Robot Wars: The Front Line in CyberWarfare

It wasn't all that long ago (only 5 or 6 years) when cyber attacks were cosy affairs – spotty geeks sitting in their front room, seeing if they can bypass security through a newly-discovered buffer overflow exploit. Remember the heady days of angst-ridden trolls hell-bent on flooding a social media site with messages of hate and ill-will? It was a simpler time.

It used to be manageable, in the main, for security teams to identify and block even quite concerted attacks (for the age). Those days are well and truly gone. Today, a full battalion of automated bots continually bang on the doors of business, government, education and charity networks. To make matters worse, just a single prepared assailant can activate a horde of marauding attacks.
Around the globe there are hundreds of these well-funded, highly organized and automated attacks traipsing their way happily from network-to-network, planting droppers and seeking out targets for a full-blown attack.

Once a victim has been selected; a vulnerability detected; the onslaught can be relentless. The attack is automatically orchestrated from a highly-focused set of distributed robots, all peering unremittingly for every morsel of intelligence data. Every misconfiguration is ruthlessly exploited.

Detection, Detection, Detection

In defense, organizations put in place all manner of detection systems, well beyond the days of simple perimeter fencing. There are a huge number of threat detection systems growing everyday – some are good, many are excellent, and a few are top-notch. These systems do a great job at identifying threats, reconnaissance, vulnerability scans, and all manner of threats.

These systems however, all have one thing in common: they all rely on someone to watch them wave their red flag in protest at the latest breach. The trouble is, with modern robotic attacks, the detection density is so high, there are simply not enough humans on the planet to deal with such volumes of alarms, let alone security professionals.

The 'detection storm' becomes itself a kind of denial-of-service: keeping security experts busy chasing a plethora of alarms. And so, alas, the 'noise filter' step kicks in, tuning down the 'clatter' of detection so that you believe only significant attacks are getting through. This too is beautifully exploited by the obedient cyberbots, 'hiding in plain sight' as so-called 'noise'.

The Perimeter is Dead – Long Live the Perimeter!

It's fair to say the bulk of security investment across the world has been in perimeter technologies (from firewalls, to web and email proxies, layer7 URL filters and threat intelligence networks). These have had a great and positive impact in increasing the security of organizations' IT infrastructures.

Of course the perimeter is not dead – in the cyber war trenches however, the landscape has changed significantly. Perimeter attacks remain the staple diet of many attack vectors, but now they have been accompanied by a new flavour of exploitation: the attack on data. This is due in no small part to the rise in perimeter defenses, coupled with the truly measureless volume of uncontrolled and unprotectable data spewing out of organizations at a prodigious rate.

Cyber criminals are taking full advantage of this 'low-hanging fruit' – organizations are defenseless to protect their data once it has entered the wild. Somewhere between 95% and 100% of your employees will have their business email sitting comfortably in an address book inside their personal email. When that address book is compromised, or someone in their contact list has their email compromised, how can an organization protect against this? In short, it's simply not possible.

Social media sites are another emblematic habitat for leaked data. They are littered with proud declarations of employees' noble job titles and CVs, leaking huge tranches of company information into the Cloud. Determined cyber criminals find it easy to build large and valid lists of email addresses, names and internal targets. Even from just one leaked email address, a company's email format is made evident. A quick scan of employees belonging to the same organization exposes a targeted and accurate picture of a company's internal personnel structure.

Coupling Cloud-based data-leakage mining with perimeter and vulnerability reconnaissance is just one example that attackers employ to bypass the perimeter. These types of data-driven attacks only need partial penetration to be able to cause great damage. Just a small but valuable part of your network need be exploited for your Financial Director to get a call for a Ransomware demand.

The Elephant in the Room is Really a Robot

A large lumbering mass of threatening grey matter is easily handled by security organizations. But when that beast turns out to be a robotic cyborg, capable of concealing and morphing itself behind your perimeter, it quickly becomes much more dangerous. Everyone knows it's there somewhere, but the thought of acknowledging its upgraded status as a mechanized cyber-spy is difficult to fit in with the traditional method of dealing with cyber threats on a manual, case-by-case basis.

