Red Team testing: Essential KPIs and metrics

Received (in revised form): 24th January, 2024

Richard Hollis
Chief Executive Officer, Risk Crew, UK

Richard Hollis is the Chief Executive Officer (CEO) for Risk Crew, a product-agnostic information security risk management and cyber security testing consultancy headquartered in London. He possesses over 25 years of hands-on skills and experience in designing and delivering creative, cost-effective Red Team exercises to a wide variety of international enterprises across all sectors. Richard is a celebrated public speaker and seasoned trainer. He has presented to hundreds of audiences across the world on a wide variety of cyber risk management and testing topics and techniques. As a recognised industry authority, he has published numerous articles and white papers and appeared on national and international broadcast news shows as well as being cited in a wide range of press including the BBC, MSNBC, Radio 4, *Financial Times*, *Time* magazine and various others. He is also a regular contributor to industry podcasts and publications such as *Wired*, *SC*, *InfoSec* and *Security Penetration Testing* magazines.

Risk Crew, 169 Tower Bridge Road, 5 Maltings Place, London, SE1 3JB, UK
Tel: +44 (0)203 653 1234; E-mail: richard.hollis@riskcrew.com

Abstract This paper first discusses the diverse types of Red Team engagements and established exercise and attack frameworks and the importance of documented rules of engagement, then details the difference between metrics and key performance indicators (KPIs), and finally identifies the essential data points to be captured and compared in your next Red Team Test.

KEYWORDS: Red Team testing, Blue Team testing, Purple Team testing, TTP, exercise framework, attack framework, rules of engagement, KPI, metric

INTRODUCTION

While Red Team testing is a significant financial investment, most businesses fail to measure its actual return. Has it improved the organisation's cyber security posture and if so, how? Is there more value in conducting Red Team testing than conducting security penetration testing, and if so, why? Answers to these fundamental questions are easily obtained by including the requirement to capture essential metrics and key performance indicators (KPIs) in your next test.

APPLES AND ORANGES

First things first. Conducting Red Team testing is not the same as conducting security penetration testing. They are two entirely different approaches to testing and as such, the return on investment (ROI) is incomparable. Penetration testing (or ethical hacking) seeks to identify (and exploit) security vulnerabilities in the business's technology systems. Red Team testing seeks to specifically duplicate a threat actor's methodology to identify (and exploit) vulnerabilities in the business's people, process *and* technology based upon open-source intelligence. Apples and oranges.

Red Team testing is far and away a more holistic approach that results in the following returns.

Identifies more weaknesses

The most obvious benefit of conducting Red Team tests is the ability to identify

holes and vulnerabilities in your defences. While security penetration tests are design to identify vulnerabilities in your technology systems, Red Team testing engagements are designed to identify vulnerabilities across your business in people, process and technology associated with your information assets, which, if exploited, could expose these assets to a threat actor. This type of testing allows the business to identify controls in people and processes that may not be effective, which, of course, is critical in understanding the threat to the business. Red Team testing allows you to know what you do not know.

Confirms more strengths

Most businesses believe that the primary value of security testing is that it identifies 'what is not working'. While this is true, what is often overlooked is that security testing also clearly identifies 'what is working'. Red Team testing will confirm the effectiveness of the controls you have implemented across the business to prevent unauthorised access to your systems and information. Red Team testing allows you to confirm what you think you know.

Challenges your defenders (Blue Team)

Can your business withstand a real-world attack? The answer largely depends on the capability of the staff you are relying on to identify, respond and contain an attack. The staff entrusted with defending your business are your first line of defence. Are they ready? Are they capable? Do they have the right tools? The right training? Red Team testing serves as a valuable training exercise in confirming that your defensive team is 'fit for purpose'. Like holding a fire drill, Red Team testing confirms your team's readiness for the real thing.

Improves your response

The point of all security testing should be to be able to confirm that you can actually identify an attack. But the second objective is to be able to verify if you can quickly and appropriately respond to one. Response times are critical in understanding your organisation's vulnerability to a real-world attack. Red Team engagements confirm attack response times based on your defences and defender's capability. These KPIs should be captured and evaluated again in future tests. Practice makes perfect in Red Team testing.

