How can Enterprises Safeguard their Data and Digital Resources?

Dr. Thomas King from DE-CIX explains how an IX platform can shield networks from a range of malicious and accidental security incidents.

As enterprises exchange more and more data along their digital value chains, they need to take a closer look at how to protect their connections to trusted partners and to their digital resources in the cloud. Dr. Thomas King, CTO at DE-CIX, the world's leading Internet Exchange (IX) operator, explains how an IX platform can shield networks from a range of malicious and accidental security incidents.

As operators of the Internet infrastructure, we are tasked with making sure that we provide a service which our customers are happy to trust. Although operating Internet Exchanges is clearly a B2B business activity, in the end what is important is the trust of the end user sitting in front of his or her laptop or cell phone. For Internet infrastructure operators, it is vital to our own business and to the business of our customers to strengthen and maintain trust in the Internet. For an Internet Exchange (IX), continual research and development, security audits and certifications, and developing and maintaining best practices are as central to the process of ensuring network security as the provision of additional security services to shield against willful and accidental damage.

Increasingly, we see that interconnection customers today look not only at the kind of on-top security features an IX provides, but also how the interconnection platform is operated. There is a growing interest, especially from customers in the enterprise sector, in topics like ISO 27001 or the German Office for Information Security's BSI IT-Grundschutz certification. They want to see evidence that operations follow certain criteria and best practices, to know that an IX is operating in a secure and reliable way. Enterprises are used to vetting their business partners on security and policy-related topics. Large enterprises rely heavily on their infrastructure partners to provide services and ensure low risk levels for their own operations. Some of the largest customers connect not only at one location to the IX offering of DE-CIX, but connect to most if not all of the 29 locations DE-CIX is operating globally.

Regardless of the kind of network, all have a basic need for routing security, like being effectively shielded from IP hijacks through the use of RPKI (Resource Public Key Infrastructure). Protection against DDoS attacks is a different story, because not all networks are necessarily interesting targets for attackers. Customers who run or host game servers for their end users, for example, usually feel the brunt of a lot of DDoS. This is a group of customers who use blackholing heavily, and can benefit from the new Blackholing Advanced service.

At the same time, enterprises tend to have a greater interest in security services, because their operations and products often exist in real-world spaces. Take car manufacturers: digital automotive services are becoming more and more important, and they need to make sure that the cars don't become inoperable or defective due to either an attack or a misconfiguration.

Shielding networks from vulnerabilities over an IX

1. **DDoS attacks**

Probably the best-known attack type that can be mitigated at an IX is the volumetric DDoS attack. The goal of a DDoS attack is to stop a certain destination from communicating with the Internet. For instance, you

have a web shop hosted on a web server, and your competitor hates your shop because you're more successful. A DDoS attack on your webserver will mean that your web shop is no longer accessible to your customer – and all the customers go elsewhere to shop (e.g. to your competitor).

When it comes to DDoS, amplification attacks have been very strong in the last couple of years, and a new emerging threat is ransom DDoS. However, despite increasing growth in the number of attacks and the volume of attacks, my impression is DDoS attacks are currently not developing as aggressively as they have done previously. There has been a lot of work done to fix new vulnerabilities as they emerge. Added to this, the free and open availability of DDoS mitigation services from companies like Cloudflare or FastNetMon has also helped to solve the issue to some extent. Don't get me wrong, though: we are still seeing these attacks on a daily basis.

So, how do we mitigate DDoS? We used to use standard blackholing. If you have blackholing, you can protect an IP address so that you stop traffic being sent to it while it is under attack. The good thing is that there is no collateral damage for the networks in the firing line. But the disadvantage is that the destination still is unable to communicate, meaning that, through the mitigation measure, the attacker has ultimately achieved the original objective.

With DE-CIX's new patented Blackholing Advanced service, we can go a step further and not only limit the data being sent to an IP address, but we can limit it to certain TCP and UDP protocols . Because if you talk about amplification attacks, we can look at which TCP/UDP source and destination ports specifically need to be blocked. We just block this particular port, and all the other ports are still accessible, meaning that the network can still communicate.

