Sense of Security

WHITEPAPER

Achieving cyber resilience by reducing your susceptibility to attack

The reason why a DDoS mitigation effectiveness test needs to be part of your vulnerability management program.

Written by
Murray Goldschmidt, Chief Operating Officer

What is a Denial of Service Attack?

A denial-of-service attack has the objective of preventing legitimate users from accessing specific computer systems and services.

Denial-of-service (DoS) attacks typically flood servers, systems or networks with traffic in order to overwhelm the victim's resources and make it difficult or impossible for legitimate users to access them.

These attacks can also be more targeted and may not require large volumes of traffic if a specific component is more susceptible to outage through a crafted attack.

A **distributed denial-of-service (DDoS)** attack occurs at scale and generally is operated through a network of compromised computers on the internet, all controlled and orchestrated by an attacker.

What does DDoS look like today and in what direction is it hitting?

Today, more than ever, organisations are susceptible to outages that are caused through attacks launched at denying their ability to operate their business and service their clients.

Mitigation technologies have achieved greater penetration in the market and the cost of mitigation has come down. This is common with large scale cloud offerings and Content Delivery Networks (CDN's) now being very accessible. However, organisations remain exposed.

Why? It seems counter intuitive and against the grain. Are we really getting such poor return on investment from our cyber defence spending? The answer, unfortunately, is yes.

And the reason for the pretty dismal state-of-affairs (at least in relation to perimeter and infrastructure security) is because today's modern organisation has a pathetic fixation on buying whatever is sold to them because it sounds like it might be good, or because their auditor is looking to see evidence of some control in place to meet their compliance with a particular standard or policy.

The main issue relates not to the technology that was acquired, but rather with the deployment, the configuration, the integration into an active operational network and the ongoing management of the environment which, more often than not, is a nest of internal teams and third-party suppliers.

The main problem is that people are buying products in isolation (and even managed services for that matter, such as a managed Cloud WAF) but expecting them magically to work in synchronisation and in harmony with everything else in the eco-system.

Of course they never bother configuring them properly – accepting the vendor supplied defaults – and then never test their system end-to-end, never bothering to review the path that an actual transaction takes from consumer to server and back.

Add on top of this the expectation related to the "people" component. We are still largely relying on people to identify when things are going wrong, to know how to respond to it, identify the root-cause, and then isolate the problem and fix it before it becomes a massive problem that rapidly snowballs into catastrophe. **Yes – that's the reason why so many companies get hacked and DoS'd.**

Why a DDoS testing program really needs to be an essential part of your vulnerability management program.

Firstly, lets look at what a **Vulnerability Management Program (VMP)** is intended to achieve, and generally what they consist of.

Simply put, a Vulnerability Management Program needs to allow you to identify vulnerabilities and fix them before the vulnerability is exploited.

There can be many types of vulnerabilities at every layer in your system. You need to identify as many as possible, understand their risk, make informed decisions about the implications, determine where you will focus your scarce resources, fix the nominated issues and then validate that it has actually been resolved (and that you didn't create another problem while fixing the first).

And why are we doing all of this? Because we have an objective to be Cyber Resilient. That means to be able to avoid attacks where possible, **but if one occurs, we need to be able to operate through it and come out on the other side still in business.**

Vulnerability Management is an umbrella term that is inclusive of configuration management, patch management, vulnerability identification and remediation. It intersects with incident management and response and threat management.

Vulnerability Identification also includes automated means such as scanning, and more manual detailed forms of assurance such as penetration testing across all layers (people, networks, API's, web apps, cloud platforms etc).

Scanning is becoming more and more valuable when augmented with threat intelligence feeds, and even more valuable when this becomes a dynamic-asset-management system telling us exactly what is on the network, what software it is running, what it is vulnerable to and how to fix it.

You want to be able to avoid attacks where possible, but if one occurs, You need to be able to operate through it and come out on the other side still in business.

Penetration Testing becomes more valuable when expensive testers aren't finding the low hanging fruit, but using their intellect to find things that relate to business process rather than poor configuration.

So the idea is to tend towards continuous monitoring, leveraging high frequency assessment by having scanning integrated in an automated sense into our ecosystem. That will deal with identifying the problems faster. But what about fixing them? Well this is where we need to rely more on technology, by working towards intelligent condition matched automated responses that are self-healing. That is, the response reverts the system back to a known-good state where we are no longer exposed.

OK – so we have a bunch of techno buzzwords. But if the objective is to be cyber resilient, then why isn't testing our susceptibility to Denial of Service attacks not part of the mix? Very interesting indeed. In fact, you will be challenged to find a penetration testing assignment where DoS testing hasn't been specifically excluded. Sounds insane. But it's true.

