Red Team Assessment Whitepaper

CyberCX

The business case for conducting Red Team assessments

Introduction

Risk management is a discipline with an extensive heritage. At its core, it requires a high degree of visibility so executives and boards are able to make informed decisions.

A key challenge facing anyone involved in contemporary risk management is how to appropriately assess cyber risk. With technologies and threats constantly evolving, adopting a dynamic approach to risk assessments is the optimal way to ensure you have the visibility you require.

In this whitepaper we will present the case for conducting offensive red teaming activities as part of a broader dynamic risk assessment.

We will explore the planning and execution of red teaming activities, particularly within the context of the critical infrastructure sector. You will also gain invaluable insights into aligning risk management with technical assessments that can be applied to almost any organisation.

01

Risk assessments – what type?

While a board has many obligations including financial, legal and social, their primary focus, simply put, is risk management. Good risk management is directly related to the quality of the data that is presented by the executives.

Cyber security is largely a new domain for many executives, directors and stakeholders to consider. One of the issues is how to describe cyber security risks effectively. Historically, one of the most popular approaches has been through Risk Assessment. Risk Assessment itself is a specialised area, and organisations are likely to have internal staff dedicated to such a role or may go to the market to appoint a specialist firm to conduct a risk assessment.

Traditionally, risk assessments have been largely paper based. The Assessor will determine a scope relevant to the objective, and then undertake the assessment using a methodology described in one of the various approved international standards on risk management.

There are, however, many different approaches to risk assessment, and the standards are not prescriptive on exactly what or how to undertake such an assessment.

For example, an Assessor may request various documents, review the previous risk assessments and the controls that were nominated, and then determine if the controls are in place and working. The Assessor also needs to ensure that the risk appetite of the organisation is current and accounted for. A business cannot mitigate all risks. The business needs to make a call on what level of risk is accepted and where controls are required to bring higher levels of risk down to acceptable levels. This is risk appetite.

The problem here is that the scope of the assessment is integral to the overall outcome. Choose a narrow scope and you will only assess part of your business, and the possibility exists that there are risks that you have not adequately accounted for. This problem is exacerbated when the organisation is tasked to determine its Cyber Risk Profile – that is the extent to which it may be subject to incurring losses as the result of a range of cyber oriented attacks. These could include, but not be limited to, Denial of Service (DoS), hacking, data loss, malware, operational down time (unavailability of systems), regulatory or legal compliance matters etc.

The reason why the problem is exacerbated, is there are so many vectors through which an organisation is exposed to cyber risk, and the threat landscape changes at high frequency.

Using a traditional approach that includes documentation reviews, interviews, analysis of previous assessments etc., while consistent with the standards, does not mean the Assessor is reviewing the ability of the organisation to defend against current and future threats.

Take the organisation's exposure to social media borne attacks as a good example. Social media exists on the periphery of most organisations. The business does not operate its own social media platforms, rather it consumes these as a service (e.g. LinkedIn). Further complicating the issue is overlap in cyber space between personal and business. A targeted attack on an employee (through their personal social media activity) could have a negative influence on the overall corporate security exposure. Simply put, the result of a single employee's breached social media account, such as Twitter or Facebook, could lead to an overall compromise of a corporate network. This is not far-fetched. There are many examples of such a domino-effect attack. Ironically, Twitter themselves succumbed to such an attack as far back as 2009[1].

Another good example of risks not commonly assessed through paper-based reviews is Supply Chain Risks. An organisation may be at risk through their relationships with other entities in their supply chain.

Consider a business that has outsourced its IT Management to a third-party company. They require remote access and administrative accounts in order to deliver their services (backup, server and workstation management, help desk functions etc). Now if the third-party is breached, the attacker could use the established privileged access channels, and the associated established privileged access credentials, to ultimately compromise the targeted end company. The Wipro breach of 2019 is likely to be a landmark case for this attack vector. There are also many other forms of supply chain risks (services, software development, hardware development, all forms of outsourcing etc).

As highlighted above, to effectively perform cyber risk assessment, you must think about the complex web of overlapping systems and supply chains that are now essentially part of the footprint of your business.