The second innovation of Blackholing Advanced is that it's no longer simply a binary switch between "data is flowing" or ""no data is flowing. We can also limit how much traffic is going to the destination aka. rate limiting – rather than hundreds of gigabytes of traffic, we can reduce it to just 10, 15 or 20 Mbits, so that the destination is not completely overwhelmed. The destination can still handle the load that is coming in, they can sort out the garbage, and allow legitimate requests to get in and be answered. Communication and service is thus still possible.

2. IP Hijacking

Another risk to networks in the Internet is routing insecurity through IP hijacking. To give you an example of how this works: Let's say you, as a malicious actor, want to wiretap the traffic that goes to an IP destination somewhere on the Internet – perhaps, to come back to our previous example, a particular web shop – because you want to steal the credit card details of the shop's customers. You can start announcing the IP space of the web shop, and if you do it right, you can receive all the requests which go to the web shop. You can either drop the traffic so that the orders from the customers don't get answered, or you can just pass it on to the web shop, having gleaned the information you wanted. This kind of IP hijacking can occur either by accident or on purpose. There have been incidents in the past where people have presumably done it on purpose – rerouting traffic from a bank, for instance, or also from the Bitcoin blockchain. But other incidents have certainly been accidental. YouTube was taken offline by Pakistan Telecom in 2008, because someone misconfigured something. They completely overloaded the network, because Pakistan Telecom Network was not big enough to handle all the load of queries going to YouTube.

With the increased number of networks and amount of IP space connected to the Internet, the increasing dependency of society on digital infrastructure, and also the value of the data being shared, it stands to reason that we can expect IP hijacking – whether it is malicious or unintended – to be growing. There are

simply more players in the field. In fact, the Internet Society MANRS project found that from the year 2019 to 2020, there was in fact close to a 40% increase in IP hijacking incidents, which is certainly worrying.

Technologies like RPKI Origin Validation and IRR filtering, which we provide at the DE-CIX's route servers, can be used to mitigate the problem. The function of RPKI is origin validation because it makes sure that it is not so easy to accidently announce the wrong IP space through a typing mistake or similar. It makes it possible to check whether you are allowed to announce this IP space, and if not, we can filter out the announcement very easily. IRR (Internet Routing Registries) filtering, on the other hand, is used to prevent the propagation of incorrect routing information. This filtering is already deployed in the Internet infrastructure for years, whereas RPKI Origin Validation has only become available recently.

Added to this, there's the forthcoming BGPSec, a standardization activity that is ongoing at the IETF. If you were to have origin validation based on RPKI together with BGPSec, which also uses part of the cryptographic building blocks of RPKI, then you would have full safety against hijacks. However, BGPSec is still in standardization, and unfortunately it has one major drawback: It is very resource intensive on the Internet routers. From what people say, I think it's still at least a couple of years away from deployment, if ever – so it's certainly not a short-term fix.

3. ASN hijacking

Another security issue is ASN hijacking. Every network that wants to be part of the Internet needs an Autonomous System Number (ASN). By hijacking someone's ASN, you can pretend to be somebody else. This can be used maliciously, mainly for sending unwanted stuff like spam and carrying out DDoS attacks. We have seen ASN hijacking in particular with companies that have registered an ASN but, for whatever reason, are currently not announcing it to the Internet. And it's very difficult to really ascertain who is behind the number – the legitimate owner, or a malicious actor. It looks as if the legitimate owner is behaving badly, which can result in them being blocklisted or far worse reputation problems. Therefore, I really encourage companies to keep an eye on their AS number even if they are currently not using it.

What else networks should do to optimize communication

From a different angle, although it is not new, bidirectional forwarding detection (BFD) has for some reason not yet become well-established. This is a shame, because it is very interesting for optimizing network communication. Basically, if two networks or pieces of infrastructure have a link and want to make sure that data is flowing in both directions, there's the so-called BFD protocol. Without BFD, it takes a couple of minutes to detect that there is an issue with a link, and in the meantime data you thought you were exchanging is being dropped on the floor, so communication is not happening. With BFD, an issue can be detected in a range of seconds or even milliseconds, so the parties can stop sending data over the broken link, and take an alternative route. Networks would be well advised to use BFD so that they can easily detect any issue and automatically reroute traffic. Early next year, we plan to implement a feature to support BFD for our route servers.