Identifies your external needs

By evaluating the capability of your defenders and benchmarking their response times, Red Team testing also highlights weakness in in-house resourcing such as real-time monitoring, incident response, containment or digital forensics. Red Team testing will identify those areas where expertise is not in-house. Additionally, by subjecting the controls implemented in the people, process and technology across the business for effectiveness, Red Team testing identifies areas that can most cost-effectively be outsourced to topic-specific professionals, such as information security awareness training. Red Team testing confirms when you need a specialist.

Confirms the effectiveness of your products

Businesses assume that the cyber security products and solutions they have invested so heavily in actually work. The answer may surprise you and certainly should be included in the scope of your Red Team testing. Assessing the effectiveness of the cyber security controls provided by your vendors is a major benefit in conducting Red Team testing. Think of it as an independent quality assurance of your security spend. Red Team testing can confirm your cyber security product budget was well spent. Is there a more significant benefit than that? Red Team testing confirms product effectiveness.

Benchmarks your policies

Finally, if you are like most businesses, you have not updated your policies since the day they were written. Given that the only constant in cyber security is change — technology, people, process, risks, threats, vulnerabilities, tools, etc. — the assumption that the businesses security policies do not need updating is a crucial blunder.

Good news is that your security policies can easily be assessed to ensure they are current and applicable if you include them in scope of your Red Team testing. Good Red Team testing can illustrate both non-compliance and where policies are missing or incomplete.

You do not get these benefits from security penetration testing. These and many other returns are only derived from conducting good Red Team testing and are easily demonstrated by implementing clear and concise KPIs and metrics across the players, rules, exercise and attack frameworks used in the testing.

THE PLAYERS

To begin to clarify what you want to measure and why, it is critical to know the distinct types of teams and exercises (see Figure 1).

Red Team

A Red Team is a group of professional ethical hackers who research, design and execute a series of coordinated technical, physical and social engineering attacks on an organisation's information processing systems and associated people and operating facilities, to simulate how a threat actor could bypass existing security controls to obtain unauthorised access to its systems or information assets or other specified objectives. They play the role of offence.

The objective of the Red Team is to subject the target organisation to the tactics, techniques and procedures (TTPs) used by malicious threat actors in order to confirm that the organisation's existing information and cyber security controls are fit for purpose. Red Team KPIs and metrics then generally measure the effectiveness of your controls against threat actors.

Blue Team

A Blue Team is a group of designated stakeholders within the business responsible for identifying, minimising and managing the impact of a cyber security attack from a malicious threat actor. This could be security operations centre (SOC) staff or any group of individuals responsible for detecting, assessing and appropriately responding to a cyber security attack across the business. They play the role of defence.

The objective of the Blue Team is to identify and counter the TTPs used by malicious threat actors to exploit vulnerabilities in the business's systems, people and operating facilities. It is essential then that the Blue Team have a view of all the security controls implemented across the business to prevent unauthorised access. Blue Team KPIs and metrics measure the effectiveness of your monitoring and incident response capability.

Purple Team

The term Purple Team is used to describe a joint exercise wherein a Red Team and a Blue Team work against each other in unison, side-by-side, transparently to improve the cyber defensive security posture of the business. It is an open book test where the offensive team identifies the attack vector and the TPPs used in advance to the defensive team to evaluate their capabilities for responding appropriately.

The primary objective of a Purple Team exercise then is knowledge transfer. By working together transparently with the attackers, the Blue Team are given the opportunity to enhance their understanding

