Solution Showcase

Cybersecurity at the Edge of the Internet

Date: August 2017 **Author:** Jack Poller, ESG Analyst

Abstract: Cloud and content providers distributing data across the Internet face a growing threat landscape. Recent escalations in DDoS attacks have corralled an army of Internet-connected devices to generate aggregate bandwidth of 1 Tbps, and have taken down major Internet sites. EdgeConneX, with over 30 "edge" colocation facilities, can help service providers to counter cyber attackers and ensure availability and resiliency through redundancy.

Overview

Content distribution networks (CDNs), Internet service providers (ISPs), cloud service providers (CSPs), cloud applications, and others have always had to worry about cybersecurity. Why? Because it is very easy and effective for cybercriminals to use distributed denial of service (DDoS) to attack service providers.

Like most other attacks, DDoS has evolved over time, resulting in a treacherous threat landscape. While the number and severity of attacks has risen every year, this past year has seen the rise of mega attacks targeting major sites. Twelve attacks have been recorded achieving greater than 100Gbps throughput, with five of those exceeding 200Gbps. Arguably the most well-known and largest attacks occurred in the fall of 2016, and were executed by the Mirai botnet, a collection of more than 100,000 compromised Internet-connected video cameras, consumer routers, and other devices.

At least 37 attacks have been attributed to Mirai, with 70% US-based targets and an average throughput of 57 Gbps. A Mirai mega attack achieved 620Gbps throughput, and shortly thereafter, in a separate attack, 1Tbps. Mirai was also used to attack the Republic of Liberia's Internet infrastructure, as well as US DNS service provider Dyn, resulting in the inaccessibility of several high-profile websites such as GitHub, Twitter, Reddit, Netflix, Airbnb, Spotify, and many others.

DDoS has been used to target media and entertainment, gaming, software, technology, banking, and other major industries. Globally, almost three-quarters (73%) of organizations suffered a DDoS attack, and 85% of attacked organizations reported multiple assaults. It can take an average of three hours to detect and start to repel a DDoS attack. Almost half (49%) of all targets could lose more than $100,000 per hour, with one-third (33%) exposed to losing more than $250,000 per hour. More than half of all targets also suffered a cybersecurity breach while under DDoS attack.[1]

Are DDoS attacks the only concern? Cybercriminals have many tools in their attack toolbox, and their targets include:

- DNS services. One in five DDoS attacks were DNS-based; the Mirai attack against Dyn directly targeted Dyn's DNS service. DNS services can also be interrupted with spoofing, also known as cache poisoning, where corrupt DNS data is used to divert Internet traffic away from the correct servers.

[1] Source: Neustar, Worldwide DDoS Attacks & Protection Report, October 2016.

This ESG Solution Showcase was commissioned by EdgeConneX and is distributed under license from ESG.
© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Ironically, DNS servers can themselves be used to generate DDoS traffic with DNS amplification. The cybercriminal sends a DNS query with a forged IP address. The DNS server responds to the forged IP address belonging to the target of the attack. Thus, small queries result in large responses that can overwhelm the target of the attack.

- HTTP/S. HTTP and HTTPS are the lingua franca of the Internet, and are a common cybercriminal target. In addition to zero-day vulnerabilities (a compromise that becomes active before it is publicly reported, giving the target zero days to respond), HTTP servers can suffer from flood attacks, whereby the attacker crafts an HTTP POST request with parameters that trigger complex server-side processing. The goal is to exhaust server-side resources.

  HTTPS uses encryption to protect end-to-end communication. However, older encryption schemes can be vulnerable to exploitation of protocol weaknesses and brute force attacks with modern computers. SSLv2, a cryptographic scheme that dates from the 1990s, is now known to be badly insecure, yet is often still used.

- Servers and applications. Server operating systems and applications are getting more complex to meet the need for more features and functionality. It is virtually impossible to have 100% security with hundreds of millions of lines of code, and 2016 saw a 34% increase in the number of reported server vulnerabilities.

  Because of the downtime risk, administrators are often loath to patch operating systems and applications to address security. A vulnerability in JBoss, a common middleware utility, was patched in 2010. Yet there are still more than 3.2 million servers with unpatched and vulnerable JBoss implementations. This vulnerability was used by the SamSam ransomware malware that specifically targets the servers and systems used to run hospitals.

