


2019 QRATOR LABS

CYBER SECURITY REPORT





CONTENT

4 I Closely watched events of 2019

- 4 TCP SYN-ACK Reflection/Amplification and other protocols
- 7 *DNS-over-HTTPS versus DNS-over-TLS*
- 7 *TCP-acknowledgement bug*
- 7 BGP Optimizers
- 11 Banks
- 12 E-commerce

13 II Re-Optimizing the Future of BGP Routing

- 16 *Qrator Labs presence map*

16 III Current and Future Development of Qrator Filtering Network

- 16 *Processing Logic*
- 17 *HTTP/2*
- 18 *Containers and Orchestration*
- 19 *Yandex.Cloud and Ingress 2.0*
- 20 *TLS 1.3 and ECDSA certificates*
- 21 *Smaller Upgrades*
- 22 *IPv6*
- 22 *Antibot*
- 23 *Hardware*
- 24 *Future Updates to the Network*

TCP SYN-ACK reflection/
amplification:

1X-5X FACTOR

SERVERS.COM
UNDER ATTACK FOR

3 DAYS

THE ATTACK MAINLY
CONSISTED OF THE

CLDAP

amplification
(with a significant
portion of fragmented

UDP datagrams)

AND THE

SYN/ACK

amplification traffic,
with other sorts of UDP
amplification present periodically.

The SYN/ACK amplification
traffic peaked at around

**208 MILLIONS
OF PACKETS
PER SECOND.**

2019 BGP optimizing incidents: three separate.

MOST SEVERE:

June 24, DQE - Allegheny - Verizon - Cloudflare route leak, resulting in losing

15%

OF CLOUDFLARE'
GLOBAL TRAFFIC.

TWO HOURS

before the route leak was fixed.

LONGEST ONE:

MARCH TO APRIL 2019

TRAFFIC

ATTRACTION

attacks with more-specific.
We saw AS263444 taking
prefixes from routes,

SPLITTING

IT INTO TWO PARTS

and announcing the route with the same
AS_PATH for two new prefixes.

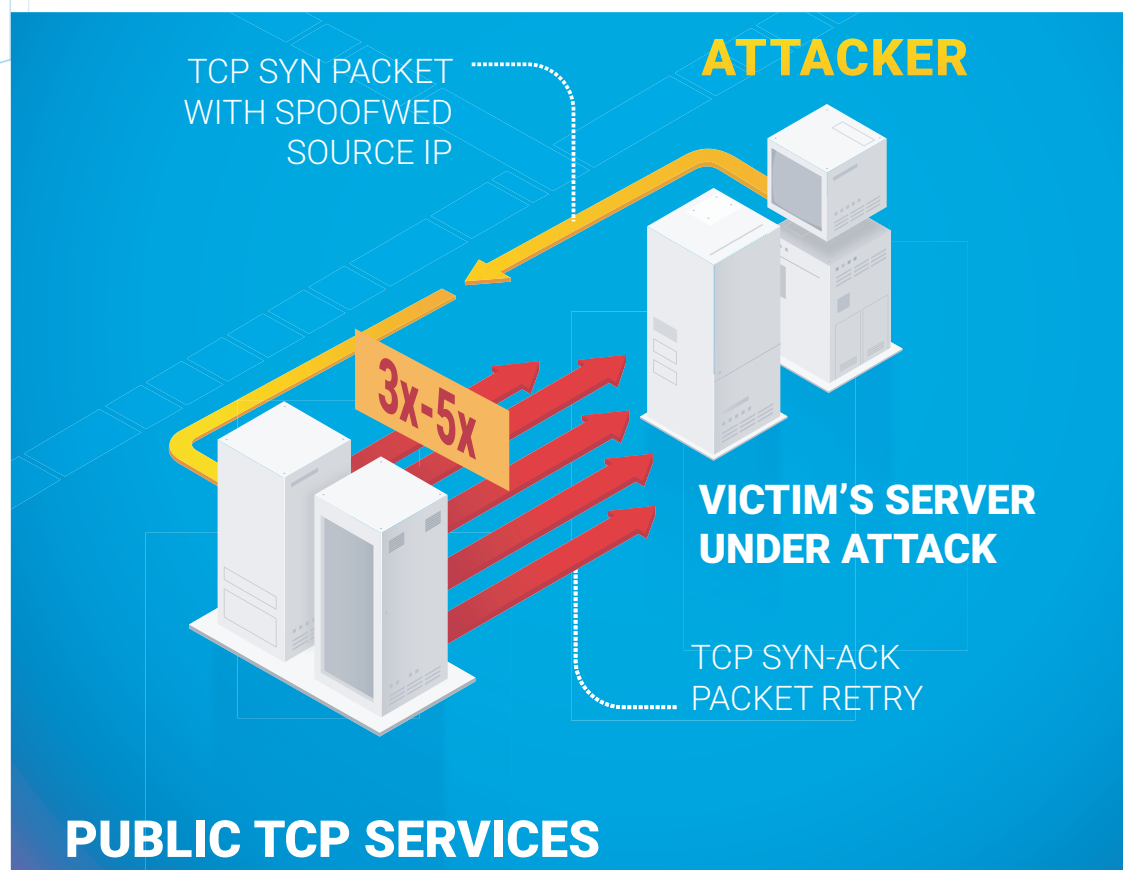
CLOSELY WATCHED EVENTS OF 2019

TCP SYN-ACK REFLECTION/AMPLIFICATION AND OTHER PROTOCOLS

Continuing growth of the IoT market means that, when misused, it could be exploited to generate significant attack bandwidth – as happened when the Web Services Dynamic Discovery protocol was employed to wreak havoc in the middle of the year. The Apple ARMS protocol that hit the news on October 19 with an amplification factor of nearly 35.5 was also visible in attacks on the Qrator Labs filtering network.

New findings were reported about amplifiers (PCAP), as well as recent at-

tacks exploiting a well-known TCP amplification vector: SYN-ACK. The main difference between this particular vector and a typical UDP amplification is that the amplification of the SYN-ACK vector does not send an answer that is multiple times bigger to a request – it only attempts to connect multiple times, thereby making a sizable amplification factor virtually. Because public clouds on the Internet answer spoofed requests, attacks employing SYN-ACK amplification vectors have become the most prominent danger on the Internet.




The hazard peaked when the big cloud hosting provider Servers.com turned to Qrator Labs for help with mitigating SYN-ACK amplified DDoS attacks.

It is quite interesting to note that the previously employed response in form of dropping all UDP traffic, which virtually mitigates most amplification attacks, doesn't help at all against the SYN-ACK amplification vector. Smaller independent Internet companies have much more difficulty dealing with such attacks which require more complex mitigation efforts.