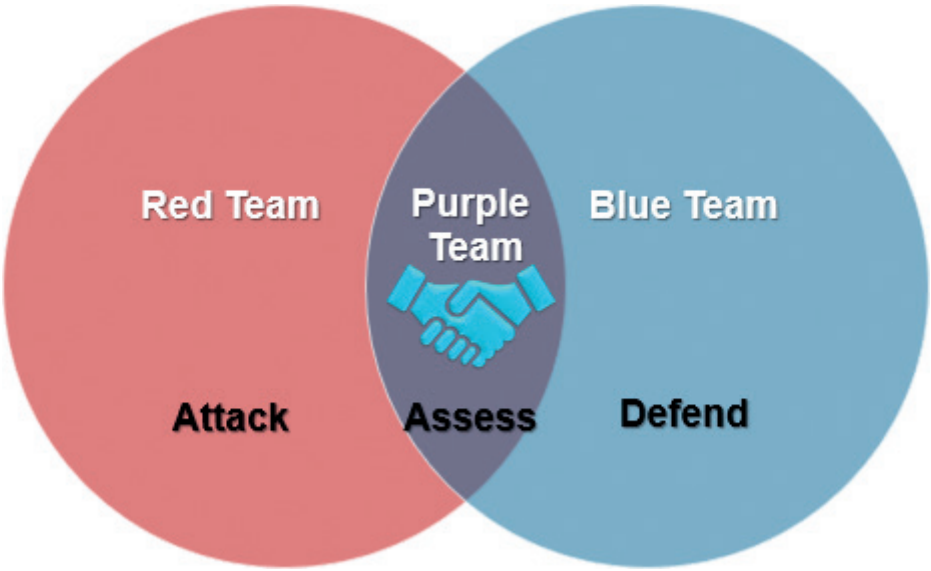

Figure 1: Team relationships and purposes

of threat actor TTPs: what they look like, where and when they are identified, what vulnerabilities they are exploiting and the actual impact they may have on the business. KPIs and metrics selected for a Purple Team exercise should then be able to assess and measure both strengths and weaknesses and provide quantifiable data used for a maturity roadmap.

THE RULES

Regardless of the type of Red Teaming exercise chosen, all engagements should follow clear and documented rules of engagement (ROE). The ROE establishes the responsibilities, relationships and guidelines among the Red Team, the customer, the system owner and any stakeholders required for the execution of the exercise.

The ROE details and documents the purpose of the test, its specific goals and objectives, along with which attack or exercise frameworks should be used in testing. It identifies the threat actors and associated TPPs to be emulated, defines and documents the exact scope and depth of the testing, and documents key stakeholders and points of contact on all sides along with applicable escalation procedures and contact details.

Red Team testing ROE establish the ground rules: exercise goals and objectives, flags to be captured, what tools can and cannot be used in testing and how far testers can go. The ROE establishes what is 'in bounds' and what is 'out of bounds'. It is all documented in minute detail and should be agreed on by signature and attached to the service level agreement (SLA) approved for the engagement.

Most importantly, the ROE is the document that specifies the precise KPIs and metrics selected to be captured during the exercise. It is all spelled out here.

EXERCISE FRAMEWORKS

Every Red Team engagement should follow a recognised framework for conducting the exercise. Exercise frameworks delineate the phases of the exercise and associated activities within each to ensure a comprehensive approach is undertaken.

Red Team exercise frameworks may be industry and geographic specific and you should ensure that you identify the

framework applicable to your organisation. The financial industry, for example, is notable for requiring Red Team testing to comply with specific established frameworks such as the following:

- G7 Fundamental Elements for Threat-led Penetration Testing.[1]
- CBEST Intelligence-led Testing — Bank of England.[2]
- Threat Intelligence-Based Ethical Red Teaming — TIBER-EU.[3]
- Framework for the Regulatory Use of Penetration Testing and Red Teaming in the Financial Services Industry — Global Financial Markets Association (GFMA).[4]
- Financial Entities Ethical Red-Teaming — Saudi Arabian Monetary Authority.[5]
- Red Team: Adversarial Attack Simulation Exercises — Association of Banks of Singapore (ABS).[6]
- Intelligence-led Cyber Attack Simulation Testing (iCAST) — Hong Kong Monetary Authority (HKMA).[7]

Exercise frameworks establish best practices for preparing for, executing and closing down the tests. They are intended to provide potential buyers with confirmation that the testing supplier possesses the qualifications, skills and experience to deliver effective testing and follows a methodical, step-by-step approach to ensure a detailed, comprehensive and repeatable test.

Figure 2 is an example of the Tiber-EU exercise framework.

The buyer is primarily responsible for determining the framework to be used for the exercise, so you need to do your research. Understanding the framework to be used for testing and the individual components and associated activities occurring within each phase is critical in identifying when and where the KPI or metric you have selected should be collected.