Organizations are trying to cope with the threats posed by cybercriminals while operating within the constraints posed by the global cybersecurity skills shortage. According to ESG research, 45% of organizations report that they have a problematic shortage of cybersecurity skills and 32% lack the appropriate IT architecture/planning skills that are critical in coping with threats posed by cybercriminals (see Figure 1).[2]

Figure 1. Top-ten Areas of IT Skills Shortage

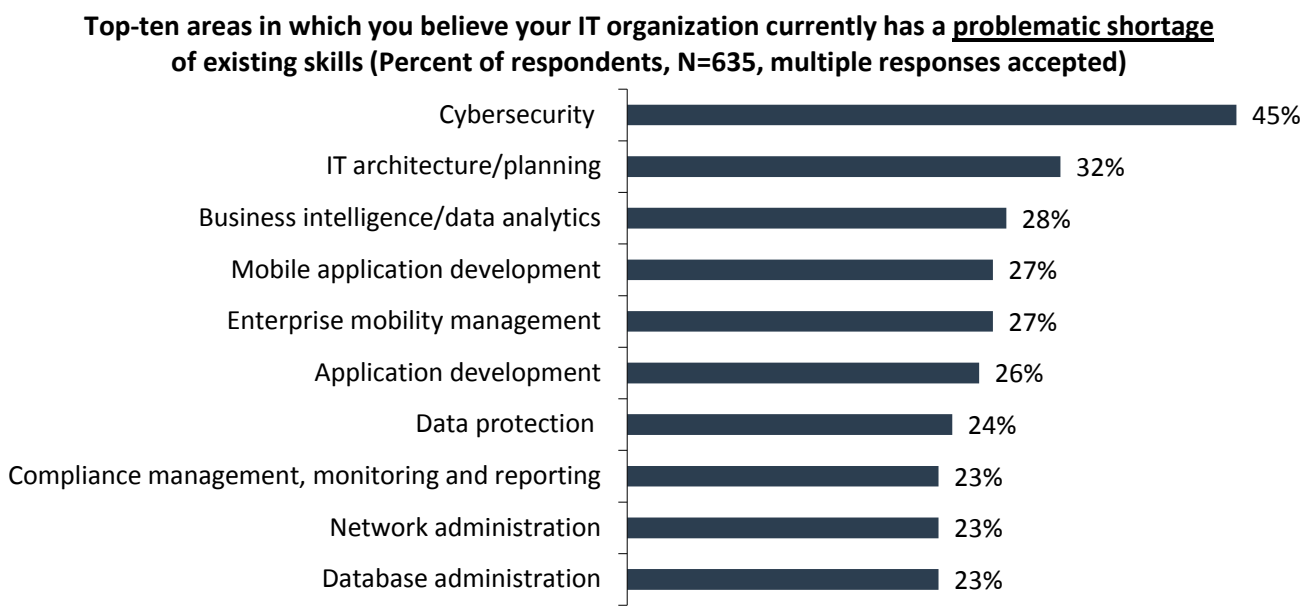

Source: Enterprise Strategy Group, 2017

[2] Source: ESG Brief, *2017 Cybersecurity Spending Trends*, March 2017.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

What's Needed?

Organizations continue to invest in the latest developments in IT, hoping to realize their potential value in combatting cybercriminals. A paradigm guiding cybersecurity and system architects is confidentiality, integrity, and availability, known as the CIA triad. Confidentiality reflects the need to limit access to data (need to know), integrity ensures that the data is accurate and trustworthy, and availability is the guarantee of reliable access to the data by authorized people.

System architects focusing on the CIA triad should consider the following when evolving their network architectures:

- Resiliency and availability through redundancy. Organizations need to create redundant data centers spread across multiple geographies and multiple network providers. This redundancy forces attackers to spread their DDoS bandwidth across all data centers and networks, lessening the load each data center must handle. Should an attacker focus their attacks on one specific data center or network provider, cybersecurity professionals can reroute users to geographically proximate data centers and networks.

  Redundancy focuses attackers on the edge, where requests can be inspected for malicious payloads. Cross-site-scripting, SQL injections, or other attacks can be filtered before they can infiltrate the origin server. With more redundancy, the architecture is better able to circumvent attackers' objectives. This ensures availability of services, minimizes the effects of attacks on service latency, and helps to maintain connections, service uptime, and service level agreements (SLAs).

