Purple Team Playbook:
Threat Modeling for Security Testing

Authors
Felisha Mouchous, MSc (Royal Holloway, 2019)
Daniele Sgandurra, ISG, Royal Holloway

Abstract

The reality with information security is that we cannot completely mitigate the threat of an attacker getting into our networks. Organisations can, however, *control* how they prepare and react to attackers by understanding how they operate. Threat modeling and security testing provide a way to first identify the threats and then simulate how an attack can take hold. In order to fully understand the threats, analysts strive to have a full picture of attackers' capability so that they can be as proactive as possible. To do this, the *red team* (attackers) and the *blue team* (defenders) in an organisation can work together to simulate attacks and test their defences.

In this article we propose a purple team testing framework (*Purple Team Playbook*) to allow organisations to leverage existing data on threats, attack techniques, defences and assets. This would help analysts in organisations to identify where they have gaps in their defences and understand how they can simulate threat actor behaviour. The Purple Team Playbook encourages analysts to understand and create new ways to test an organisation and, in the long run, allows organisations to save costs. [a]

[a] This article is published online by Computer Weekly as part of the 2020 Royal Holloway information security thesis series `https://www.computerweekly.com/ehandbook/Royal-Holloway-Purple-team-playbook-threat-modeling-for-security-testing`. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full thesis is published on the ISG's website at `https://www.royalholloway.ac.uk/research-and-teaching/departments-and-schools/information-security/research/explore-our-research/isg-technical-reports/`.

Larger Attack Surface: Fighting a Losing Battle?

As technology evolves, the security threat landscape grows. Just over twenty years ago the Internet was not available to everyone, so the landscape was much smaller, and we were less connected. In 2019 researchers from Kaspersky uncovered a new attack framework, called 'TajMahal', and they reported that this attack had remained undetected for five years. Due to the complexity of TajMahal, the researchers believe that these attackers could have also compromised other organisations, but this is yet to be determined. This reinforces the fact that organisations need to be evolving their testing and defence capabilities to try and keep up with sophisticated attackers (see Figure 1 for an illustration of growing threat landscape). If they do not, they could be under attack and not know it.

It might seem that organisations are fighting a losing battle when it comes to security: it is not a case of *if* they will come under attack, it is a case of *when*. This can lead to a greater financial toll on an organisation as they struggle to keep up. It can also be difficult to justify the time and monetary resource to be put into security as we only know how well our defences work once someone stress tests them. Unlike physical threats, such as damage from a fire where insurances typically cover the costs, it is much harder and can be more complex to recover from a cyber incident. With this in mind, we need to find a way to understand what our threats are and test how well we can stand up against attackers.

Figure 1: Organisations can draw parallels with the rise of smart devices in the home to see how rapidly the threat landscape is growing. Seemly innocuous appliances, such as refrigerators, can now provide a doorway into a home network if not protected.

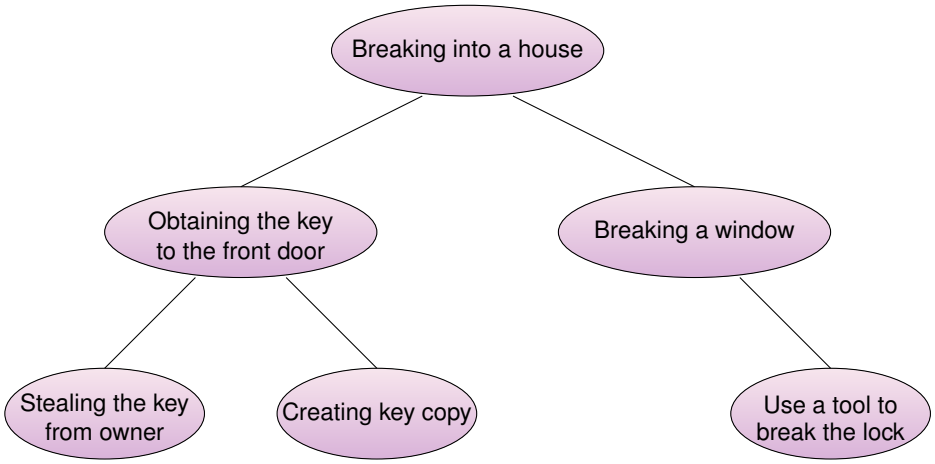

Figure 2: Attack tree example.