ATTACK FRAMEWORKS

Additionally, every Red Team should follow a framework for the attacks conducted in the exercises. Whereas exercise frameworks provide a phased approach for conducting the overall test, attack frameworks establish a process for conducting each and every individual attack in the exercise.

Attack frameworks define the step-by-step activities testers should implement in all the stages of the attack, from gaining access and escalating privileges to data exfiltration, obtaining command and control and exiting the target systems.

There are several industry-recognised attack frameworks such as:

- *MITRE ATT&CK*: Commonly acknowledged as the industry standard, it defines the terminology for TPS.[8]
- *Lockheed Martin — Cyber Kill Chain*: This popular framework details how threat actors work and the steps they perform during a breach.[9]
- *Unified Cyber Kill Chain*: This academic framework brings together several differing cyber kill chain methodologies for a more unified approach.[10]

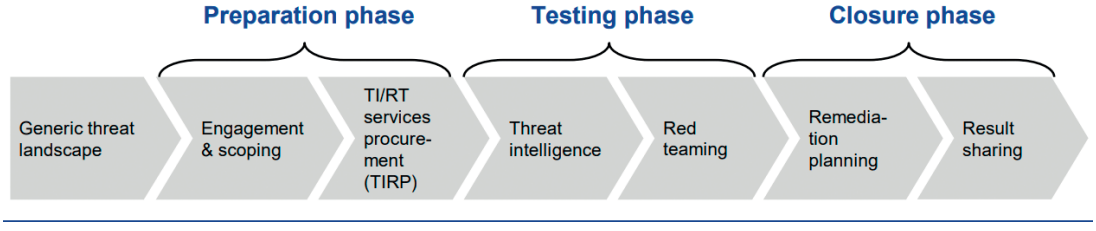

Figure 2: Example Tiber-EU Red Team exercise framework

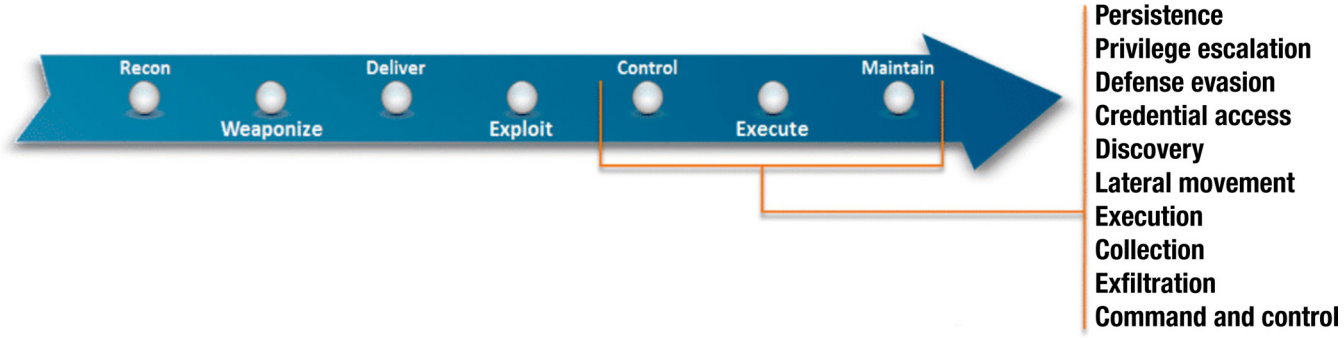

Figure 3: Example MITRE ATT&CK framework

Figure 3 is an example of the MITRE ATT&CK framework.

Attack frameworks are usually determined by the test provider unless required for compliance to a specific exercise framework. Similar to the exercise framework, knowing and understanding the specific attack framework to be used in the testing allows you to determine where to place a KPI or metric.

ESSENTIAL KEY PERFORMANCE INDICATORS

So, what is a KPI? A KPI is a quantifiable value used to track progress towards a key business goal. It is a measurement that provides a high-level perspective for making a strategic decision. So, KPIs provide direction towards achieving desired results and therefore can help a business make better-informed decisions.

