical online services from DNS attacks. Our dedicated and fully redundant network of globally distributed proxy servers resolves incoming DNS queries quickly and reliably. Our Security Operation Centers deploy advanced, traffic scrubbing technologies around the clock to keep your DNS service up and running, guaranteeing uninterrupted uptime.

Attacks on DNS servers can critically impact online business operations, harming a company's reputation and causing financial losses. Unfortunately, DNS attacks resulting in downtime for web services are common for companies doing business on the Internet.

Run-of-the-mill DNS hosting providers are usually not experienced in DDoS mitigation, and thus offer no remedy in the event of an attack on your DNS domain. Since they cannot react effectively, such providers may simply drop your DNS to protect their other clients, causing your DNS entry —and your business—to disappear from the online world. .

Patching DNS vulnerabilities

Virtually all network applications use distributed resources. Under a tree-like system of delegation and a series of repetitive, recursive processes, DNS plays the role of middleman, forming a traffic bridge connecting clients and servers.

Similar to SMTP relays or web proxies, recursive DNS servers accept messages from clients and forward them on to other servers when necessary. Ideally, a name server would only accept queries from local or authorized clients.

However, many will accept DNS queries from any source, and many DNS implementations also enable recursion by default — even when authorized data is only supposed to be handled by the name server. These circumstances can leave servers vulnerable to DNS attacks.

Protection against all major DNS attack types

- DNS amplification: This most common form of attack occurs when a DNS response is significantly larger than a query. For instance, if an "open resolver" receives a query packet containing a large buffer advertisement from a "spoofed" requesting IP address, its reply can be quite large. And that can result in a denial-of-service attack on the DNS server

- NXDomain: A DNS server is flooded with queries for a non-existent domain (a.k.a. NXDomain). The recursive server doesn't know the domain doesn't exist until it receives responses from the queries it initiates. The process consumes valuable server resources and overloads the cache. As a result, legitimate DNS queries are dropped or delayed.

- Phantom domain: Attackers set up phantom domains that don't respond to DNS queries. A recursive DNS server is forced to wait for responses, which consumes server resources, resulting in delayed or dropped responses.

- Random sub-domain: Randomly generated attacks target sub-domains on a legitimate domain. To resolve the domains, a recursive DNS server spawns concurrent queries that inevitably hit the limit, while authoritative DNS servers experience DoS.

- Look-up domain: Attackers set up domains that establish TCP-based connections with a recursive

DNS server and keep the connections alive with random responses. The server is tied up and eventually exhausts its resources.

How it works

Nexusguard fetches zone records from your DNS servers and caches them in our globally distributed scrubbing centers. Your DNS queries are forwarded to the Nexusguard's cloud and resolved by our DNS caches. Consequently, we absorb all DNS attacks and your DNS servers never need to respond to any malicious DNS query.

Thanks to the advanced filtering technologies in DNS Protection, you need only point your DNS name server to Nexusguard's DNS virtual IP, and the service will filter malicious traffic and respond to all client queries. DNS Protection also supports GRE tunneling if you wish to retain control of your own DNS IPs. The tunnel solution is stable and reduces the risk of a real IP being attacked.