The Empirical 'Line-in-the-Sand'

Most Security teams consist of, on average, between 2 and 25 security professionals whose job is to monitor and investigate detected breaches as they arrive on an organization's doorstep. The level of experience of the team varies, but history shows the expertise is generally quite high, with at least some portion of the team being very knowledgeable indeed.

The primary limiting factor here is one of scale. Your security professionals can be highly skilled, but there's a limited number of hours in a day, and so can only handle a set number of incidents at a time. This number is typically between 20-80 incidents per day, per team.

As such, the very common practice is to 'tune-down' detection systems to fit into this number: eliminating 'noise' from the daily view. But how can anyone be truly confident that today's noise is really noise? When viewed from this perspective, it's easy to see how tuning your valuable and expensive detection investment to fit your team's size is not going to deliver visibility of the true scale of attacks on your organization. Indeed, this practice exacerbates the scalability problem by continually submerging your highly trained security staff in never-ending 'data-repression' exercises; forever attempting to keep the volume of alarms to a manageable level.

This is, of course, not a sustainable way to proceed securely into the future, where the barrage of attack vectors becomes ever more dense, and the need for them to be actioned quickly ever more acute.

In an average month, at least 20% of an organization's IT assets are likely to be attacked (or seen to be attacked) significantly enough to raise an alarm. If assets tend to the 'user-laden' side (e.g. lots of desktops and/or mobile devices) the number will be higher. If your organization does not see this many alarms, either the detection is insufficient or the detection rate is tuned down (or both).
Either way, security teams will not be able to confidently identify complex and long-running attacks, leading to large gaps in security visibility. Of this 20%, only about 1-1½% may turn out to be 'actionable' security alarms (actionable meaning something needs to be done about it straightaway by a security member).

So who needs the rest of it? Surely that's just noise and false positives? Ah, well. Put like that..you can already see the danger of the 'silent assassins' that live in the noise and false positives. These are the drones; the expendable clone-bots whose purpose is simply to gather information, get in the way of 'real' Security, and maybe even plant a dropper here and there if it gets the chance.
These 'bothersome pests' form the embryo of a 'killer-attack' in waiting – incubating and silently waiting for one to not be discovered or be ignored as noise. In order to truly minimize the 'actionable' attacks, it is imperative that the 'mosquitoes' be detected, analyzed and scrubbed. But there are so many? How can you deal with such an offensive?

Defense of the Clones

The most vitriolic of cyber attacks are spawned by robotic mechanoids. It stands to reason the best way, indeed the only way to stave off these marauders is by employing benevolent automation to 'hoover-up' these swarms of vermin: some of which will be carrying highly toxic weaponry if left unchecked.
For those familiar with the 'Automation-fabric' – this is a highly effective instrument – coupled with comprehensive detection - to protect your hugely valuable business networks from 'death-by-a-thousand-cuts'...

Delivering an automated 'immune-system' will bring the necessary scalability to your existing security processes whilst simply not allowing any security alarms to go uninvestigated. The automation-to-fight-robot-attacks model frees your security staff, not only from dealing with endless noise and false positives, also highlighting severe attacks, whilst delivering liberation from the 'security tune-down trap' of forever spending all day dampening detection systems just to survive until tomorrow, when the process starts all over again.

Scale to Achieve

Using automation to enable your security team and your security investment to achieve the scalability your organization needs to protect itself, in conjunction with a solid detection strategy, is a highly effective weapon in the new frontier of the cyber-robot wars.

About the Author

Peter Sturge leads the Software Development and Security Research Teams as CTO at Honeycomb Technologies Ltd., and has over 25 years of cyber security experience in network, IT and software security systems. Before co-founding Honeycomb Technologies in 2007, he cut his Cyber Security teeth as Senior Security Consultant at Integralis' Security Operations Centre (now NTT), and before that as lead software developer, writing and securing enterprise network solutions for Computer Associates, HP and Novell.