Since the purpose of conducting Red Team testing is to assess your business's ability to withstand a real-world cyberattack, the goal of the KPIs you select should be to assess the business's performance in identifying, minimising and managing the attacks simulated in the exercises. The data collected from these KPIs should provide information on which decisions can be made to obtain that outcome. To meet that requirement, a good KPI has three criteria: it is clearly defined, relevant to the business, and able to clearly show performance in achieving the stated goal.

Having defined the objective of a KPI, Figure 4 shows the five essential ones to include in a Red Team exercise.

Mean time to detect

Mean time to detect or discover (MTTD) is a measure of how long a problem exists

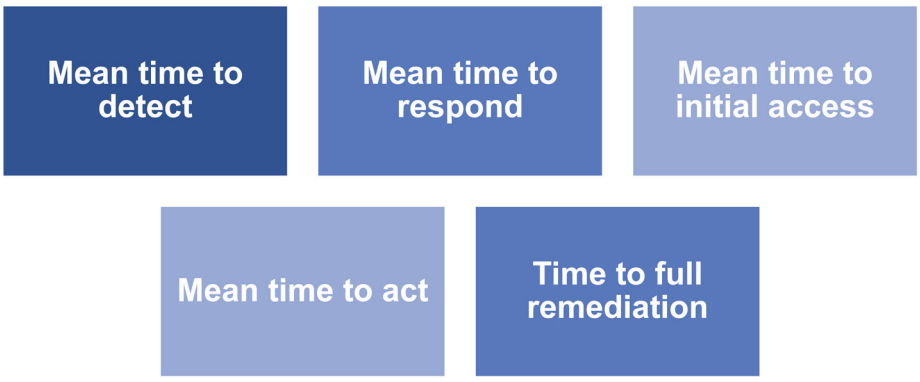

Figure 4: Five essential KPIs

before the business becomes aware of it. It is a critical Red Team KPI for obvious reasons — if the business is not aware of a security incident, it cannot respond and manage it.

For Red Team testing, the KPI should measure the duration of time between the start of the Red Team TTP activity and the identification of the activity by the business (Blue Team). To calculate the testing MTTD, add all the TTPs detection times associated with the test and divide it by the number of attacks. Tracking this KPI allows you to show (document) improvement in the business's incident detection capability.

Mean time to respond

Mean time to respond (MTTR) is a measure of how long a problem is detected before the business responds to assess it. This is an essential Red Team KPI as it is critical that a business responds to assessing a security incident before it escalates.

This KPI should measure the duration between the time the Red Team TTP activity is detected and the time it takes for the business (Blue Team) to respond and start their assessment of the activity. To calculate the testing MTTR, add all the TTPs response times associated with the test and divide it by the number of attacks. Tracking this KPI allows you to show (document) improvement in the business's incident response capability.

Mean time to initial access

Mean time to initial access (MTTIA) is a measure of the length of time it takes the Red Team to gain initial access to your systems after initiating the attack. This is an essential Red Team KPI as it indicates the strength of your perimeter security controls and your end user's vulnerability to social engineer attacks such as phishing.

This KPI should measure the duration between the time the Red Team TTP activity is initiated and the time it takes for them to obtain unauthorised access to your systems. To calculate the testing MTTIA add the duration of start and stop times for each of the attacks implemented by the Red Team intended to obtain unauthorised access — and divide the sum by the number of attacks. Tracking this KPI allows you to show (document) improvement in perimeter and phishing controls.

Mean time to act

Mean time to act (MTTA) is a measure of how long it takes after a problem is responded to before the business takes action to address it. This is another essential Red Team KPI as it is critical that a business act as soon as possible to address a security incident before it escalates or disrupts operations.

This KPI should measure the duration between the time the business (Blue Team) responded to a Red Team TTP activity and the time a solution to address the TTP is implemented. To calculate the testing MTTA, add all the TTPs reaction times associated with the test and divide it by the number of attacks. Tracking this KPI allows you to show (document) improvement in the defence posture of the business holistically.

Time to full remediation

Time to full remediation (TTFR) is the measure of how long it takes after a vulnerability identified and exploited by the Red Team is either accepted or remediated and documented on the business risk treatment plan. TTFR then is the duration of time between the end of the testing and completion of testing of the remedial control implemented to address the vulnerability identified in testing.