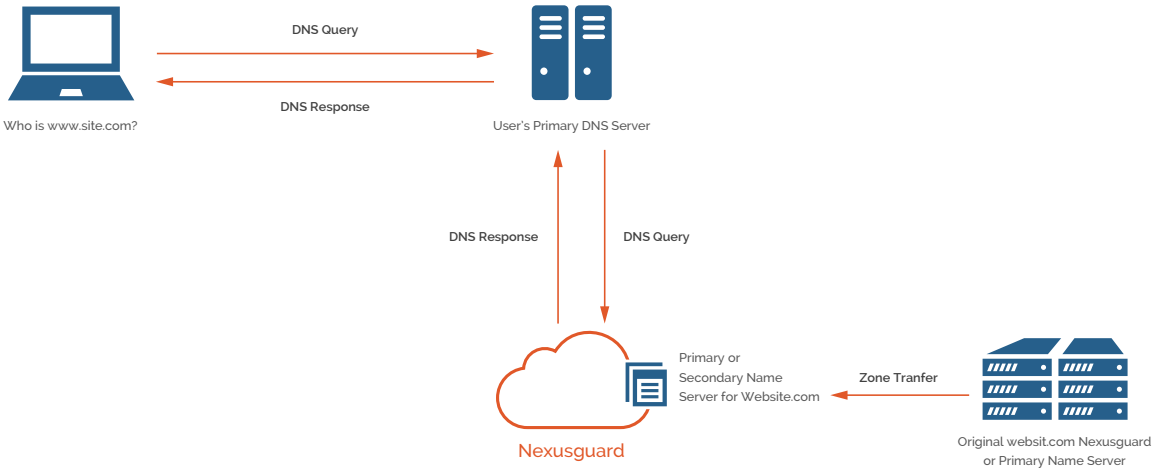

Global DNS attack mitigation capacity: 26.5 million qps (as of July 2015)

Content and networking optimization

Overview

With the explosive utilization of multimedia applications and internet-based digital media distribution, businesses face the challenge of presenting a seamless interactive media experience to their customers. In order to maintain and improve their informative websites or eCommerce activities, resources and infrastructure are put to the test.

With Nexusguard's globally distributed data centers, you can benefit from lower latency and higher capacity with the caching and load balancing capabilities. During peacetime, the dynamic and static content caching mechanism makes the most of the caching capacity and offloads the server from excessive HTTP requests.

All traffic going through the cloud is compressed and cached, which is effective in speeding up the delivery of high-traffic websites or online services access by users around the world. In other words, the closer the CDN server is to the user geographically, the faster the content will be delivered to the user.

The excessive capacity also serves as the last layer of protection to absorb the final bit of attack traffic, if any, that has slipped through the preceding layers. The load-sharing traffic services support multiple backend configurations. Automatic, backend failover is also implemented in the event of a backend server failure.

Custom caching rules

A "Purge Cache" option is also available, which enables you to purge your entire site or a specific resource on our server to enable immediate update of new content in the cache (e.g., redesign of site, page, etc.). Users can also create custom caching rules to explicitly control caching per URL and resource types, all of which can be managed on the Customer Portal.

Visibility and control

Delivered through a dedicated Customer Portal, we deliver full network traffic visibility, advanced analytics and real-time threat detection for all traffic into your network. This enables you to identify, detect and respond to threats swiftly. A comprehensive report of all activity lets you conduct forensic investigations, perform proactive incident response and even make informed business decisions.

It harnesses and combines the power of Big Data with visual and analytic technologies to deliver real-time meaningful business and security intelligence to empower business leaders to make informed decisions. It monitors your website from three perspectives: business status, threat protection and service performance.

The intuiative, comprehenisve Customer Portal boasts the following features and benefits:

Protection—Live Attack Mitigation Map & Dashboard

- **Threat Radar** - Displays a live map of the number and source of current threats and attacks.
- **Protection Trends** - Presents historical views of past threats versus legitimate requests and bandwidth use, providing IT security with the intelligence needed to formulate dynamic strategy implementations.
- **Threat Analytics** – A detailed understanding of the attack techniques used against particular web assets provides insight into potentially existing bottlenecks and weaknesses in application infrastructures.