1. 2018 Cost of a Data Breach Study: Australia, Ponemon Institute LLC, July 2018

02

Why does all of this matter?

If the Risk Assessor is not familiar with the techniques that attackers are actively using today, then it is unlikely that the assessor will be covering the full scope of risks that the organisation is exposed to. While this alone does not make the traditional risk assessment irrelevant, it certainly negates the relevance of relying on it for your business' overall risk identification method. If you do not identify the risks you will not be able to apply controls to mitigate them. These traditional risk assessments are still important as part of an overall risk management program, and they are still required under many of the standards. However, given the changing nature of cyber risk, assessments must now be augmented with other methodologies in order to assess risk more dynamically.

03

Welcome Dynamic Risk Assessments

We need to adopt assessment methods that are going to give a higher degree of assurance that we are identifying realistic vectors through which the business may be subjected to attack. At CyberCX these dynamic risk assessments often take the form of red team assessments. Red team assessments are conceived on the premise that we are not stifled by prescriptive rules, a spreadsheet listing the methods we must use, or having meetings with any number of stakeholders who may not disclose the true state of affairs for fear of losing credibility.

We have a systematic approach to conducting the assessment. We will profile your business identifying all the relevant information about your suppliers, your staff and executives, your facilities, your email, web sites, remote access and financial systems. All of this will be profiled through publicly available information. We will devise an approach to assess all these vectors to determine those most susceptible to attack – that is, the methods that attacker will most likely adopt if targeting your business.

The reason why we call it a dynamic risk assessment, is because throughout the process we utilise constant feedback loops. This allows us to tailor our assessment to each business as we learn more about the strength of controls, and hone in on weaknesses that we identify which may yield wider value to an attacker.

Dynamic risk assessments take into consideration vulnerabilities in your supply chain. A breach at one of your trusted suppliers can have a domino effect through your organisation.