We have all these great ambitions to be Cyber Resilient, but try and scratch beneath the surface, you will find resistance to actually determining what's wrong.

Is it because we like living on the edge of our seats, enjoying the exhilaration of not knowing when the next attack is coming and if we will survive? Or is it that people just don't want to ask questions when they don't really want to know the answer. If things were more obvious and clear it would remove the defence of plausible deniability. So we just exist, day to day, without really having any confidence about our resilience into the future.

What about director's duties, managing enterprise risk, maximising shareholder returns while conducting business in an ethical manner? Seems like vulnerability management is in the same boat as climate change. A mission statement crippled by an executive-endorsed-semi-functional-engine-room.

</rant>

OK – so we have established the DDoS Testing is seldom included in a Vulnerability Management program, because it will uncover what is essentially an inconvenient truth. But for those that are bold enough to accept our fallibility and genuinely want to improve their cyber security posture, the good news is that testing the effectiveness of your DDoS Mitigation techniques can be readily undertaken, it won't break the bank, can be made repeatable and will enable you to tend towards automated responses rather than relying on manual intervention.

The other good news is that all the tech that you have already invested in will do the job for you. In all likelihood you don't need to buy anything else, you just have to configure it, test it and know how to manage it.

Where are the attacks heading?

Let's cut to the chase. They are occuring at the application layer now, because most of the large scale volumetric attacks are arrested well before they get to your servers (whether cloud or on premise). Our CDN's and cloud scrubbing services do a good job for volumetric attacks. Because they are easy to identify and suppress through global networks that can soak up all the traffic. But application layer attacks are another beast altogether.

A single application layer attack can occur well below the threshold for any volumetric defence to trigger, yet have the ability to render a large scale enterprise web application unable to serve its customers.

Add in the fact that application layer attacks generally occur through a TLS (formerly called SSL) encrypted channel, visibility remains a major weak point. We have to understand where the TLS is terminated (and possibly re-applied) in order to inspect it. Given that TLS is a cryptography based technology, it is itself subjected to crypto attacks because crypto is very process intensive. So now we need to add crypto attacks into the mix as well.

It would surprise you to know how many people are relying on DDoS scrubbing centres to protect them at the application layer. These platforms simply can't. They aren't inspecting the traffic at the application layer!

And what about a Layer 4 Proxy. No good either. TLS is terminated above layer 4. The Proxy will just shift the traffic from one system to the next.

But what about my Web Application Firewall (WAF) you cry? Isn't that designed to protect me at the application layer? Well the good news is that tech such as this probably can help you, but not in the default configuration that is shipped by the vendor that you never bothered to change.