TCP SYN-ACK amplification is not new, but it is not widely known as an attack vector. A malefactor sends SYN packets to all the public TCP services available on the Internet with a forged

source IP-address, and each server replies by trying to connect multiple - usually 3 to 5 - times. For a long time, this vector was not considered exploitable and only in 2019 did we see attackers become capable of generating enough bandwidth to overwhelm even large infrastructures. Interestingly, this new attack type is not about "amplification" itself, but simply the resending of packets in response to a delivery failure. For those seeking other vectors, we point to QUIC protocol, which is currently ripe for the same type of abuse, where QUIC-enabled servers answer client requests with responses that are much larger (IETF draft "recommends" sending a reply that is no more than three times bigger than the application).

UNIQUE AMPLIFIATORS, 2019




Amplification is no longer about 100x amplification – we see that 3-5x is enough. Confronting this problem involves eliminating the spoofed traffic as a category; there should be no way for anyone to imitate another endpoint and overwhelm it with traffic from legitimate content providers. BCP38 does not work – and developers of new transport protocols, like QUIC, do not appreciate the danger coming from even small amplification. They're on the other side.


Networks need a tool to blackhole or at least rate-limit spoofed traffic, and such a tool requires intelligence about the legitimate source of the request. The cloud networks belonging to large companies like Amazon, Akamai, Google and Azure represent ideal targets for amplification attacks – they have very efficient hardware capable of fulfilling all attacker needs.


Such attacks are also harsh to troubleshoot in the modern Internet, where, as we mentioned, frontend and backend applications and libraries are so profoundly interconnected. If you use open source software in your development stack for cloud hosting and you get assaulted with an SYN-ACK amplification from the same place – you are in big trouble. Repositories not working and configs not updating while you are being blocked (because of the spoofed requests) from communication with those repositories/containers – this a situation no one wants to find themselves in. We saw this multiple times in 2019 with companies pleading for help, seeing unimaginable critical dependencies for the first time in their history. Further development of the BGP protocol is needed to combat spoofing within the TCP/IP stack. Routing tables are

% CHANGE 2019 TO 2018

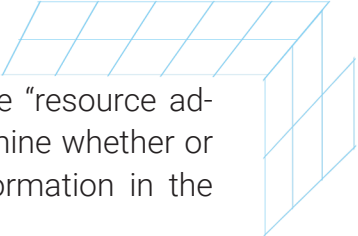


INDUSTRIES YoY DDoS ATTACKS DYNAMICS



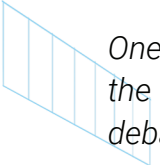


entirely different from prefix tables, and we need to enable the network to understand when a packet is illegitimate or malicious and disregard it – in other words, provide authentication at the network infrastructure level. Attention should not be focused on the “destina-




tion address” but on the “resource address” instead to determine whether or not it matches the information in the routing table.

DNS-OVER-HTTPS VERSUS DNS-OVER-TLS

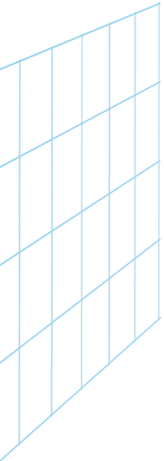


One of the hottest topics of 2019 was the DNS-over-HTTPS vs DNS-over-TLS debate. While controversial on the server side due to the legislative differences (EU's GDPR versus US' state and federal laws), it is already being supported by developers of the main browsers: Google Chrome and Mozilla Firefox. We could also speculate whether or not DoH




and DoT adoption lowered the number of amplifiers over the Internet. However, it is much more likely with the over-TLS version due to persistent connections between servers serving DNS queries. On everything else we have to wait for the market to decide.

TCP-ACKNOWLEDGEMENT BUG




In June 2019 Netflix identified several TCP networking vulnerabilities in FreeBSD and Linux kernels. The vulnerabilities specifically relate to the Maximum Segment Size (MSS) and TCP Selective Acknowledgement (SACK) capabilities. The most serious, dubbed “SACK Panic,” allows a remotely-triggered kernel panic on recent Linux kernels.



Disabling SACK functionality can lead to increased delays; however, it will protect servers from possible denial of service attacks - a temporary decrease in TCP/IP performance, according to Qrator Labs, is a reasonable way to neutralize a severe vulnerability. Patches for these vulnerabilities are available for both operating systems.

BGP OPTIMIZERS



BGP incidents are long-lasting; on a current scale of magnitude, hijacks or route leaks that spread far enough have the most poison in their longevity and distribution. That is probably because network development is progressing at

a considerably slower rate than other areas of development. That has been true for quite a long time, and it is a legacy that has to be abandoned. Money must be invested in network software and hardware, and into peoof course, it's

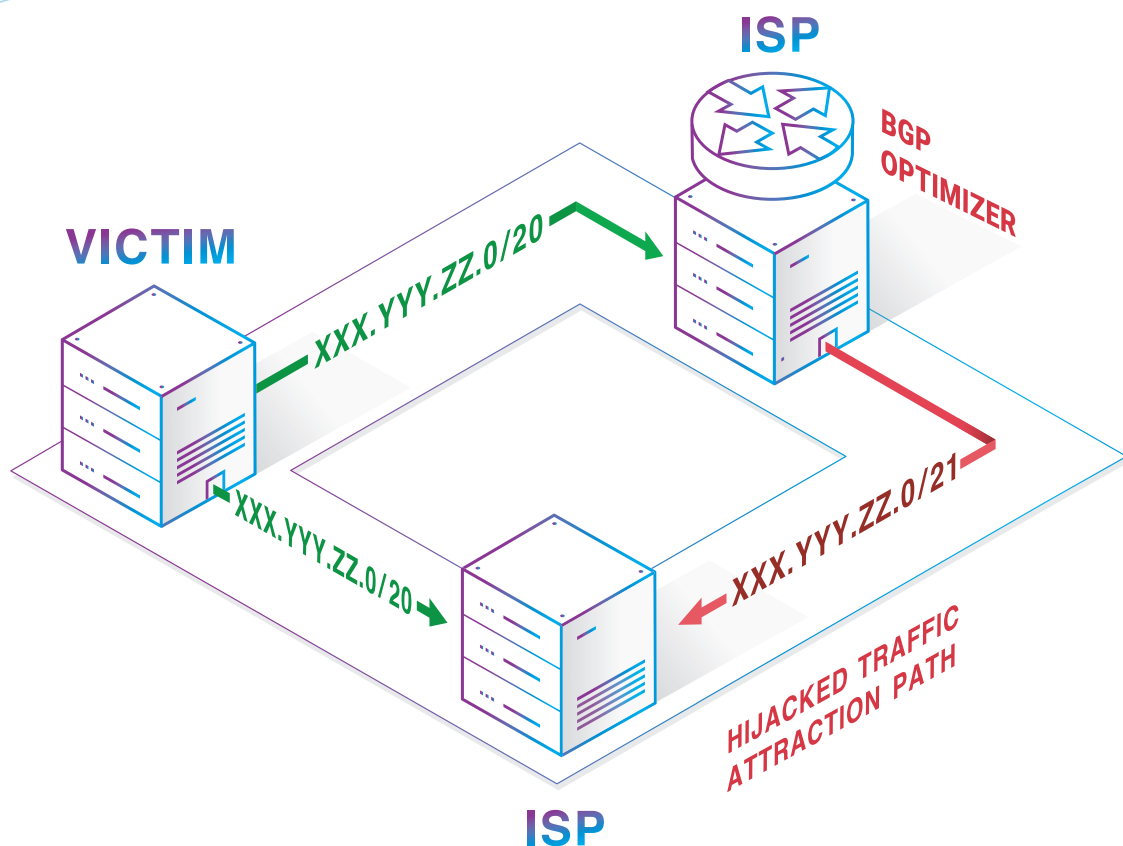
BGP incidents are long-lasting; on a current scale of magnitude, hijacks or route leaks that spread far enough have the most poison in their longevity and distribution. That is probably because network development is progressing at a considerably slower rate than other areas of development. That has been true for quite a long time, and it is a legacy that has to be abandoned. Money must be invested in network software and hardware, and into people fine-tuning BGP filters.