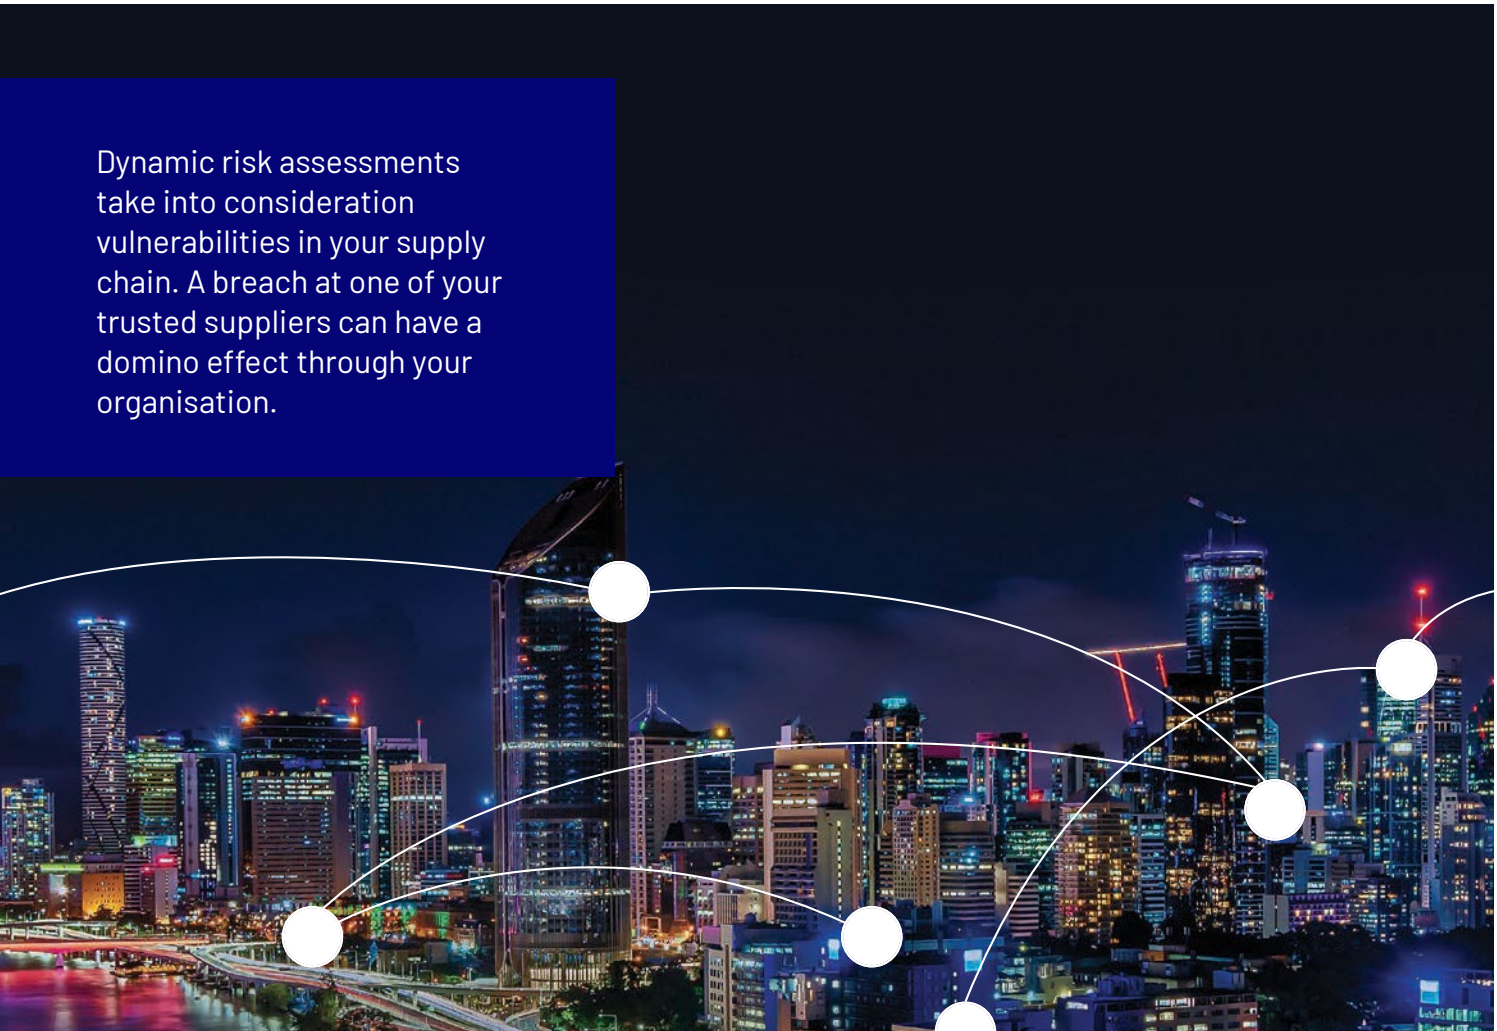

04

Red Teaming

Case Study on an owner & operator of critical infrastructure

The target

A large organisation providing critical services and handling private and extremely sensitive data.

The goal

Through a consultative approach, our client defined a number of objectives to test the resilience of their existing controls to protect their systems. The client has invested heavily in technology, insourced and outsourced services and believed they covered all cyber security bases. The red team assessment is intended to validate the effectiveness of the current controls. These were the goals defined by our client:

1. Gain access to the network
2. Compromise the domain
3. Locate, access and exfiltrate the primary datasets
4. Gain access to the main Transport Layer Security (TLS) keys responsible for various encryption functions
5. Compromise isolated systems responsible for delivering critical services

The overall red teaming assessment approach

During the scoping and planning phase the assessment was broken down into three main phases,

- Reconnaissance phase
- Persistence phase
- Results

Reconnaissance phase

The mission was to gather intelligence on the target, from which to launch attacks, including:

Enumerate the external perimeter as much as possible (IP space, DNS records, Exposed services, technologies in play, SaaS/PaaS etc);

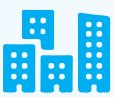

Perform physical reconnaissance of the corporate offices and identify entry points for gaining entry; and

Perform intelligence gathering on as many employees as possible;

Identify possible WiFi networks used by the client.