Know Your Enemy?

Threat modeling is used to identify, represent and prioritise threats in a uniform way. Threat modeling allows organisations to understand what threats they need to protect themselves against. As it is impossible to protect ourselves from every security threat, it is best to identify what feasible threats we are likely to face. One of the simplest ways to threat model is to use a method called **attack trees**. This is like creating a mind map for ideas, in this context used to understand the ways we might come under attack. An example could be to map out all the ways we can break into a house (Figure 2). In an attack tree, the nodes below the root tree (*breaking into a house*) would map out all the possibilities an attacker could use to achieve this goal, where one could be that someone steals the keys. However, if a similar threat modeling technique is used in an organisation, there is a risk that the tree will get very complex and lose its meaning, and so it may not be an effective model to use practically.

Another example of a threat model is the Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service and Elevation of Privilege (**STRIDE**) model. The way the model works is that we will have a system and use each word in the STRIDE model to determine if the system is affected by any of these issues. For example, let us consider the previous example of breaking into a house and “spoofing” as a threat. In this scenario, if we had a card entry door system and someone were to clone

or steal a valid entry card, then this would be a valid threat for a company and would need to be taken into consideration when the product is being designed and created.

> **Hard to truly know your enemy**
> A difficulty with threat modeling is that it's hard to think like an attacker: we can ask an analyst to map out all the ways they might rob a bank, for example. However, depending on the employee's experience, they may give you different ways that are helpful or could be highly theoretical and not feasible. Furthermore, more than likely, they are not real criminals so they can only base their assumptions on their experience so your model could be lacking in real threats.

Overall, threat modeling is an important activity during the development process as well as for established systems to understand what controls should be in place. It is clear that time and expertise needs to be invested for threat model to be really useful and valid for an organisation. In addition, as models can be highly subjective, we need to go a step further and play out these threats in the form of security testing to understand the practicalities of what we identify.

Impersonate Your Enemy?

We can now use the models of the attackers gathered in the threat modeling phase to test the security our systems in practice. **Security testing** mainly deals with secure code reviews, vulnerability assessments, penetration tests and red team tests. In particular, a **penetration test** aims to simulate an attack on an application in a controlled environment so that its security posture can be improved - this is not to be confused with *vulnerability assessments* that only go as far as identifying a vulnerability and not exploiting them.

Red team testing goes a step further by looking at the organisation as a whole and its security posture rather than a single system. Red team testing is objective based, for instance, red team members will use any means necessary to achieve their objective that would not normally be permitted during a penetration test, such as physical intrusion and social engineering.

The **Blue team** which is composed of analysts from the Security Operations Centres (SOC) will not know that this test is taking place. One of the goals of this type of testing is to see how the blue team responds to the red teams' attacks: this effectively tests an organisations incident response capability and identifies areas they need to improve on.

Purple team testing involves the red team conducting a test like they usually would, but instead of the blue team being unaware of this, they work together to enhance the results of the test. This effectively allows the blue team to create defences based on the red teams simulated attacks, and the red team to improve their attacks based on the blue team counter-defence. Organisations will need to have a strong understanding of what they want out of this type of testing and be clear about what their potential threats are, so that they are simulating the correct attacks.

There are many challenges that security testing can pose, such as a reliance on third parties that may lead to internal analysts being overwhelmed by the findings and not remediating appropriately. Another challenge is that testing is not continuous and only gives us results at a snapshot in time. Finally, there is also a significant time and cost overhead to conducting each test.

An Approach to Purple Team Testing: Purple Team Playbook

In the event of emergencies, we have processes and plans in place for what to do - a playbook of sorts. Just like having these incident response playbooks and disaster recovery plans, *we advocate organisations to have a playbook to document their enemies' behaviours and discover how resilient they are to them*.

Our framework, the **Purple Team Playbook** (PTP), addresses threat modeling from the perspective of the internal red and blue teams in an organisation to be used in the security testing process. The purpose of conducting various types of security assessments is to ensure that we can protect and detect if an attacker has breached an organisation. Many sophisticated attackers such as **nation state** actors rely on stealth and are harder to detect, so security testing needs to evolve in order to match the pace they are working at. By centralising the knowledge in the PTP and getting internal red team testers and blue team testers working together, we can better understand where an organisation is currently and identify where the gaps are in testing. This framework also allows an organisation to fully understand what data they need to hold in order to threat model effectively. They could of course purchase commercially-supported tools that may do the same thing; however, they still need to design and decide what data they want to leverage and what they want out of the tool. Instead, PTP shows organisations how they can use their own data to understand what threat actors are targeting them and how they can use this to security test their systems.