BGP optimizer-related incidents showed that the BGP statistics everybody relies on upon contains many issues. The point is that everything can be changed in the BGP announcement received from a peer before being announced further – the protocol is very flexible. That is what optimizers exploit: more or less specific prefixes together with malfunctioning filters and local pref. AS-path is

the second thing that gets exploited. Prepends (being a part of the AS-path) are the third. If someone de-aggregated a prefix into two more specifics – he would usually win the path challenge and get the traffic. Optimizers take the valid route and announce more specific prefixes – it is pretty straightforward. And it works, wrecking half of the Internet along the way.

BGP optimizers exist because a lot of companies want to control outbound traffic flow automatically without considering incoming traffic at all. Accepting routes that should not exist is a huge mistake because such routes don't exist at their origin.

Many articles were written during 2019, including ours, about the risks of "BGP optimizing". In the case of Verizon,



not a pleasant task to create a new filter policy for each new customer. However, one particular customer was not unique. We know that Verizon doesn't have any prefix filter because hundreds of problem-causing routes came from their customer AS396531, a stub – AS with only one connection to other AS. Moreover, it hadn't set a prefix limit for that link. There wasn't even a basic check for the presence of other Tier-1 operators in the AS_PATH from the customer's direction (this type of inspection can be set by default and doesn't need to be changed).

There were several discussions about this incident in the press. Besides possible actions for Verizon, they also described the benefits of RPKI and continue to support strict maxLength in ROA records. But what about maxLength? We know that during strict max length, all of the more specific become Invalid according to ROA validation. We also know that there exists a "Drop Invalid" policy. There is a draft that said that maxLength should be equal to prefix length.

Cloudflare follows this best practice. However, there was a small problem. Verizon did not have the "Drop Invalid" policy. Maybe it did not have RPKI based validation at all. As a result all the more specifics were re-announced even further, even though they were considered Invalid from Prefix Origin Validation and attracted all the traffic themselves. The main problem was that while they were still Invalid – Cloudflare couldn't simply announce those more specifics by themselves because its upstreams would drop those routes.

Route leak could have been eliminated once it was discovered. Simple AS_PATH manipulation of the form: ASx AS396531 ASx (where ASx is the origin AS number) during route creation could help these routes to avoid leaks on their path by applying a loop detection mechanism, while the ISP awaits an answer from the third party. Although every party has to do this manually and keep such policies in mind.

Again, in reality, the most common method was used: writing a complaint. And this resulted in another hour of delay. Communication could be painful and we cannot blame the Cloudflare – they did their best under the given circumstances.

What do we end up with? Sometimes we are asked how difficult it is to use BGP in organizing network attacks. Here is a possible example. You are a beginner attacker. You connect to Tier-1 or another huge ISP that does not have any fil-



13385546

BGP OPEN PORTS, 2019

ters. Then, you choose any service you like as a target and take prefixes of this target service and start to announce more specific of them. Also, you drop all data packets that are redirected to you. You thereby create a black hole for this service for all transit traffic on this ISP. The target service would lose real

traffic – the more revenue. So, if you announce subnet prefixes of these routes with the same AS_PATH, you'll get the rest of the traffic. Also, the rest of the money.

Could ROA help? Possibly yes, if you decide not to use maxLength at all and you don't have any ROA records with intersected prefixes in them. For some operators, it's not an available option.

Regarding other security mechanisms, ASPA wouldn't help (because AS_PATH is from a valid route) with this particular hijack type. BGPsec is also ineffective due to the deployment rate and remaining likelihood of downgraded attacks.

Thus, we have a clear profit motive for an attacker and a lack of protection. A great mix!

What should be done? The obvious and the most radical possible step – review your current routing policy. Besides, it would help to divide your address space into the smallest pieces (without intersections) that you wish to announce. Sign a ROA only for them without using a maxLength option. Current ROV can and would save you from this attack. But again, for some operators, this is not a reasonable approach due to the exclusive use of more specific routes.

With Radar.Qrator we can try to monitor that kind of situation. To do so, we need basic facts about your prefixes. You can contact us, set a BGP session with our collector and provide information about your view of the Internet. We'll highly appreciate it if you send us a full-view (so we can find propagation scope of other

2019



61668527

Almost x2.5 compared to 2018

BGP STATIC LOOPS, 2019

money because of this DoS that would affect a considerable number of customers. It would take at least an hour to identify the cause and another hour to resolve the problem, and that is only if the incident was unintentional and all the parties are willing to solve it.

In March 2019, there was another case, that at the time we did not connect to any type of BGP optimizer. Though it still deserves a bit of attention.

