Testing Cybersecurity Effectiveness

An Adversarial Perspective

Prepared by:

Dynetics, Inc.

Huntsville, AL

April 12, 2019

ASSESS effectiveness

OPTIMIZE investments

TEST resilience

CERTIFY results

Effective Cyber Risk Management!

D-19-11545

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

Executive Summary

Over the last three years, Dynetics has conducted simulated cyber attacks (Adversarial Simulations) for 20 organizations including very small businesses, Fortune 1000 companies and government agencies. Each organization was investing in traditional cybersecurity (firewalls, anti-virus, spam filters, etc.) but to various degrees and with various levels of in-house cybersecurity staff and expertise. The results, however, were consistent:

- **100% of organizations were breached via social engineering.**
- **The time to breach perimeter defenses ranged from 2 minutes to 23 days.**
- **18 of 20 Adversarial Simulations demonstrated the potential for catastrophic impact.**
- **The time from perimeter breach to full compromise ranged from 40 minutes to 60 days.**

OF THE 18, NONE WERE AWARE OF THE COMPROMISE UNTIL NOTIFIED BY DYNETICS!

While these results may seem dire, they should not be viewed as an indictment of the IT staff charged with implementing cybersecurity. The results, however, do show that traditional cybersecurity approaches alone are no longer effective against motivated hackers—hackers who are willing to target an organization and be intentional about their objectives. Therefore, as a result of conducting these Adversarial Simulations, Dynetics offers the following perspectives:

- **The primary reason a successful cyber attack leads to significant or catastrophic impact is the inability to detect malicious activity on endpoints (workstations and servers) behind a firewall.**
- **The time available to detect an intrusion and avoid catastrophic financial impact is directly related to the effectiveness of internal cybersecurity controls.**
- **Perimeter defenses are necessary to stop random, unsophisticated attacks, but are ineffective against social engineering tactics used by motivated hackers.**
- **Most IT staff are unfamiliar with techniques motivated hackers use to defeat traditional cybersecurity and avoid detection.**
- **Most IT staff have too many responsibilities to effectively monitor for intrusions.**
- **Effective cybersecurity can be achieved with layers of controls that include internal protection and detection in addition to perimeter protection.**

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

Anatomy of an Attack/Adversarial Simulation

An Adversarial Simulation is a "threat faithful" simulated cyber attack that aggressively tests the effectiveness of existing controls and demonstrates the potential impact of a real cyber attack. Testing in this way gives a customer a glimpse into what the outcome of a real attack might be but in a controlled environment where the "adversary" is working for the customer. This approach to finding gaps is much less stressful than finding out during a real attack.

An Adversarial Simulation begins with a Rules of Engagement call with a select few customer participants to establish a timeline (typically 90 days), identify any constraints, and identify data/systems that will be targets of interest for Elite Ethical Hackers. For most organizations, targets typically include customer/employee data and/or intellectual property.

After completing the rules of engagement (ROE) call, Adversarial Simulations proceed along the Attack Lifecycle (Figure 1) with little to no customer interaction. Elite Ethical Hackers plan a targeted attack against the organization, usually by gathering open source intelligence about the customer from publicly available resources (e.g. website, LinkedIn, Facebook, etc.). This information is used to craft an attack scenario (typically phishing) that has a high likelihood of success. In the majority of cases, the first attempt successfully breaches the network perimeter. After breaching the network perimeter, the Hackers' first action is to establish persistence so the attack can survive a reboot or logoff. Next, the Hackers begin looking for objectives and identifying which key individuals might have access to the data/systems of interest. Hackers look to escalate their privileges and pivot to other workstations until they finally gain access to the data/systems that were identified to them by the customer. The advantage of an Adversarial Simulation over typical penetration tests is that an Adversarial Simulation follows each phase of the Attack Lifecycle and accurately mimics breaches reported in the news.

