NETSCOUT

| WHITE PAPER |

As Enterprises Increasingly Embrace Edge Computing

Two Use Cases Underscore The Need For Intelligent, Automated Edge Protection

SECURITY

TABLE OF CONTENTS

As Enterprises Increasingly Embrace Edge Computing, Two Use Cases Underscore The Need For Intelligent, Automated Edge Protection

While much of the world's population continues to strive toward life returning to pre-pandemic form, enterprises are increasingly acknowledging that the changes wrought by the pandemic are permanent – especially in terms of enterprise IT and security.

Studies show that the pandemic accelerated the rate of digital transformation for enterprises by about five years in a matter of eight weeks. Moreover, 97 percent of IT leaders see the digital initiatives they started in 2020 continuing into 2022 and beyond. Doing so requires organizations to move away from traditional networking technologies to maintain a high-quality user experience for employees, vendors, customers and others supported by the enterprise network.

One of the biggest changes that enterprise networks have undergone is the move away from traditional hub-and-spoke architectures to edge computing, a distributed architecture in which processing and data storage happen closer to the data source. The massive increase in the number of devices that organizations support is driving the need for near instantaneous data transfers to and from those devices. But modern applications and services can't absorb the latencies involved in sending and receiving data between end devices and a data center somewhere in the cloud.

Moving resources to the edge reduces network latency and increases speed of data communications over the network. Edge computing enables enterprises to increase network performance by reducing the need to send captured data from the network periphery back to a central system. Almost half of enterprise IT leaders say edge computing reduces operational expenses by reducing reliance on costly bandwidth to connect locations, as well as by reducing data redundancy.

The benefits are so compelling that 90 percent of industrial enterprises will employ edge computing by 2022. Likewise, more than half of enterprise-generated data will be created and processed at the edge by 2023 – up from less than 10 percent in 2019. In fact, 36 percent of enterprises already have edge computing tools in production or are piloting initiatives, while another 25 percent are actively researching the technology. By the end of 2023, more than half of all large enterprises will use edge computing for six or more use cases.

But for all of the benefits that the edge provides for enterprises, it also introduces more complications in terms of the risk that has to be assessed and managed. The scale and exposure of threats increases significantly at the edge, where the attack surface is expanded and exposure to threats is increased. Those threats are clearly evident in terms of distributed denial of service (DDoS) campaigns, data theft and leaks, third-party vulnerabilities, and intrusions into the enterprise network.

Not surprisingly, 47 percent of enterprise IT leaders say that investing in security and privacy are a top focus for technology investment. In particular, there are two use cases that best demonstrate why enterprises need intelligent, automated protection at the edge.

Almost half of enterprise IT leaders say edge computing reduces operational expenses

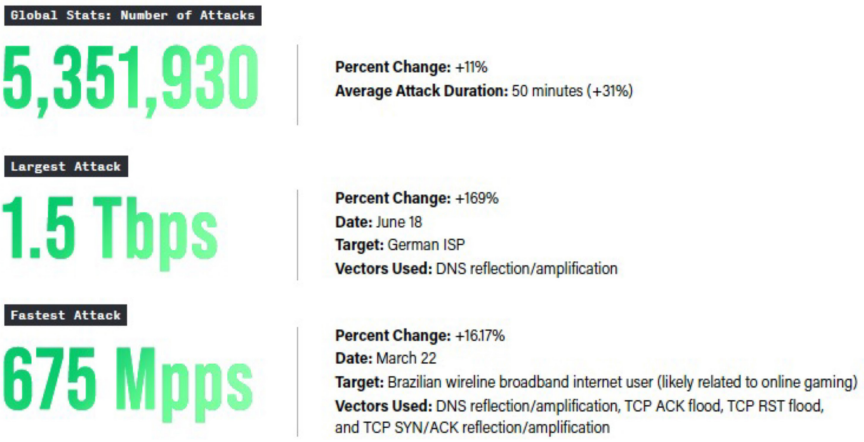

Figure 1.

The scale and exposure of threats increases significantly at the edge

Use Case 1: Inbound DDoS Attacks

Cyberattackers haven't had to change much in their bag of tricks to effectively threaten enterprises as they move to the edge. Instead, attackers have simply applied more focus and increased the number of attacks they've long used – especially DDoS attacks.

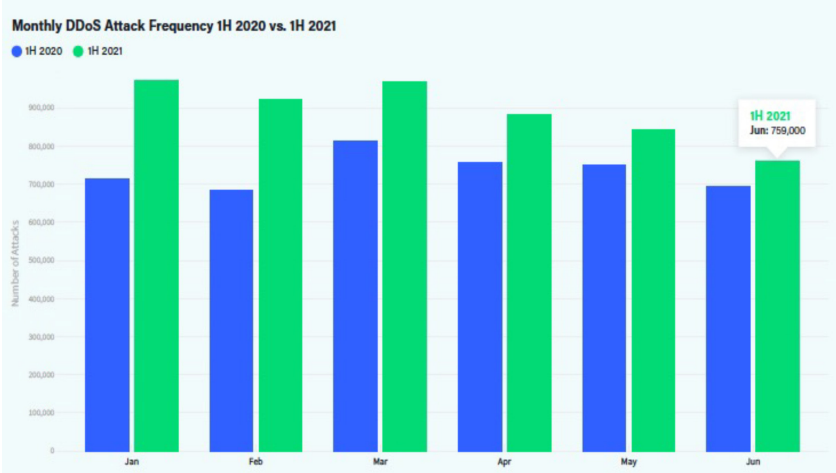

Figure 2.