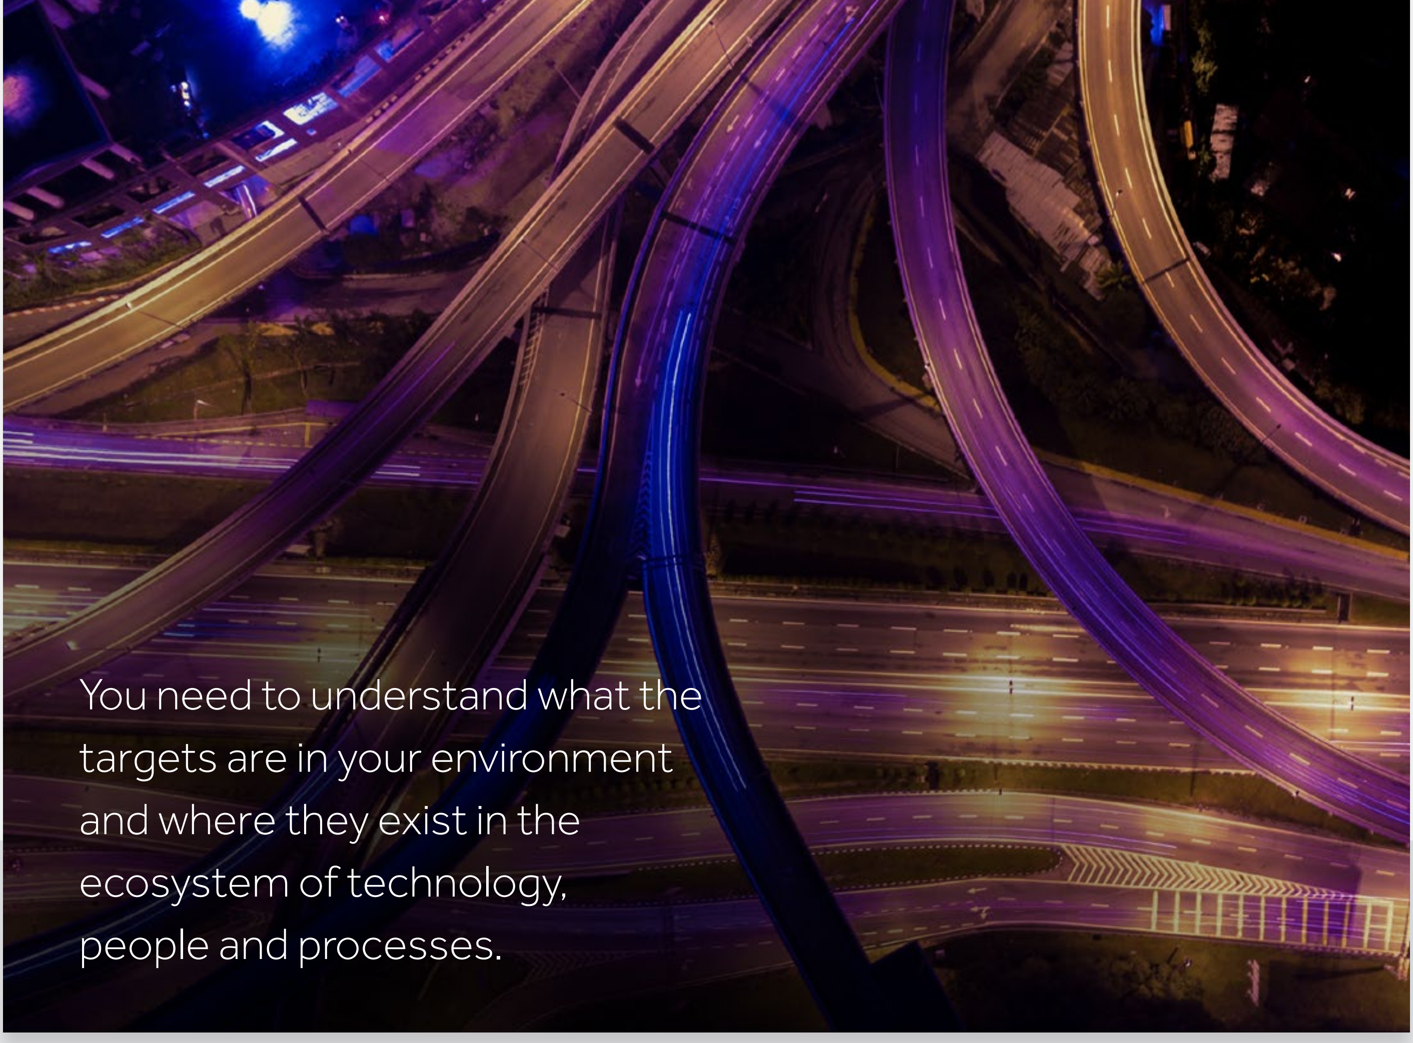

You need to understand what the targets are in your environment and where they exist in the ecosystem of technology, people and processes.

And what about my Next Gen Firewall – yes the one with about 50 different license and subscriptions they sold me? No – that won't help you – you were a fool to buy anything labelled NextGen or AI. It's the same as buying that all-in-one power tool on late night TV where they throw in the second one for free and a set of steak knives. The thing smokes when you drill through a piece of soft pine.

So what we are really saying?

My tech is ok, if I know how to use it? Basically – yes! We need to understand what the targets are in our environment (we will explore targets a bit more in the next section) and we need to appreciate that this target operates in an eco-system.

You don't just have an isolated web server for example! A web server relies on DNS to enable people to route to it. It relies on a router somewhere at the junction to the internet. It relies on a load balancer that probably terminated the TLS and chose which web server in the pool to send the connection to. And it also relies on a number of other downstream systems such as authentication servers to authenticate and authorise the people using the site, storage systems to hold the content and email systems to send out confirmations.

There is also a help-desk system somewhere and a contact centre, and a remote access VPN to allow your support staff to manage the platform over the weekend and late at night. Our cross-section of the enterprise needs to identify all, yes all, of these components.

So the target exists in an ecosystem of technology, people and processes.

In this paper lets take a deeper dive into the technology component.

Technology

Each type of technology will have some security capability. We need to understand where in the technology stack this operates so we know how it will contribute to our resilience objective. And then we need to check if it has actually been configured to deliver the security function that you want from it.

Tip:

our assessments have identified that attacks that are successful due to misconfiguration occur m ore than a staggering 70% of the time.

So what is a target?

A target can be an **IP Address**. It can be a service (e.g an email server) and it can also be a protocol. For example, a Web Server may serve content on HTTP and HTTPS. That's two possible targets. Why? Because the security controls are often applied at the service level. A web server could serve HTTP content but be rendered unable to serve HTTPS.