As you can see, these five simple KPIs produce critical information for both benchmarking the effectiveness of your current cyber security capabilities and identifying specific actions required to strengthen them.

ESSENTIAL METRICS

While they are often confused, a metric is not a KPI. A metric is a quantifiable measure used to track progress and evaluate success of a specific action. Metrics measure a specific activity or a process and provide granular information associated with that activity which can be used for tactical decision making. A good metric should provide information that is comparable, understandable, illustrates a ratio, and most importantly, facilitates behavioural change.

Where are metrics used? That depends on what you want to measure. The success of a specific action undertaken by the Red Team? Specific components of the Blue Team's defensive actions? Specific areas where a Purple Team can improve its capabilities?

Red Team testing metrics are often used to track the progress and performance of the TTPs associated with the attacks, which can be used to identify strengths or weaknesses. The following are examples of these types of metric:

- *TTP maps*: Attack TTPs can be mapped to the testing framework task used (ie Mitre ID). This allows a heat map to be created which identifies the Blue Team's strengths and weaknesses in detection capabilities for each attack.
- *Kill chain step maps*: Similar to the above, this also measures the detection and response capabilities strengths and weaknesses of the Blue Team.
- *Hostname and host types*: Documenting these allows the teams to see which operating systems have which strengths and weaknesses.
- *Type of detection expected/missing*: Allows the team to map which controls are performing well and/or which are underperforming.
- *Action timestamp*: Documents when TTPs are initiated by the Red Team. The timestamp should be synchronised with the Blue Team (and/or SOC's security information and event management [SIEM]).
- *Detection timestamp*: Documents when TTPs are detected by the Blue Team. The timestamp should be synchronised with the Blue Team (and/or SOC's SIEM).
- *Response timestamp*: Documents when TTPs are detected by the Blue Team. The timestamp should be synchronised with the Blue Team (and/or SOC's SIEM).
- *Business unit reaction*: Identifying reaction times from specific parts of the business (IT versus marketing) may provide insight into which departments need more collaboration with security.

There are many other metrics you can capture during an exercise to provide actionable data necessary to facilitate change, such as the following examples:

- Number and type of vulnerabilities found.
- Number and type of vulnerabilities exploited.
- Number of unidentified devices found.
- Number of unknown connections found.
- Security patches missing.
- Percentage of successful phishing attacks.
- Percentage of successful vishing attacks.
- Percentage of successful tailgating attacks.
- Number of incidents identified by staff.
- Number of incidents reported by staff.

So which metrics are 'essential'? Your essential metrics should be determined by the current maturity level of your business's information security management system (ISMS) (see Figure 5).

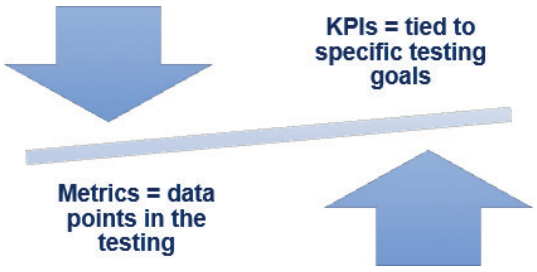

Figure 5: Essential metrics testing

Any metrics you select should be discussed and coordinated prior with your testing provider to ensure the activities are included in the test and the data points are captured, documented and presented to the business in a suitable format.

MATURITY MODELLING

What should be emphasised is that the KPIs and metrics you do select, while useful in providing current data for immediate decision making, are more valuable when used as a yardstick on which to mature your programme.

Most businesses conduct the same test over and over rather than adjusting the scope to benchmark it against conditions understood in the last test and 'raise the bar' — improving the outcome and continually maturing their security posture.

Collecting metrics and KPIs is an invaluable method of identifying and quantifying the changes and improvements made (or not made) over time. Information gleaned from your KPIs and metrics is the key to maturity modelling. Yet many organisations fail to retain and use it for this very purpose. Do not make the same mistake.