Consider, for instance, that almost 5.4 million DDoS attacks were waged in the first half of 2021 – an 11 percent increase from the same period in 2020, as shown in **Figure 1**.

If that growth trajectory continues – and there's no reason to believe that it won't continue when looking at attack frequencies in **Figure 2** – enterprise IT and security will face a record-breaking 11 million DDoS attacks during 2021.

While a great deal of focus is placed on stopping large DDoS attacks, protecting the edge requires emphasis on smaller attacks – especially application layer attacks.

Application-layer attacks are usually low-to-mid volume since they have to conform to the protocol the application is using, which often involves protocol handshakes and protocol/application compliance. As such, these DDoS attacks are primarily launched using discrete intelligent clients and can't be spoofed. Such attacks target specific vulnerabilities or issues within an application, resulting in the application not being able to deliver content to the user.

To better understand application layer attacks, it helps to examine three examples:

- Slowloris is an application layer DDoS attack that uses partial HTTP requests to open connections between a single computer and a targeted web server. The goal is to keep those connections open for as long as possible to overwhelm and slow the target.
- In a Slow Post DDoS attack, the attacker sends legitimate HTTP POST headers to a web server. In the headers, the sizes of the message body that will follow are correctly specified. However, the message body is sent at a painfully low speed – sometimes as slow as one byte every two minutes. Since the message is handled normally, the targeted server will do its best to follow specified protocol rules, resulting in the server subsequently slowing to a crawl.
- TCP state exhaustion attacks attempt to consume the connection state tables that are present in many infrastructure components, including load balancers, firewalls and application servers themselves. Even high capacity devices capable of maintaining state on millions of connections can be taken down by these attacks.

Determined attackers use these types of attacks to monitor the results of their efforts. When they ascertain vulnerabilities, they make modifications to overcome the network's defenses. These attacks also create problems for enterprise IT and security, as maintaining an ongoing list of known attack patterns is impractical due to scale issues and the rate at which such lists have to be updated. Likewise, it's unwise to maintain a long-lived set of payload patterns because they also have a high risk of causing collateral damage.

When looking at DDoS trends over time, attacks are cyclical in nature. Attackers develop new DDoS attack types and vectors that are then used to launch a new wave of attacks. As defenders become more proficient in stopping these new DDoS attacks, the attackers develop new types of attacks, and the cycle repeats itself.

Successfully detecting and stopping application layer attacks requires intelligent DDoS attack protection at the network edge to maintain the availability of business critical applications.

Use Case 2: Compromised Devices

The second use case centers around compromised devices, the best example of which is found in the expanding Internet of Things (IoT) infrastructure and how attackers are using botnets against it.

The global enterprise IoT market is expected to reach $58 billion by 2023, a 26 percent compound annual growth rate (CAGR). The expectation is that there will be more than 15 billion IoT devices connecting to enterprise infrastructure by 2029.

Not surprisingly, 60 percent of enterprise executives believe the IoT will play an important role in the digital business strategies of their company. They cite four top benefits of the IoT, including gaining competitive advantage (47 percent), creating new business models (43 percent), meeting changing customer expectations (34 percent) and improving quality (33 percent).

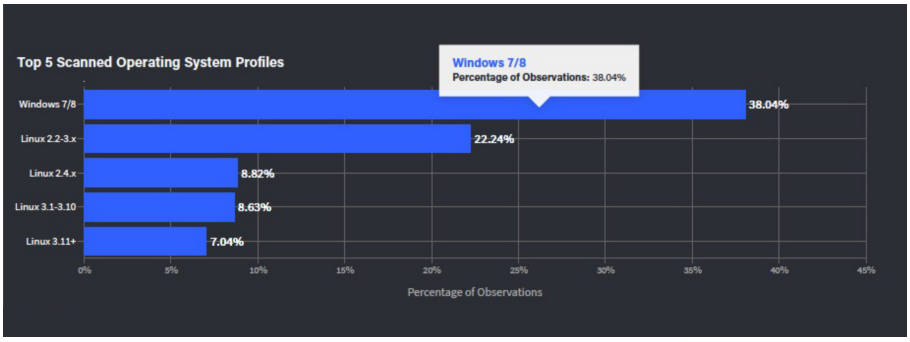

Figure 3.

Successfully detecting and stopping application layer attacks requires intelligent DDoS attack protection at the network edge to maintain the availability of business critical applications.

But the massive influx of these devices creates a massive headache for corporate IT and security teams. In fact, it's not uncommon for IT organizations to find IoT devices on their networks that they did not install, secure or manage. Such devices can be hacked in minutes, but these breaches can take months to discover.