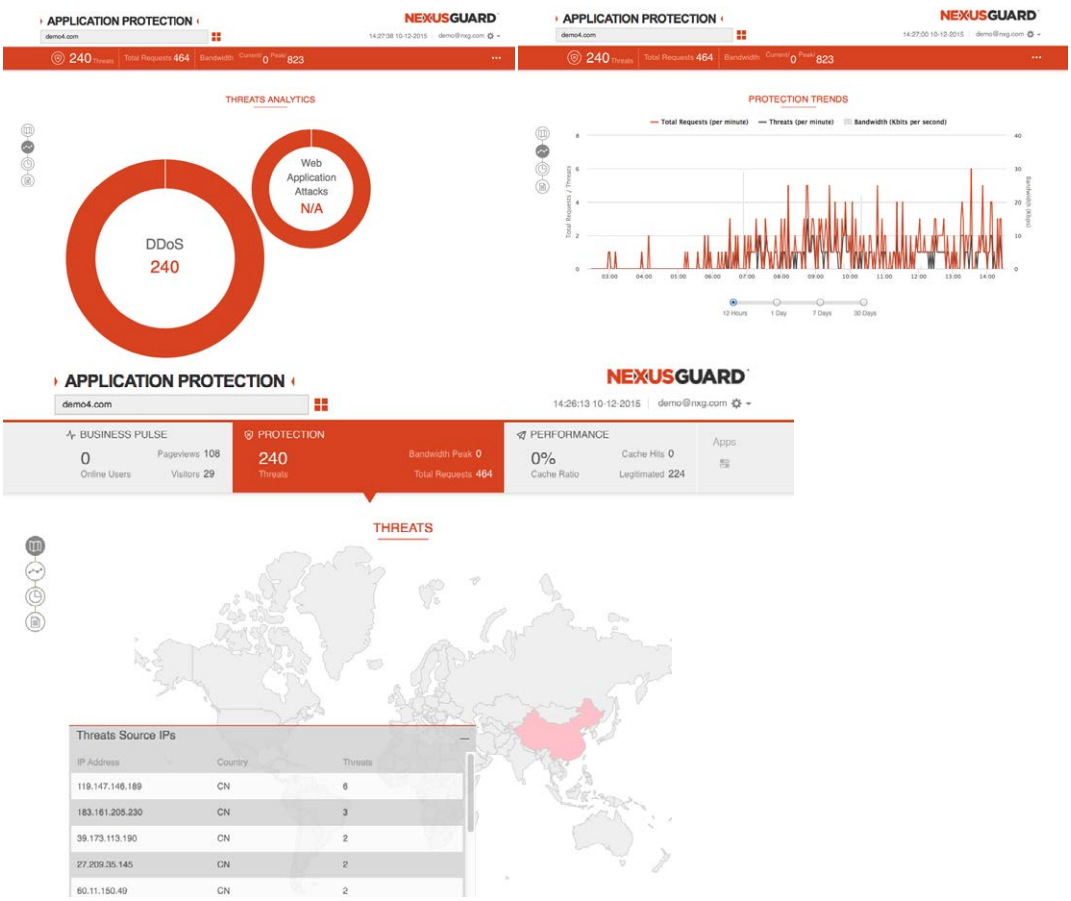

Business Pulse —Visualized Web Analytics

This module provides business related data and analytics to provide users with a better transparency of their business activities.

- **Live Visitor Information** – The number of browsing users online, including their individual IP address and location, the time it takes for pages to load for each visitor, and the resources each individual is accessing. Provides true visibility of business status, especially in the event of an attack.
- **Website Trend** – Graphic representations and historical views of the number of visitors and pages viewed allow managers to better predict visitor trends and make any necessary changes.
- **Website Analytics** - The top three most popular content sources are displayed along with other useful website-related statistics such as browser breakdowns, search engine referrals, and more.

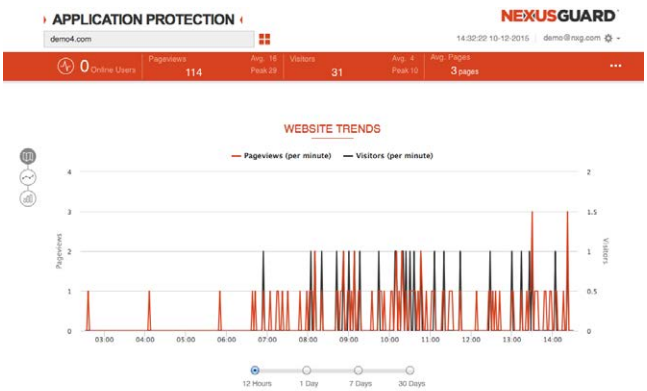

Performance—Web & Service Performance

This module displays site performance and optimization information.

- **Threat Radar** - Displays a live map of the number and source of current threats and attacks.
- **Data Center Status** - Scrubbing Centers provide deep visibility into the operational status of web services.
- **Caching Performance** – An historical view and understanding of caching performance enables enhanced website and traffic performance optimization.