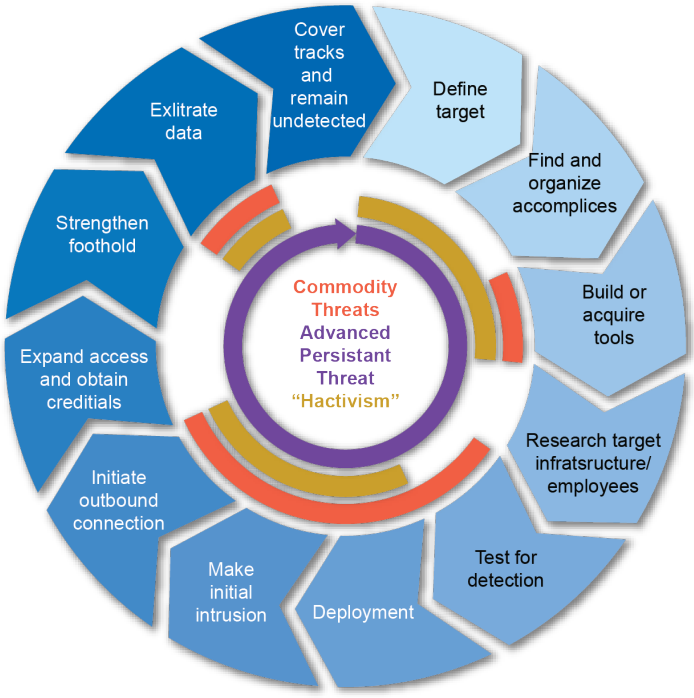

Figure 1. Attack Lifecycle

KEY POINT

The time to move through the phases of the Attack Lifecycle can take months to years. The recent Starwood breach was active for four years before it was discovered. To provide beneficial feedback in a reasonable timeframe, Adversarial Simulations are necessarily constrained to about 90 days. However, this still allows time for the Hackers to discover ways to defeat controls and avoid detection using the same techniques as malicious hackers.

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

Results from Testing Cybersecurity Effectiveness

Over the past three years, Dynetics has performed 20 Adversarial Simulations for organizations and was 100% successful in bypassing perimeter security using social engineering tactics such as phone calls and emails with "malicious" attachments or links. Typical employee targets for social engineering included Customer Service, Human Resources, Sales, and Technical Support. 18 of the 20 Adversarial Simulations resulted in full compromise of the environment, demonstrating potentially significant or catastrophic impact through the exposure of customer data, employee data, and compromise of the key information assets.

The common thread among all the social engineering attacks was the nature of the targeted employee's job. The role of each targeted employee demanded they spend most of their time interacting with unfamiliar individuals. This makes it extremely difficult to address this problem with training alone. For example, a popular training theme is "Never open anything from someone you don't know!" However, when the job requires interfacing with the unknown public, it's unrealistic to expect employees to make the right decision every time. Some common themes Hackers use are social engineering scenarios:

- **Customer Service**
  "I'm moving to your area. Can you tell me if these addresses are serviceable?"
  Email Payload: malicious link to property addresses
- **Human Resources**
  "I would like to be considered for the plant manager position that you have open."
  Email Payload: malicious attachment with resume
- **Sales**
  "I am moving my business to the area and looking at office space. What are my options?"
  Email Payload: malicious link to office space
- **Technical Support**
  "I am having trouble getting to this website. It worked yesterday. Can you get to it?"
  No payload required (just a simple phone call)
- **Any Role (Gift Card)**
  "Fill out this short survey and you will receive a free gift card."
  Email Payload: malicious link to survey

Once the perimeter breach occurred, Hackers escalated their privileges within the environment by taking advantage of vulnerabilities, weaknesses or configuration mistakes. Common issues across the 20 Adversarial Simulations included:

- **All employees had administrative access to all other workstations**
- **Insecure logon/startup scripts could be modified**
- **Clear text passwords were stored in Active Directory**
- **Clear text passwords were stored on file share**
- **Trivial passwords**
- **Reuse of local administrator account passwords**

Once privilege escalation occurred, the majority of organizations had no restrictions on lateral movement behind the firewall. This allowed the Hackers to easily move to discover key information assets.

KEY POINT

The actions taken by the Hackers to move through the Attack Lifecycle occurred on endpoints (workstations and servers) and could have been detected long before critical data/systems were compromised. Unfortunately, in most cases, no detection controls were deployed. In other cases, indicators that were collected were not monitored and therefore not acted upon.

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

Case Studies

The following case studies provide more details on how three organizations were compromised and what the impact of the breach could have been had these been real attacks.

Case Study 1

Dynetics placed a phone call to the main customer service number under the false persona of "Stan," who was seeking to move to the area. This phone call established some level of trust because now the victim was expecting an email from "Stan" with a list of addresses to lookup. An email was then sent with a "malicious" link and once the link was opened, "Stan" had internal access to the network just as if he were an employee. "Stan" then used his access to enumerate as much as possible within the environment looking at their Intranet site and file shares. On one of the file shares, "Stan" found a readable file that contained clear text passwords. One of those passwords gave "Stan" root access to a Linux server in the management network. With root level access on the server, Stan enumerated another account on the Linux server and extracted the password hash. This hash was cracked in less than 2 days. This account password was reused within the service provider's Microsoft Windows network and had Active Directory Domain Administrative privileges. Using this account, "Stan" was able to fully compromise customer, employee, and accounting data. This could have led to identity theft of customers or a potential blackmail scenario where the attackers try to extort money in exchange for not releasing the information. A key takeaway from this event was that the duration between "perimeter breach" and "data breach" was 60 days. Had there been a good detection strategy in place, "Stan" would have been caught long before any real damage was done.