The biggest threat against IoT devices stems from botnets that harness vulnerable devices for DDoS attacks. Two of the most prolific IoT botnets are Gafgyt and Mirai, which accounted for more than half of the total number of DDoS attacks that occurred in 1H 2021.

These IoT-based botnets have been used to launch all types of DDoS attacks on a worldwide basis. Because of their size, these botnets are used to launch volumetric attacks, which are designed to overwhelm internal network capacity with significantly high volumes of malicious traffic. These DDoS attacks attempt to consume the bandwidth either within the target network/service or between the target network/service and the rest of the Internet.

Attackers also have combined compromised IoT bots with application layer attacks. For example, an IoT-based botnet consisting of compromised Mikrotik routers was used to launch several high-profile HTTP and HTTP/S-based application-layer DDoS attacks.

Studying data from honeypot networks provides a clearer understanding of how botnets are used by attackers against enterprises. Botnets that are tracked in this way typically have a common device profile, which reveals the most common operating systems (OSs) that bots use to propagate attacks.

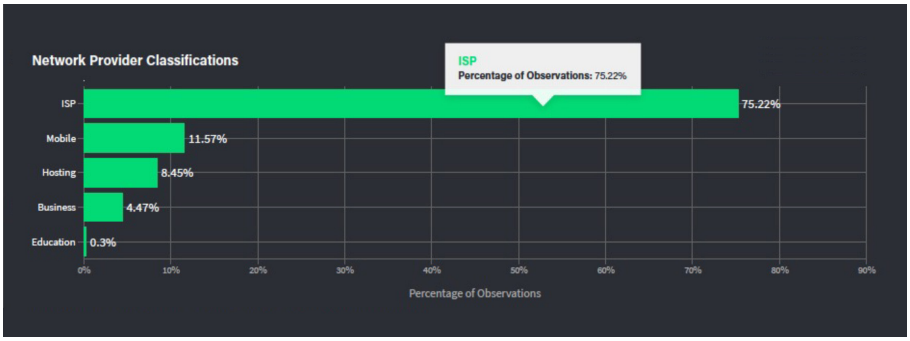

Figure 4.

As shown in **Figure 3**, the top three OS profiles are Windows 7/8, Linux 2.2-3x, and Linux 2.4x. By using this information to determine the kinds of devices on the network, IT and security can then match those devices against common username and password combinations to ensure they aren't susceptible to brute-force attacks.

By gaining insight into the device profile and operating system type, it's then possible to break down the network types to ascertain where botted nodes reside. As shown in **Figure 4**, the top three source network profiles are ISP, mobile, and hosting, where device control is nearly nonexistent. That lack of control means those ISP and mobile numbers really represent compromised subscribers.

IoT-based botnets have been used to launch all types of DDoS attacks on a worldwide basis.

How To Combat Both Use Cases

The proliferation of these two use cases against enterprise networks requires that IT and security personnel to apply defender automation skills that match those of attackers. Moreover, protecting against such attacks requires the ability to look at outbound traffic in order to identify internally compromised devices and indicators of compromise (IoCs). That translates to a need to view defense from both the outside in, as well as from the inside out.

It's also important to remember that the network edge includes private and public clouds, partner networks and remote users. As such, blocking traffic at the network edge requires complete confidence that you're not blocking legitimate traffic. Ensuring that the edge is protected without blocking legitimate traffic requires a solution that features:

- **Next-generation network edge cybersecurity stack:** The solution should be fronted with stateless threat detection and mitigation technology designed to protect the stateful network cybersecurity stack itself, as well as the network and services behind it.
- **Ability to conduct rapid, highly contextual cyberthreat investigations:** that enable the confidence to block at the network edge.
- **Automatic blocking of inbound DDoS attacks:** Specifically, the solution should block TCP state exhaustion attacks that threaten the availability of stateful devices like firewalls, VPN concentrators, and load balancers.
- **Automatic blocking of outbound IOCs:** This should stop compromised internal devices from communicating with outside command and control (C&C) infrastructure that has been missed by the firewall or existing cybersecurity stack.
- **Stateless packet processing technology:** This protects the network, services, and stateful network cybersecurity stack from threats.
- **Integration into existing cybersecurity stack and processes:** To achieve more effective threat detection and response.

LEARN MORE

Learn more about how to ensure intelligent, automated edge protection or contact us today:

www.netscout.com/solutions/omnis-smart-edge-protection

NETSCOUT

Corporate Headquarters
NETSCOUT Systems, Inc.
Westford, MA 01886-4105
Phone: +1 978-614-4000
www.netscout.com

Sales Information
Toll Free US: 800-309-4804
(International numbers below)

Product Support
Toll Free US: 888-357-7667
(International numbers below)

NETSCOUT offers sales, support, and services in over 32 countries. Global addresses, and international numbers are listed on the NETSCOUT website at: www.netscout.com/company/contact-us

SECCWP_025_EN-2101 10/2021