- Redundancy promotes rapid detection and response. Service providers need to inspect traffic to separate the good from the bad. Correlating traffic metadata—traffic type, source, destination, packet timing, and more—can help an organization to quickly detect an attack. Once identified, cybersecurity professionals can decide on the appropriate response, from filtering bad traffic to shunting users to other proximate sites. Rapid detection and response limits potential damage and reduces the risk of a breach attempt hidden inside a massive DDoS attack.

- Standardization reduces the attack surface. A standard environment reduces the number and variety of software and systems that must be secured, and minimizes the workload, especially critical with today's cybersecurity skills shortage. With standard components, the cybersecurity team can focus on ensuring systems are always up-to-date, and that security holes are patched as soon as they are identified.

- Redundant ancillary services provide availability. System architects need to apply the same focus to their ancillary services as they do to their primary services. A CDN or cloud service may have 100% uptime but will provide an unsatisfactory user experience without a functioning DNS system. After all, you can't use a service if you can't find its address. Architects need to protect their DNS servers with redundancy, ensuring DNS availability in the face of attack. The DNS system should also follow the policy of least privileged access, guaranteeing that only those with the need are authorized and allowed to make changes.

Resiliency and availability through redundancy can be used by CDNs, ISPs, and CSPs for risk mitigation and accelerated risk detection and response, as well as to reduce the workload for the cybersecurity team.

Redundancy, Availability, and EdgeConneX

Redundancy is a key component for ensuring availability of services, and one third of the CIA triad. CDNs, ISPs, and CSPs leverage redundancy to alleviate cybersecurity risks and to speed time to incident detection and response. These service providers also need to ensure that their data centers have the requisite power capacity, system density, and connectivity to house and distribute content to local ISPs throughout a metro area.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

This is the exact mission of EdgeConneX, a Herndon, Virginia-based "Internet of Everywhere" company. Their Edge Data Centers® serve metro areas by directly connecting into the local ISPs, and are mostly located outside traditional peering hubs, easily enabling multiple deployments in a distributed environment.

CDNs such as Akamai, web performance and security companies such as Cloudflare, CSPs, ISPs, and other organizations operating at the edge of the Internet can and do use multiple deployments with EdgeConneX to address:

- DDoS attack mitigation. EdgeConneX Edge Data Centers provide physically secure colocation space. Each facility has all the requisite redundant power and cooling, and access/peering agreements with multiple backbone network providers. With over 30 locations, EdgeConneX enables CDNs, ISPs, and CSPs to redundantly disperse services throughout the continental United States and Europe. Service providers leverage both physical redundancy and network redundancy, enabling load sharing and alternate routes and services when faced with high-bandwidth DDoS attacks.

- Ancillary service protection. Rather than rely on centralized servers for ancillary services such as DNS, authorization, and authentication, organizations can place servers in each Edge Data Center. This redundancy helps to ensure DNS, authorization and authentication, and other ancillary service availability, without which customers cannot use the primary services of CDNs, ISPs, and CSPs.

- Attack detection and response. Extending beyond traditional peering hubs, EdgeConneX network partners provide interconnection between Edge Data Centers. Aggregating traffic monitoring across all locations helps organizations to rapidly identify DDoS and other cybersecurity attacks. Once detected, data center throughput information can be used to balance the load or transition users away from a specific data center suffering an attack.

Given its capabilities, EdgeConneX continues to gain traction with service providers and CDNs globally. Architects seeking to leverage redundancy for cybersecurity would be well served to assess how EdgeConneX can help.

The Bigger Truth

Cybercriminals have always targeted CDNs, ISPs, and CSPs. This past year has seen a drastic rise in the intensity of attacks, with mega attacks now attaining 1 Tbps, more than most service providers can withstand on their own.

Organizations moving to the edge of the Internet should look to focus on availability through redundancy as a key component in their cybersecurity and system architecture. EdgeConneX's colocation facilities can help service providers to implement redundancy in their primary and ancillary services quickly and easily. EdgeConneX's customers are thus able to mitigate the risks of DDoS attacks, enhancing end-user quality of experience, and maximizing return on investment.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.