Imagine, that you are a transit provider, announcing routes to customers. If they are multihomed, you will get only a part of their traffic – however, the more

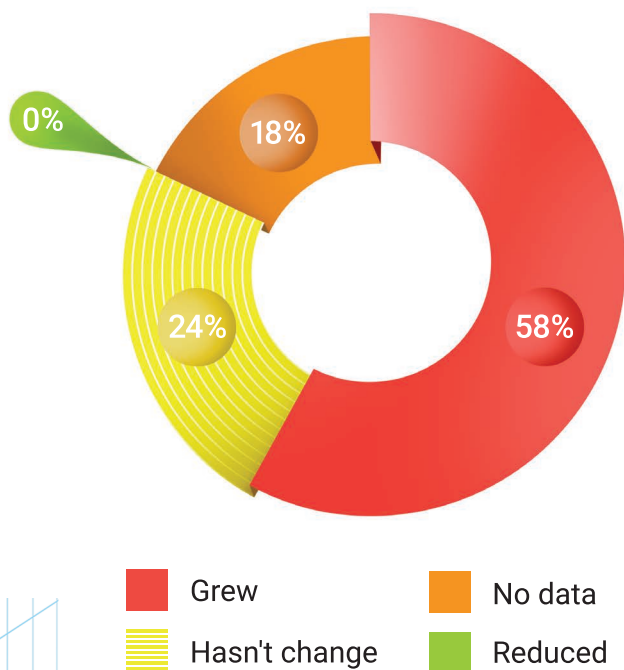
incidents), but the whole list of routes with your prefixes should be enough for the beginning. If you have already set a session with us, please check that all your routes are sent. We should have reliable data about your prefixes so that in the future we can automatically detect that type of attack on your address space.

When you discover such a situation, you can also try to mitigate it. The first

approach would be to announce routes with these more specific prefixes yourself. In case of a new attack on these prefixes – repeat these steps. The second approach – punish an attacker by restricting access of your routes to their AS. You can do this by adding their ASN in AS_PATH of your old routes and thus avoiding their AS by exploiting BGP route loop detection mechanism for your benefit, which we described earlier.

BANKS

DDoS threat level change in finances, 2018



serious consequences of security breaches.

Most financial institutions surveyed consider hybrid solutions to be the most effective means of countering distributed denial of service attacks.

The dynamics of the last two years clearly indicate that the information security field is growing at an enormous pace; over the past 2 years, most banks have increased investment in information security. Cybersecurity has already become visible at the level of company management. Business leaders are beginning to pay more attention to the implementation of security policies, and the position of Director of Information Security has acquired a completely new meaning. Information Security managers are gradually transforming into key advisers for top managers of financial organizations, implementing business tactics and security strategies in accordance with the company's needs.

We've conducted a study in Russia during 2019, showing that financial institutions recorded an increase in information security and seek to accelerate such investments.

Respondent banks highlight financial and reputational damage as the most

MOST ENDANGERED INDUSTRIES, 2019



IV

BANKING



III

BETTING



II

E-COMMERCE



I

**PAYMENT
SYSTEMS**

E-COMMERCE

DDoS attacks remain a significant threat to Russian retail, especially developing digital service channels. The number of attacks in this segment continues to grow.

In some market segments, despite its overall stabilization and consolidation, confidence that competitors will refrain from lousy behavior remains at a low level. At the same time, large online stores for the most part trust their clients and do not consider personal motives of clients as a significant reason for cyberattacks.

As a rule, medium and large e-commerce businesses learn about their readiness for DDoS attacks only from passing an actual "battle test". The need for preliminary risk assessment and


project preparation is far from being universally recognized, and even fewer companies actually carry out internal reviews or audits.

The respondents consider the main reasons for hacks to be a malfunction at the store level and theft of the user base.

In general, the level of maturity in the retail sector towards cybersecurity is growing. All respondents use some form of DDoS protection and WAF.

In future studies, it is planned to include a representative sample of small online businesses among the respondents and to study this market segment in greater detail, including its perceived risks and current security level.