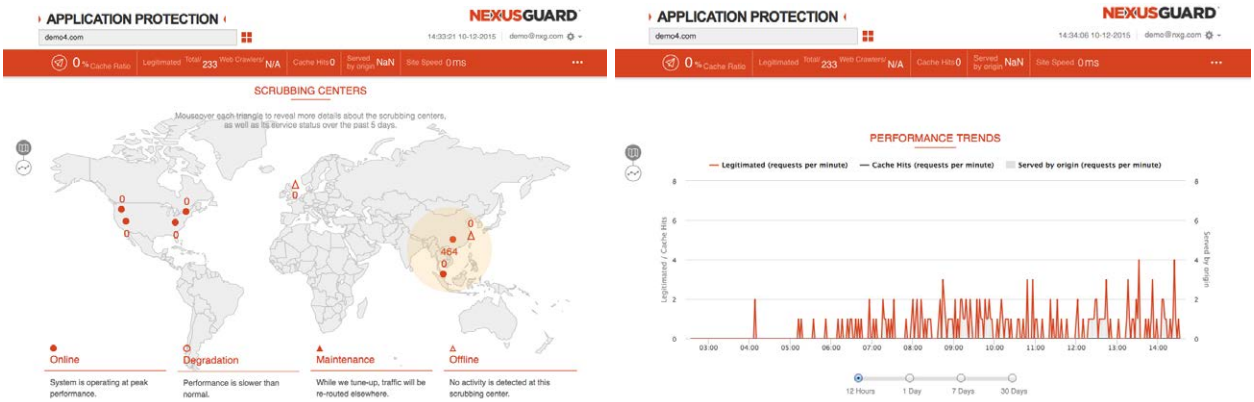

24x7 SOCs staffed with DDoS experts

Solely relying on automated filtering tools and large bandwidth for DDoS mitigation is not enough. At Nexusguard, we have two Security Operations Centers (SOCs) deployed in Asia and Americas staffed with security experts to monitor and respond to attacks and threats around the clock while providing seamless support to our customers.

Their experiences, skills and availability are essential part of our comprehensive mitigation mechanism, all being combined can flexibly and timely respond to the changing techniques used by attackers. Each SOC provides you with the best DDoS protection at all times, with local language support.

A "Purge Cache" option is also available, which enables you to purge your entire site or a specific resource on our server to enable immediate update of new content in the cache (e.g., redesign of site, page, etc.). Users can also create custom caching rules to explicitly control caching per URL and resource types, all of which can be managed on the Customer Portal.

Conclusion

For any organization relying on uninterrupted online presence, DDoS attack cause direct revenue loss, customer loss, increase expenses, and can damage their reputation overnight. Attackers nowadays employ multiple hacking and attack techniques to commit more severe cyber-crimes.

Nexusguard's integrated DDoS protection solution is designed to deliver a perfect balance of protection and performance for your websites, web applications and infrastructure. Our multi-layered defense system provides the shortest time to mitigation, stopping multi-vulnerabilities attacks instantly.

A new breed of enterprise-grade defense, our solution, supported by a massive backbone infrastructure and cutting-edge innovations, delivers:

- A robust, cloud-based network capable of protecting against the largest bandwidth and resource flood-based attacks.
- Heuristic-based, intelligent, and accurate detection and mitigation of application-based attacks.
- Static and dynamic content caching and acceleration to optimize web application performance and the user experience.
- Proprietary, progressive challenge-response mechanism for high mitigation accuracy and non-disruptive user experience.
- 100% uptime guarantee; shortest mitigation response time; always-on protection.
- Single point of contact with dedicated Technical Account Manager.
- Real-time dashboards with visually appealing maps, charts and graphs, showing live traffic, attack and threat statistics, detailed event logs, visitor demographics and many other useful analytics.

White Paper

NEXUSGUARD

Founded in 2008, Nexusguard is the global leader in fighting malicious internet attacks. Nexusguard protects clients against a multitude of threats, including distributed denial of service (DDoS) attacks, to ensure uninterrupted internet service. Nexusguard provides comprehensive, highly customized solutions for customers of all sizes, across a range of industries, and also enables turnkey anti-DDoS solutions for service providers. Nexusguard delivers on its promise to maximize peace of mind by minimizing threats. Headquartered in San Francisco, Nexusguard's network of security experts extends globally.

Twitter twitter.com/nexusguard
Facebook facebook.com/NXG.PR
LinkedIn linkedin.com/company/nexusguard

nexusguard.com
contact@nexusguard.com

20151218-EN-A4