Case Study 2

Dynetics sent an email with a "malicious" link to a survey that, if completed, would provide the survey taker with a $20 iTunes Gift Card. In less than one day, an employee had opened the survey granting Dynetics internal access to the network inside the firewall with the same level of access as the victim. Dynetics began enumerating the network and discovered that there was a script that ran regularly as SYSTEM ("the highest level privilege on a windows workstation"). A misconfiguration was discovered in the permissions of the file that allowed anyone to make changes to the script. Dynetics modified the script so that it would execute additional code, giving Dynetics control as well as continue to provide its intended function. As the script continued to run on its normal schedule, each machine within the provider's network became compromised, giving Dynetics full control of each asset. Some of these machines belonged to Network Engineers. One of these workstations had a KeePass database along with the password to the database located in the same folder. With the ability to open this password database, Dynetics was able to gain full control over the core customer facing network. This could have led to a takeover of the customer network resulting in a service outage for customers. The attacker could potentially demand money for returning control back to the provider. Similar to the first case study, Dynetics persisted for approximately one week before demonstrating impact to the customer. With a solid detection strategy, Dynetics should have been discovered before any real impact was realized.

Case Study 3

Dynetics sent an email with a "malicious" link to a survey that, if completed, would provide the survey taker with a $25 Amazon Gift Card. This company had fairly aggressive web filtering in place that blocked all "unknown" categories. Since our link was to a site that had just been registered, it was categorized as "unknown" by the target company. However, this particular user was working from home and accessed

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

the site via her own Internet connection. This gave Dynetics control over the company workstation but not to the company network. Upon accessing the workstation, Dynetics discovered that the company was enforcing a certificate-based two-factor VPN. However, the certificate was configured to be exportable. Dynetics exported the certificate and then coerced the user into providing their username and password by launching a fake authentication prompt. Now Dynetics had all the pieces to the puzzle in order to establish their own connection to the target company network. Once inside via VPN, privilege service accounts were enumerated and cracked resulting in full domain compromise. This particular company had invested in excellent preventative controls (aggressive web filtering, two-factor authentication, etc.) but had not invested in any detection. Once Dynetics found a way in, they were able to move freely and the organization was not aware they had been breached until Dynetics called to schedule an outbrief.

Conclusions

Most organizations are not prepared to successfully defend against a motivated hacker. The lack of preparedness is not due to lack of interest, lack of investment, or lack of qualified staff, but to a lack of awareness of techniques used by motivated hackers. While most will acknowledge that employees make mistakes that will lead to a breach, most are not aware of the extent of damage that can be caused if an attack goes undetected; and most have no visibility into malicious activity inside their firewalls.

By publishing this whitepaper, Dynetics is attempting to raise awareness of the steps necessary to address the limitation of traditional cybersecurity approaches. Dynetics firmly believes that businesses can be successful if they change their perspectives to assume perimeters breaches are inevitable and deploy solutions that can detect malicious activity before incurring significant or catastrophic impact.

KEY POINT

When the results of an Adversarial Simulation are presented, organizations have a tendency to fixate on the specific issue the Hackers exploited. However, motivated hackers (and Dynetics Elite Ethical Hackers) take advantage of whatever circumstances they encounter and pursue the first successful attack vector. In these 20 Adversarial Simulations, the attack vectors were different each time, and would likely be different if the Adversarial Simulation was repeated. While known issues should be addressed, it's more important to ensure you are deploying a comprehensive, layered approach that will successfully protect against and detect all potential attack vectors.

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary

About the Authors

Craig Mitchell is a Principal Cyber Analyst, specializing in network/system security and enterprise support. His role at Dynetics includes designing secure solutions as well as acting as trusted advisor for many commercial and government organizations. As an Elite Ethical Hacker, Craig conducts reconnaissance, vulnerability analysis, penetration testing, and adversarial simulations. He has provided classified support to multiple national security programs. Craig has a B.S. in Management Information Systems and holds the following certifications: OSCP, CISSP, MCITP, MCSE, CCNP + Security, GSEC, CEH.

Robert Dowling is the Cyber Risk Practice Lead for Dynetics, where he is responsible for developing cyber risk management solutions to address business and technical risks and serves as the primary interface to customers, business partners, and associations. Robert is a frequent speaker on cyber risk at many NTCA and regional rural broadband events. He supported systems and software engineering for NASA and Defense programs before moving into business development roles.

FOR MORE INFORMATION

Dynetics Corporate Communications
CorporateCommunications@dynetics.com
(256) 964-2000

© 2019, Dynetics, Inc. All rights reserved.
Dynetics Proprietary