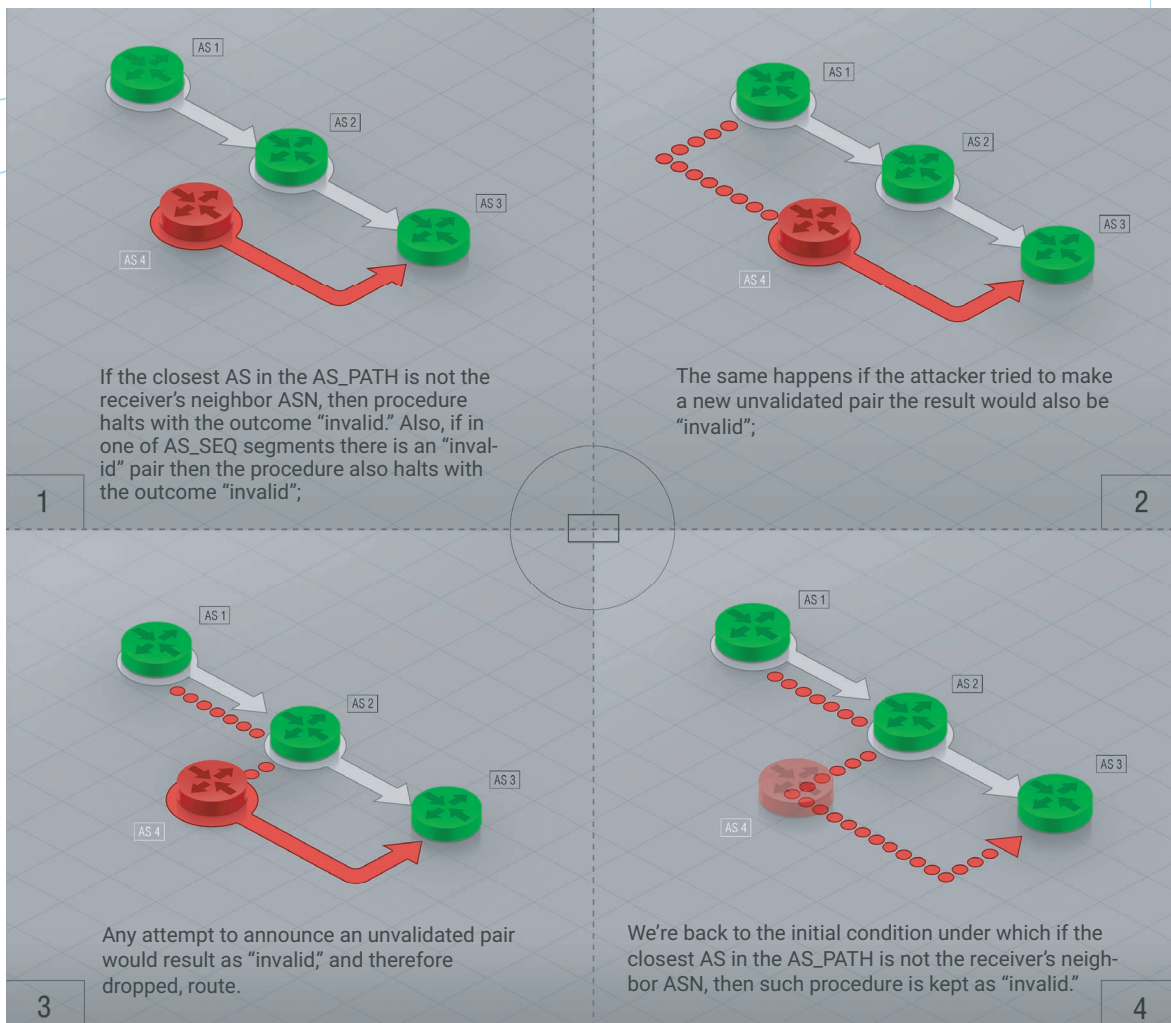
Most frequent incidents




RE-OPTIMIZING THE FUTURE OF BGP ROUTING

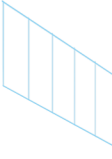
We hope that ASPA will become an RFC in the upcoming year – almost everybody understands that there is a need for the wider that the IRR/RPKI combi-

nation, yet more lightweight than BGPsec solution. 650+ sessions established with Radar. Qrator by the end of 2019



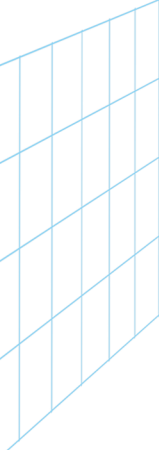


Radar.Qrator is upgrading the usability and reliability of the service and the underlying BGP relations model we serve with our subscription-based real-time monitoring.




Last year Radar.Qrator put a great deal of effort into accelerating its data processing and SLA under which we provide a real-time feed to our customers. Currently, Radar is the biggest BGP-collector and analyzer in the world, with over 600 sessions established, and we intend to deliver the accumulated data to real-time feed subscribers without delay.

Radar.Qrator is growing faster than expected, in terms of both free users and subscribers. In 2020 thanks to Radar subscribers we've begun developing several improvements that will be deployed incrementally during 2020 – one of them being new storage for incident data on each AS.




Radar.Qrator is growing faster than expected, in terms of both free users and subscribers. In 2020 thanks to Radar subscribers we've begun developing several improvements that will be deployed incrementally during 2020 – one of them being new storage for incident data on each AS.

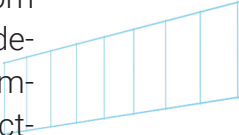
Another issue is the expected latency in the Radar web interface available for everyone free of charge. As the number of feeds grew, we were confronted with the need to upgrade our database and processes for serving web user queries. Due to the increasing number of users and connected autonomous systems we must serve each customer quickly through our web page.



We hope that ASPA will become an RFC in the upcoming year. The need for a solution that is wider than the IRR/RPKI combination, yet more lightweight than BGPsec is well understood. This year we saw that BGP misconfigurations that lead to route leaks could produce massive outages, even to the most experienced and professional operators. Surprisingly, those incidents also proved that there is no standard practice capable of defeating all possible failure scenarios.



We need the biggest ISPs in the world to support this measure to get the necessary initial traction. Participation from large communities that could help detect route leaks is another step. Simplifying, ASPA is a new object connecting the existing ROA and RPKI objects. ASPA realizes the principle “know your provider”: ISP's only need to know their upstreams to employ a secure method of defending themselves against most BGP-related incidents.



As with ROA, ASPA allows filtering routes at any connection point: with upstreams, peers and of course customers. Combining ROA and ASPA could solve most of the BGP security for almost anybody without making significant changes to the protocol itself, filtering out intentional attacks and random (human factor) errors. However, we would have to wait for the software and hardware support for the ASPA and, still, optimizing cases are out of ASPA's scope.

An essential benefit of ASPA is that it's a rather simple concept. We plan to make this draft a working protocol and an RFC with the help of co-authors and the IETF community.