A target can also be a **particular page on a website** – such as the logon page to get into the application. Why? Because the logon page relies on an Authentication server to determine who gets through to the next page. And the logging server can only process a certain number of auth requests in a given time no matter if they are successful or not.

A target can also be a **firewall** or a **logging server** or both. Here's a nifty web server DoS attack. Send lots of connections to a firewall on a port that you know it is going to block. Say TCP port 444. It will likely serve the web site on TCP 443, but you expect it to block TCP 444. And when a firewall blocks something, especially under a default deny rule, it also logs. It logs it locally, and if configured to address more sophisticated business requirements, it probably also logs it to a centralized logging system (or SIEM). And you just joined the dots yourself.

The firewall's local storage and processing capacity is not unlimited, and the number of connections that the logging platform can process from the firewall and all its other log sources is also not unlimited in hardware or software or license or subscription.

So send enough traffic that you know is going to be blocked to a firewall and you will bring down the firewall and its logging server and thereby prevent the web server from serving the site. And you didn't even attack the web server. Cunning!

And before you say none of this affects me because I am using public cloud and they have oodles of capacity... sorry. **Cloud** is no different. Nothing is unlimited. Cloud is elastic – but not unlimited.

Anyway, auto-scaling is not a DDoS mitigation solution. It is a capacity response solution. And an expensive one at that. So while auto-scaling may make the attack more difficult, it takes time to kick-in (seconds to a number of minutes) and will deliver a denial of service to your wallet if left unchecked and operating at high capacity for an extended time.

But I outsourced the responsibility to my service provider

Nice try. But the investigation into the Australian 2016 Census Fail comprehensively and forever rebutted that claim. You can outsource a service or a function, but you can't outsource your responsibility and you have to have adequate assurances in place to validate that the service provider is delivering the service that you subscribed to!

But the guy who sold it to me said...

You walked right into that one didn't you. Did you read the fine print? DoS mitigation services have as many get-out-jail-clauses as a cyber insurance policy. Take mitigation leakage as an example. Speak to a sales person and they won't ever highlight to you that DoS mitigation may be vulnerable to attack leakage.

What does that even mean? Well, in order for the mitigation to get activated a threshold needs to be reached in order to determine what was previously considered normal traffic is now considered attack traffic that needs to be blocked. So the attacker's IP goes from a whitelist to a blacklist so that all subsequent traffic from it is blocked.

But there is a timer on this. An IP Address won't be in the blacklist for ever. All the attacker has to do is figure out the transition time and keep the traffic running (probably 15 minutes say). The mitigation then stops working, the very same attacker still keeps sending the same traffic it did before, but now the attack leaks through the mitigation, until the threshold is triggered again and the traffic goes from whitelist to blacklist.

Have enough patience and do this enough times from a wide range of sources, you essentially get a continuous attack leakage condition. That is, at all time, some traffic will be transitioning from blacklist to whitelist. And that's where the attack leaks through the mitigation. Read the fine print on the DDoS Service, and it will say this service may be subject to attack leakage. You just didn't read it, or didn't know what it meant. Now you do!

So the best thing to do is to test your service. Identify at which point the mitigation kicks in, and for how long. Know where and when the attack leakage will occur. Determine if you can tune it better to your requirements. The one size fits all policy can probably be made more bespoke to your specific requirements. And even if that mitigation has some limitations, remember you have other defences that can pickup the attack around the fringes. Defence in depth means you aren't relying on one control for your overall security.

Conclusion

Resilience testing has been part of most other industries for yonks. Cars are much safer today than ever before because the manufacturers have invested heavily in developing products and testing their effectiveness. Airbags, seatbelt tensioners, crumple zones, front and rear sensors etc. Air Travel is one of the safest forms of transport, because every incident requires analysis to the forensic level that enables manufacturers to improve their products and processes going forward.

IT and Cyber should be no different.

Add a DDoS Resilience Test to your Vulnerability Management program today! Remember you can't outsource your responsibility and the old plausible deniability trick is wearing pretty thin.

Sense of Security®

Cyber Resilience,
Information Security
& Risk Management.

COUNCIL OF REGISTERED ETHICAL SECURITY TESTERS
CREST

To discuss how our security solutions can help protect your most vital assets, please call us today.

☎ 1300 922 923
+61 (2) 9290 4444 SYDNEY
+61 (3) 8376 9410 MELBOURNE

✉ info@senseofsecurity.com.au

senseofsecurity.com.au

Sydney
Level 8, 59 Goulburn St
Sydney NSW 2000
AUSTRALIA

Melbourne
Level 15, 401 Docklands Dr
Docklands VIC 3008
AUSTRALIA