Persistence phase dimensions

The mission was to gain access by any means and setup persistent access back into the network. Points of attack that were identified included

WiFi

possibly enabling us to gain access to a guest or corporate network from which we may be able to pivot our attack

Social engineering

possibly finding someone who would provide more information to enable an attack to progress further

Tailgating/physical access

possibly enabling us to gain access to their internal networks from where we can escalate privileges and pivot the attack to target systems

External network access

assessing external systems for vulnerabilities that may enable us to gain access into internal systems or may provide information to otherwise further our attacks. External systems include remote access and any cloud-based systems (such as cloud email)

Phishing

running several campaigns to gain sensitive information from employees or encourage them to download and install files that may assist us gathering information about the computing systems or possibly gain remote access into the environment

USB drops

possibly creating a scenario where someone would install our apps into the environment, bypass endpoint controls and create a remote access channel for us

Step 1

Reconnaissance results

We extensively enumerated the organisation's external perimeter – that is the footprint that defines them on the internet. Like many organisations, they have migrated to Office 365 (O365) for email and other hosted Microsoft products (such as OneDrive and SharePoint). They have a large number of employees and we recursively gathered more and more information about them. We determined that the Corporate WiFi is accessible outside of the physical tenancy. We also identified two doors that may be targeted to gain physical access into their office.

The path to persistence

Access to O365 was obtained via password spraying using usernames identified during the reconnaissance phase. O365 was protected using Multi-Factor Authentication (MFA) on the user interface. However, MFA was not enabled on both Exchange PowerShell and the Exchange Web Services (EWS) interfaces. Having access to email allowed us to impersonate users, gain credentials and we gather contextual content to use for social engineering and phishing. Overall, the perimeter of the organisation is quite solid so there were no other means of compromising them remotely.

WiFi was protected by EAP TLS, with very few users using it. Within the time allocated, an attack on the Wireless network would not be feasible even though it was broadcast beyond their perimeter.

Tailgating was trivial and allowed us to implant a small device (called a LAN Turtle) on to the network in a surreptitious location. However, this did not provide us the remote access we needed because it was determined that the organisation is running a robust deployment of Network Access Control (NAC) through 802.11X requiring enrolment into their platform in order to use the network. We had to go back a second time to see if we could bypass this. While trying to be discrete and unobtrusive we identified the room where computer repairs and provisioning is undertaken. In this room we had an inclination that some network ports may not be protected, for the convenience of the technicians in order to build and test laptops before issuing them to staff. Voila. We found one port for which NAC was whitelisted, we quickly connected our LAN Turtle to it and established a back-channel connection to our control platform (via 3G mobile internet).

This has resulted in Goal #1 being achieved.

Step 2

Gaining administrative access to the domain

While still in the office, we timed our next activity while the IT room is unoccupied at lunch time.

We gain access to unused desktop machines. While the organisation has been quite diligent with their security controls to-date, they have not been as thorough as needed. The Standard Operating Environment (SOE) does not have boot time protection. As a result, we can boot a desktop with one of our pre-compiled Linux distributions through a USB key. We can extract cached credential hashes deep from within the Microsoft Operating System, and can then perform attacks against these credentials offline. One of the credentials turned out to be a domain administrator, the highest level privilege in a Windows based domain. This was their next mistake. Like many organisations, technicians use elevated privileges to perform standard tasks. Using the domain administrator account, we were able to expand our attack. We extracted all the user accounts in Active Directory and proceeded to crack them.

This has resulted in Goal #2 being achieved.

Step 3

Up until now, all our activities went unnoticed and unchallenged. We knew what the goal target data was, but had no idea of the format or location. Internal network reconnaissance did not yield any clues. We traversed public network shares searching company documents. We discovered an internal document describing Standard Operating Procedures (SOP). Clearly our client was very diligent. Many standards require such documentation to exist, but they never expected that such a document may reveal clues to enable an attacker. This particular SOP related to provisioning access to the target data set! We now knew that it was a Linux server requiring SSH access. The SOP contained the email address of the person responsible for the access to the server. We had cracked his AD password and due to password reuse we gained access to the Linux server. Another careless mistake! Re-using passwords is commonplace, convenient but downright dangerous. We found the encrypted and password protected data files and copied them off the server.