<<<< 2020

PROCESSING LOGIC

HTTP/2

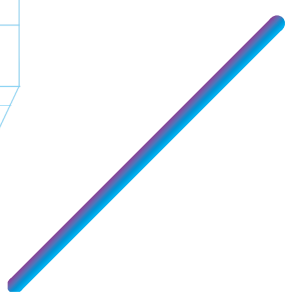
CONTAINERS AND ORCHESTRATION

YANDEX.CLOUD AND INGRESS 2.0

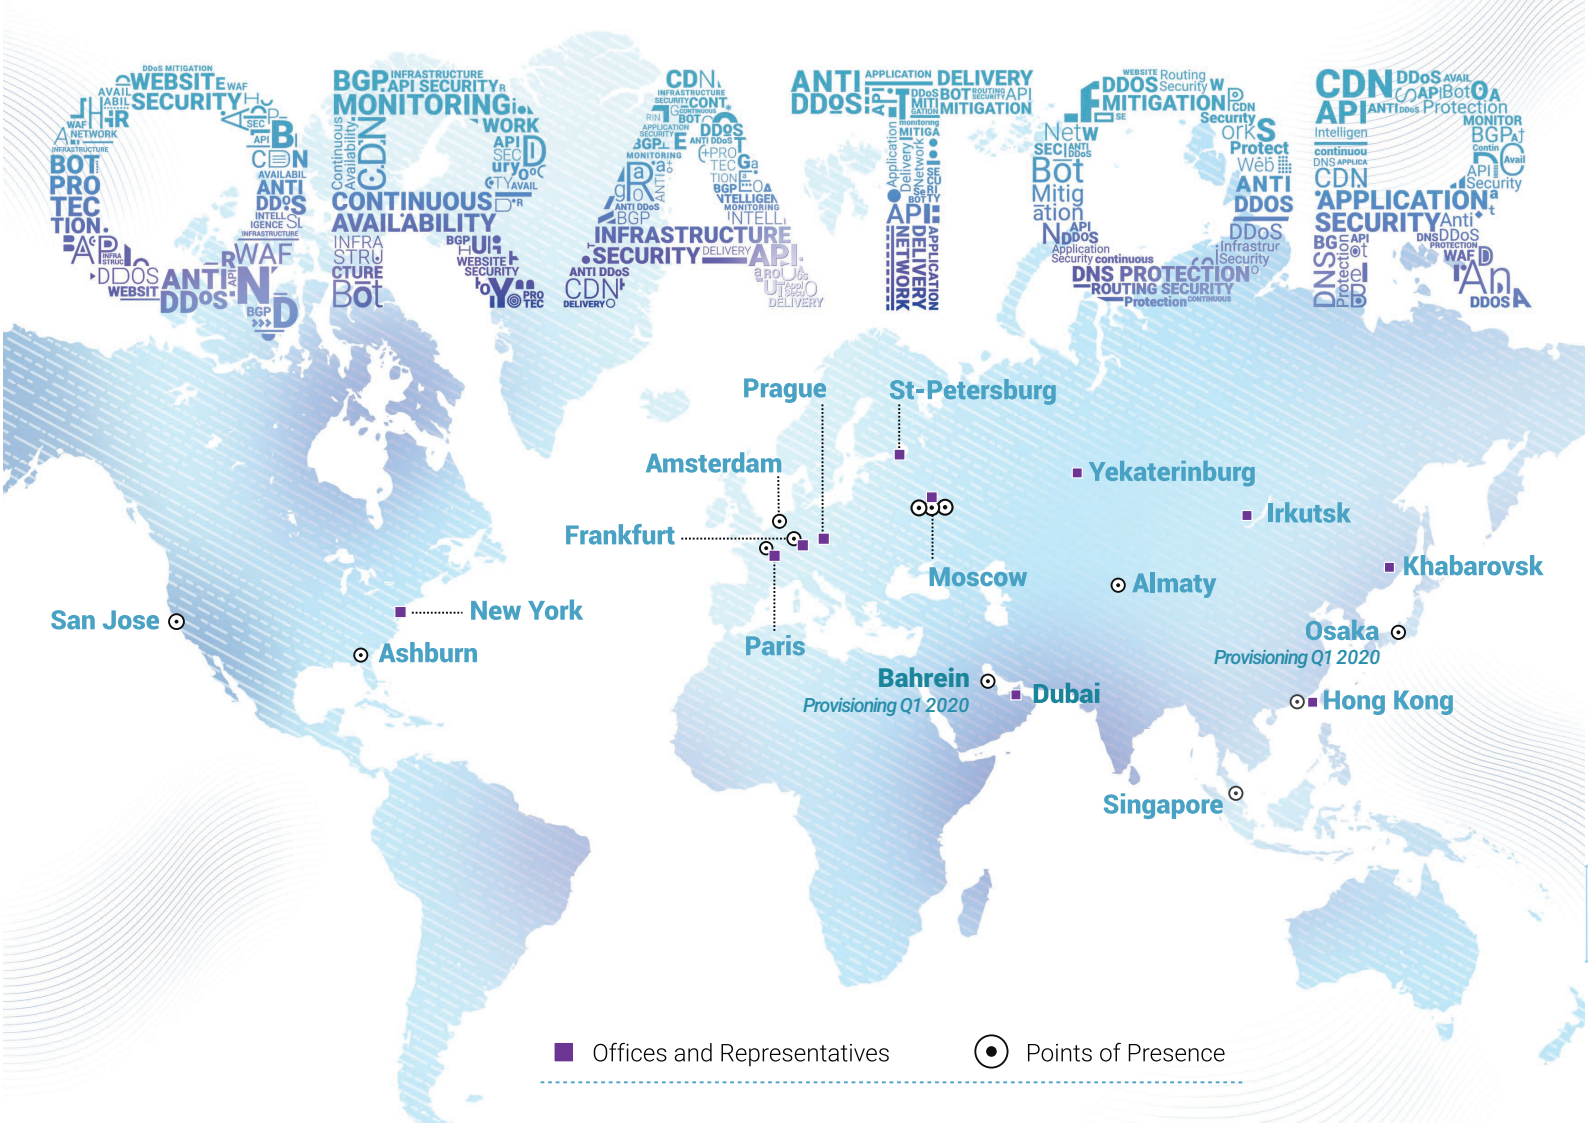
TLS 1.3 AND ECDSA
CERTIFICATES

SMALLER UPGRADES

2020 >>>



IPV6
ANTIBOT
HARDWARE
FUTURE UPDATES
TO THE NETWORK



CURRENT AND FUTURE DEVELOPMENT OF QRATOR FILTERING NETWORK

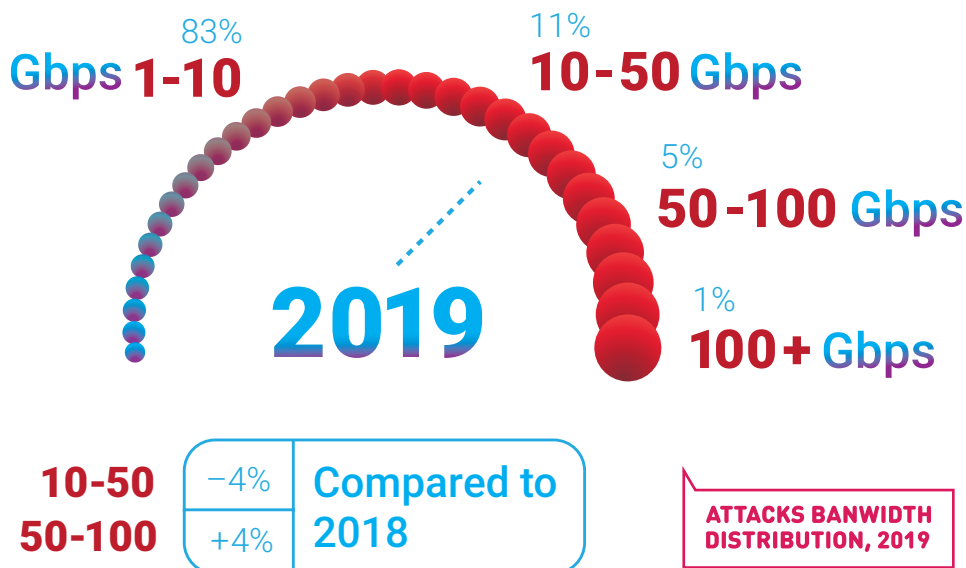
PROCESSING LOGIC

During 2018 and 2019, we've invested considerable time and effort to improving our filtering rules, which are the foundation on the DDoS mitigation service we provide to our customers. They are fully deployed and already demonstrably improving effectiveness, as well as memory and CPU consumption

at our points-of-presence. The new filters allow us to synchronously analyze requests and serve different kinds of JavaScript-challenges for bots, which helps further advance the quality of attack detection.

The primary objective for the change in the logic behind our filtering is to detect multiple classes of bots from the first request. Qrator Labs uses a series of stateful and stateless checks within the filtering logic and only after we confirm the legitimacy of the user do we send a request on to the customer's server. So, we need the filters to work blazingly fast because nowadays DDoS attacks serve tens of millions of requests per second, and some part of them could be one-time unique requests.

In 2020 as a part of this project, we are going to improve the traffic onboarding process. Our filtering network will react faster to new customers with the traffic flow we have not yet observed and modelled. As a result, customer enrollment will become smoother with less false positives at the start of attack mitigation, and we will be able to respond more quickly when urgently approached by customers in the midst of 'fights for their lives' under massive assaults and in need for quick mitigation.