CONCLUSION: THE WHOLE DOUGHNUT

I remember once hearing a Red Team testing leader say, 'I don't look at the hole in the doughnut, I look at the whole doughnut'. Simply stated, a good Red Team tester stands back and takes a holistic view of the target to find and exploit overlooked vulnerabilities lost in plain sight. This is what makes the exercise so valuable — it tests the whole and not a single component (see Figure 6).

Your selection of KPIs and metrics, then, should be aligned accordingly — measurements of both the whole and individual components of the test. Together they measure the 'big picture' and that is the entire point of Red Team testing: it validates the maturity and effectiveness of the totality of the controls comprising your ISMS.

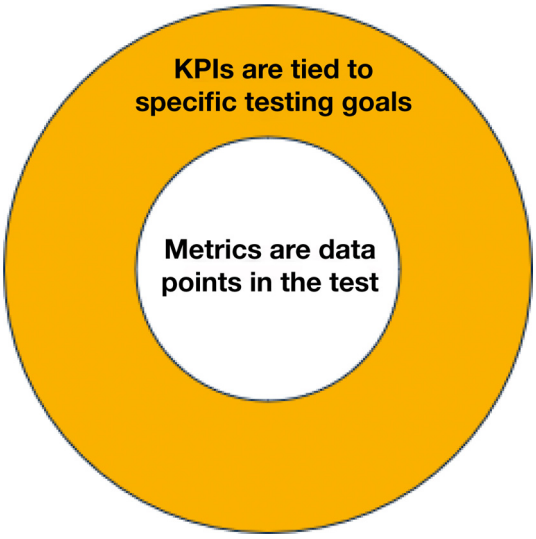

Figure 6: Testing doughnut

Before undertaking any Red Team exercise, you should decide (and document) what to measure and why. These measurements should be simple to define, easy to understand and communicate, and actionable and goal-oriented. They should produce data relevant to the business's stated cyber security objectives and be used for both long and short-term decision making. Most importantly, they should provide a foundation of information for the continuous improvement of your programme. It doesn't get any better than that.

References

1. HM Treasury (February 2023), 'G7 Fundamental Elements for Threat-Led Penetration Testing', Policy Paper, Gov.UK, available at https://www.gov.uk/government/publications/g7-fundamental-elements-for-threat-led-penetration-testing (accessed 24th January, 2024).
2. Bank of England, 'CBEST Threat Intelligence-Led Assessments', available at https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/cbest-threat-intelligence-led-assessments-implementation-guide (accessed 24th January, 2024).
3. European Union Central Bank, 'What is Tiber-EU?',

available at https://www.ecb.europa.eu/paym/cyber-resilience/tiber-eu/html/index.en.html (accessed 24th January, 2024).
4. Global Financial Markets Association (GFMA), 'GFMA Framework for the Regulatory Use of Penetration Testing in the Financial Services Industry', available at https://www.gfma.org/correspondence/gfma-framework-for-the-regulatory-use-of-penetration-testing-in-the-financial-services-industry/ (accessed 24th January, 2024).
5. Saudi Arabian Monetary Authority (SAMA) (May 2019), 'Financial Entities Ethical Red-Teaming', Webflow, available at https://www.sama.gov.sa/en-US/Laws/BankingRules/Financial%20Entities%20Ethical%20Red%20Teaming%20Framework.pdf (accessed 24th January, 2024).
6. The Association of Banks in Singapore (ABS) (November 2018), 'Red Team: Adversarial Attack Simulation Exercises', available at https://abs.org.sg/docs/library/abs-red-team-adversarial-attack-simulation-exercises-guidelines-v1-06766a69f299c69658b7dff00006ed795.pdf (accessed 24th January, 2024).
7. Hong Kong Monetary Authority (HKMA) (November 2020), 'Intelligence-led Cyber Attack Simulation Testing (iCAST)', available at https://www.hkma.gov.hk/eng/news-and-media/press-releases/2020/11/20201103-4/ (accessed 24th January, 2024).
8. MITRE ATT&CK, available at https://attack.mitre.org/ (accessed 24th January, 2024).
9. Lockheed Martin, 'Cyber Kill Chain', available at https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html (accessed 24th January, 2024).
10. Unified Kill Chain, 'The Unified Kill Chain', available at https://www.unifiedkillchain.com/ (accessed 24th January, 2024).