Further utilising the domain admin that we had achieved, we worked our way laterally to an online LAN password manager application. And here we found the key we needed to unlock and decrypt the primary datasets.

This has resulted in Goal #3 being achieved.

Step 4

Repeating the same process, we located the primary TLS keys and their pass phrases.

This has resulted in Goal #4 being achieved.

Step 5

In the network we continued to roam around and looked for more and more information that could further our attack. We identified some Design and Change Control documents that listed the separate networks hosting the goal services. The password vault allowed us access to the VPN certificates to gain access to the network via the secure VPN. That password vault was a gold mine! However, we still had some hurdles to overcome. The network, and its services, deep within the organisation were on the operational side of the business. They were not tied to Active Directory, because they are not not really corporate systems and were intended to be isolated. We were stuck because we had no credentials for this part of the network. But we were not defeated yet. We identified the administrative users for the target systems. In accordance with policies, the users stored their passwords using KeePassX password manager. Using access to the Roaming Domain Profiles we were able to recover the KeePassX key database files with the certificates. Unfortunately, we were not able to crack the passwords for the password database files. We returned to the Roaming Profiles and inspected Mozilla Firefox saved password database files. In these files we recovered web administrator console passwords for the target systems from saved Firefox entries. This was another avoidable mistake. Browsers often prompt users to save their credentials directly in the browser for convenience. But this is never as secure as an independent password safe. We subsequently cracked the cached credentials and had root access to the most valuable system in the business.

This has resulted in Goal #5 being achieved.

Living off the Land

It is important to note that all the goals were achieved without exploiting any vulnerabilities. This is a true live-off-the-land attack. We explore, find elements of interest, attack, escalate, pivot and move deeper and deeper within the network, using only the tools and systems that the company has deliberately configured.

06

Take Aways

Please think about the following take away insights for conducting your future red team assessment:

1. Context is all important. You need to understand your environment, key data and key concerns and risks.
2. A risk assessment is only as good as the scope that you are looking at. Take a dynamic view of risk assessment, and ask the right questions, not just the standard set of questions.
3. Attacks normally exploit technical weaknesses and people. It is important to consider people, process and technology in equal measure.

07

Conclusion

Whilst organisations routinely undertake risk assessments to help them develop cyber resiliency, it is essential that such assessments are based on reality. With technologies and attack vectors changing constantly, obtaining a realistic understanding of your cyber security posture, and the extent to which you may be susceptible to attack, can be a challenge. Adopting a dynamic approach to risk assessments can help you account for a rapidly evolving security landscape.

Red teaming is an effective cyber security activity that allows you to dynamically assess risk, as it is based on realistic scenarios tailored to your individual circumstances.

This case study has demonstrated the true value of red teaming and how it can be instrumental in taking a more dynamic approach to risk assessments. Whilst a traditional paper-based assessment may identify some risks (for example no boot control on the SOE, or MFA not applied to all O365 interfaces), it would be unlikely to accurately demonstrate the implications of all the risks that collectively may compromise critical infrastructure.

The next time you are asked if your business is cyber resilient, think about the approach we have outlined above. Put a red team assessment firmly in your agenda for your risk management program.

CyberCX

Written by:

Murray Goldschmidt - Executive Director, Cyber Capability, Education & Training

Willem Mouton - Principal Security Consultant

More Information

For more helpful information, refer to the free resources on the CyberCX website or feel free to reach out to the CyberCX team, who will be happy to answer your questions.

www.cybercx.com.au

1300 031 274

About CyberCX

CyberCX is Australia's leading independent cyber security services company. Unifying the most trusted cyber security brands and the experts who built them.

CyberCX delivers end-to-end cyber security services and Australia's best cyber security talent with the most comprehensive range of cyber security services to business, enterprise and government.