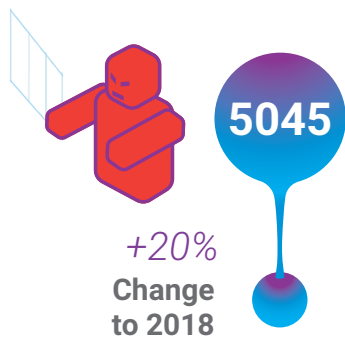


HTTP/2

The issues with HTTP/2 implementations that concerned us (DoS protocol exploitation possibilities) were mostly resolved during 2019, which enabled us to support this protocol within our filtering network.


We are actively working to provide HTTP/2 support for every client in 2020 after substantial amounts of live testing. As a part of this development in 2019, while enabling TLS 1.3 early in the year, we provided eSNI with an automated Let's Encrypt certificate.

AVERAGE ATTACKING BOTNET



ATTACKING DEVICES /
BOTNET SIZE, 2019

LARGEST ATTACKING BOTNET



CONTAINERS AND ORCHESTRATION

Modern containerization trends do not align well with our security approach. That means that we have to deal with many challenges when working with Kubernetes. Because such an orchestration tool wants to manage almost everything, which we can't permit, since we operate a distributed network where nodes communicate with each other, it would be highly risky to give any containerization stack too much power. It's a pity that Kubernetes does not necessarily have all those features in place, out of the box, or maybe it has them on a black-box level that couldn't be thoroughly inspected. However, this doesn't stop us from integrating Kubernetes with the infrastructure of the Qrator filtering network.

The future of developing and improving fault-tolerant infrastructures likely

involves integrated images, containers and ultimately, services meshed with each other to form service networks. As data volumes grow so too must the monitoring effort expand, and the easiest way to build monitoring capacity is to provide a natural and secure way for applications to communicate with each other. We believe that Kubernetes has proven to be one of the best solutions available, fitting the needs perfectly, though sometimes with additional work required in order to handle such a challenging environment as DDoS attack mitigation. Qrator Labs customers can already feel the power of our new log and monitoring system via the client dashboard.

**ATTACKS DURATION, 2019
(COMPARED TO 2018)**

1.9 HOURS OF AVERAGE ATTACK

2018

2019

1.9

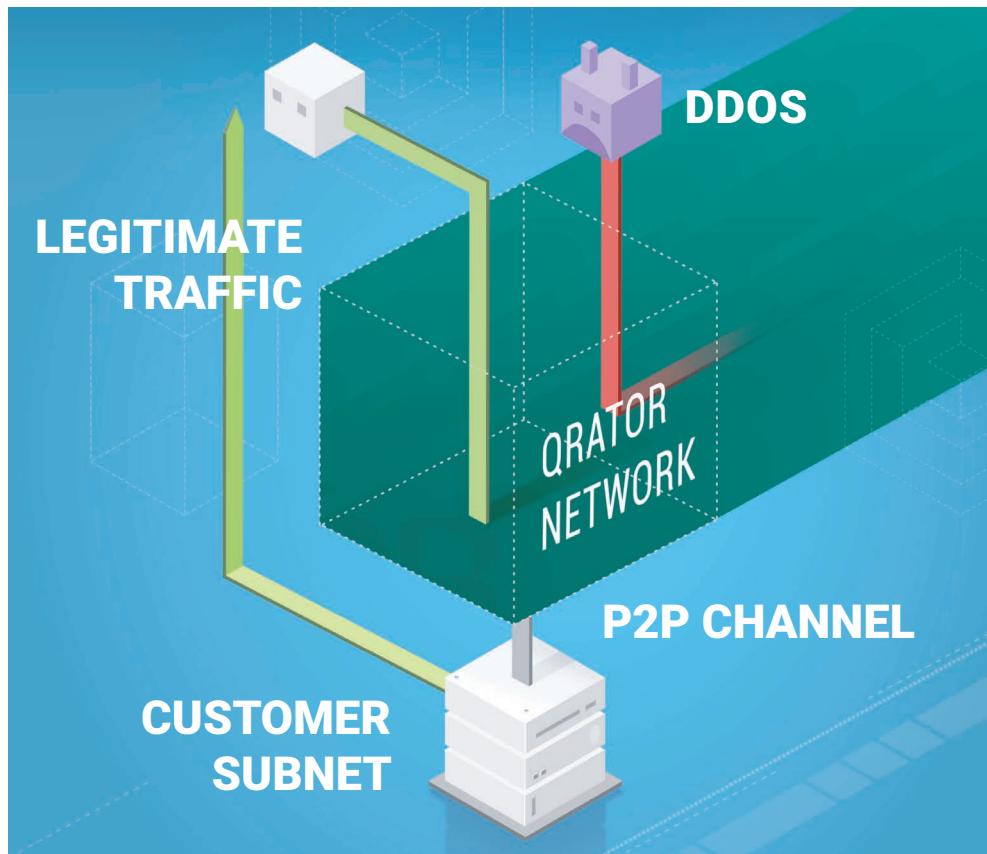
27% SHORTENING

YANDEX.CLOUD AND INGRESS 2.0

During 2019 together with Yandex.Cloud we deployed an improved version of the Ingress filtering service, which allows more precise fine-tuning and broader configuration. It is available to the end-user via informative yet straightforward interfaces. This newer

Ingress service is already deployed on our leading filtering network, ready to serve customers seeking off-ramp traffic inspection.

Yandex.Cloud made it possible for us to go through a process we never attempt-



ed before – integrating two big clouds – and it was in an unconventional manner. We were able to integrate two infrastructures with the help of Qrator Labs physical nodes that were placed in the network of a partner and worked over his traffic.

A newer, more deeply customizable version of the Ingress service is ready for use after extensive testing. With one of the largest cloud service providers in

Russia – Yandex – we’ve built a state-of-the-art inbound traffic filtering service specifically designed for businesses with excessive amounts of outbound traffic.

The new Ingress service is easy to understand and implement by the end-user. It allows for customization of all the parameters that were identified during the extensive testing period.

TLS 1.3 AND ECDSA CERTIFICATES

At the beginning of 2019, we enabled TLS 1.3, and simultaneously deprecated SSL v.3 support. Since we provide DDoS attack mitigation without decryption, we made additional improvements

to the scheme, improving the speed and reliability of the syslog transfer. We want to remind you of the benchmarking results.

Signature type	Handshakes per second
ECDSA (prime256v1/nistp256)	3358.600
RSA 2048	972.567

As you can see, for a single core of the Intel® Xeon® Silver 4114 CPU @ 2.20GHz (launched Q3’17), the overall difference in ECDSA performance, compared to the widely adopted RSA 2048 is 3.5x. Now let’s take a look at the same processor’s OpenSSL -speed results with ECDSA and RSA.

Signature type	sign	verify	sign/sec	verify/sec
RSA 2048 bit	717 µs	20.2 µs	1393.9	49458.2
256 bit ECDSA (nistp256)	25.7 µs	81.8 µs	38971.6	12227.1

With modern CPUs, you get almost the 5x difference in favour of ECDSA.

SMALLER UPGRADES

One of the smaller and stealthier features introduced in 2019 was the active upstream test. If for some reason a failure occurs in one of a customer's multiple connections under our mitigation service, we would be the first to know and react, by withdrawing the malfunctioning link out of order until the issues are resolved. We quickly recognize the problem by not only looking at the number and percentage of errors in the traffic flow but also monitoring the specialized health check interface we have implemented with our customer.


Near the end of 2019, Qrator Labs made a substantial update to the customer dashboard. While we still maintain the previous version for our customers who

are used to it, we strongly encourage new clients to try and use the updated version, which has much broader features in place, like the TLS certificate management.

While the older version applies server-side rendering, the updated version incorporates REST API and the client-side JS application. We believe that such an approach will enable us to deploy new features more rapidly. One of the first such features will be the "Analytics" section of the dashboard, which allows custom graphics that correspond to the number of data sources and types they work with. Other advanced metrics will be available.

L2-4 to L7 ISO OSI attacks comparison

71,77%
L2-4
VOLUMETRIC
& PACKET INTENSIVE
ATTACKS



28,23%
L7
APPLICATION
LAYER ATTACKS

IPV6

We are actively preparing to enable IPv6 on Qrator Labs filtering network in 2020. As a cybersecurity company, such a transition is complicated. Due to the nature of DDoS attacks, mitigation with IPv6 address space demands an entirely new approach because it is no longer possible to use any form of "list" when working with a theoretical limit of 2^{128} IP-addresses. And we're talking about TCP here, not the UDP. For us, 2020 is going to be the year of

IPv6. With the depletion of the IPv4 address pool, there is no other way to upgrade any network to meet future demand effectively. We believe that the specific challenges we face with DDoS attack mitigation could be solved effectively under the IPv6 terms. For us, this primarily means that we should be able to announce a customer's IPv6 prefix via BGP across our network and simultaneously provide him first-class cybersecurity services.

IPv4

2^{32}



4 OCTETS

IPv6

2^{128}



16 OCTETS

ANTIBOT

In 2019 we saw a steady movement of fingerprinting and anti-fingerprinting. Understandably, many Internet services would like to find a way to separate automated or bot requests from the legitimate ones and, on the other hand, there is nothing so unique in the browser's intention to provide an additional level of security to legitimate users. At Qrator Labs, we have been reminding the market for years that if the information is public and no authorization is required to get it, there is almost no way of pro-

tecting it from automated parsing. At the same time, we remind business owners that they have discretion over how to handle information on requests to their servers and related to specific user behavior. Taking a proactive approach in this regard could help evaluate and identify malicious activity under particular circumstances.

As bots produce an increasingly greater share of total traffic, we anticipate the point in time when large businesses

will create specific gates for automated interfaces and requests from good bots will be appropriately identified and served. Even now, at the beginning of 2020, we would say that a user who is not 100% legitimate by all indicators would have a hard time accessing each and every web resource out there — some would not allow you in if, for ex-

ample, you simply changed your browser's user-agent to custom. Antibot services are under active development within our company, and in 2020 we expect to be ready to activate pilots with customers who are interested in semi- to fully automatic detection and protection against a wide range of automated, or bot-originated, actions.


HARDWARE

We've tested the newest AMD processors, and we liked it so much that during Q1 of 2020 Qrator Labs will turn on a new point-of-presence in Osaka that would be built on the AMD cores. PCI Express Generation 4 available with AMD processors promises to double the bandwidth between the CPU and NIC. It will be tested and evaluated during 2020, and we will publish the results of DDoS attack mitigations in specific use-cases and stress scenarios. This channel between the NIC and CPU has become a bottleneck with the extensive growth of ethernet transfer bandwidth. The combination of additional cores, bigger cache and newer PCI Express promises positive results.


Our reason for trying the AMD CPU lies within the fact that "diversification" has always been a rule for the network architecture, though in 2020 for the first time we are building a point of presence on an entirely new CPU.

We eagerly anticipate testing the combination of PCI Express generation 4 capable CPUs and NICs at the Osaka point of presence during 2020. Hopefully, the combination will suit our needs. Asian

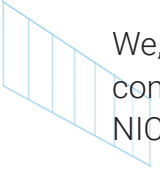





traffic is growing exponentially, and we want to explore possibilities of building DDoS attack mitigation with various hardware on board.



At the same time, we are waiting for the Ice Lake generation of the Intel processors. We will plan further development based on the test results with CPUs from those two providers.




We, along with the rest of the industry, continue to await the Intel 800 Series NIC. We hope to have it by the end of




2020 to see if there will be a place for it within our infrastructure.

Regarding switches, we still prefer Mellanox equipment over anything else, and the company has proven to be a highly reliable partner multiple times.



It is hard to say, but we should share our view that Broadcom dominates the market for networking hardware and NICs. We love Mellanox products and hope they only benefit from the NVIDIA merger, but the results are not precise yet.




FUTURE UPGRADES TO THE NETWORK



In 2020 we are ready to expand our partnerships with various services providers in areas such as Web Application Firewall and Content Delivery Networking and Optimization. Qrator Labs has always tried to integrate with multiple service providers to offer the best possible combination of security services. By the end of the year, we plan to announce a variety of new services.



In 2019 we were challenged by the WebSockets security issues. This particular technology is becoming more and more popular, and its complexity makes it challenging to secure correctly. During 2020 we will actively work



with our customers to determine our future development path. The ability to secure customer applications and use a WebSockets technology, even an arbitrary payload within one, will enrich our future development efforts.

What we do is not know-how and not inception. And we were late. We do our best to keep up by doing unique things. A little part of this lies within the fact that a group that studies academic and scientific writings that correspond to their primary activity has some extra time to prepare for upcoming challenges.



www.qrator.net **Qrator Labs CZ**

Legal address:

Nové Město, Růžová 1416/17
11000 Praha
CZECH REPUBLIC

Actual address:

Nové Město, Růžová 1416/17
11000 Praha
CZECH REPUBLIC

+420 602 558 144
mail@qrator.cz