Nation state – sophisticated threat actors that are government funded; they will be as covert as possible as they are motivated politically

What Do We Know Already?

To construct the PTP, we need to first gather all the necessary information and then model it in a way that makes sense to our organisation. This is a good exercise to understand the volume of data we actually have and build the foundation of the playbook. This part of the framework can be automated using a **data ingestion engine**, for instance, pulling data from an API (Application Programming Interface) and dropping it into the database using Powershell or an ETL (Extract Transform and Load) application. Once the data is ingested, we split the data into four datasets, namely red, blue, organisation data and a combined purple team dataset. The user can use data visualisation software and connect it to the Purple Teams Playbook Knowledge base database. This gives the user freedom to search and show the data in any way they would like and create dashboards to model the data.

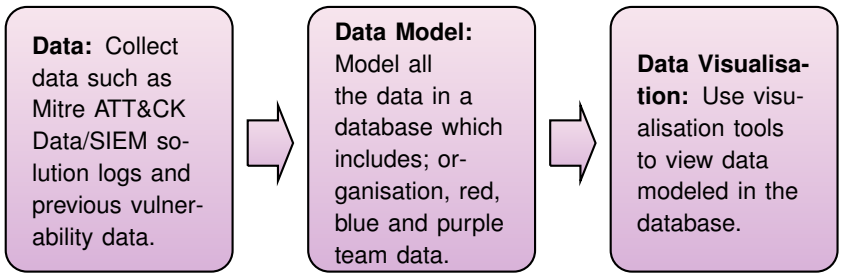

There are some things to consider when trying to build this framework: firstly, the organisation needs to be on board with the framework at an executive level to ensure there is funding and time allocated to setting up the framework. Secondly, there will need to be some form of maintenance to ensure that the framework is updated, and that access is only granted to those who need to use it.

Finding the Unknown Unknowns?

To use the framework to find the unknown unknowns, we conduct threat modeling from the red and blue teams' perspective. The data modeled in the data model part of the framework is used to determine where the gaps are in testing. As threat actor capabilities are always evolving, it is important to understand what has been tested and what monitoring and mitigations are currently in place and where improvements can be made.

Usually analysts are dependent on their own experience and on talking to the owners of the system to get the information required. By using the PTP, we have a centralised place to look at all the relevant data, for instance, data from previous tests to understand whether the vulnerabilities have been fixed. Also, by mapping tests against known threats, such as using the **Mitre ATT&CK** framework, which is a well-known framework that contains details of tactics and techniques used by sophisticated attackers, we can have visibility on the security controls that the blue team have in place for each of the techniques, which helps us understand our coverage. For a traditional penetration test, we would not normally look at blue team data as there are time constraints for each test so the applications security should be prioritised. Overall, the PTP framework allows analysts to tailor tests to specific gaps in the organisation and help uncover previously unknown threats.

Key Purple Teaming Playbook Takeaways

Finding gaps in defences: By using all the data an organisation has available, it is easier to find patterns and correlations that could be leading to security gaps. This information could be harder to get normally as information can be populated in different places and hard to join together.

Easy mapping against known threats and new threats: As the necessary threat data is all in one place, it can be easier to determine how well an organisation is doing against known threats, for instance mapping their defences against the Mitre ATT&CK framework.

Encouraging innovation: As analysts can understand easily where the gaps are and the defences are, they can get more creative like a normal attacker would to find new ways to simulate attacks on an organisation.

Cost saving: Once all the relevant data is in place, it is much easier to see what controls are and are not working. This could lead to cost savings as some controls that are costly could be replaced or removed if they are not working as they should. Also, analysts can target testing more effectively, so less money needs to be spent on identifying all the gaps.

Evolving to Fight Another Day

Trying to stay secure in this ever-changing world is extremely difficult – there is no so-called 'silver bullet' to security. The only way to try and keep up with all the threats we face as technology changes, is to try and evolve our capabilities to defend against them. Purple teaming is just one way that we can try to evolve as we have a lot of data in organisations that can be weaponised to mount our defence.

> The only truly secure system is one that is powered off, cast in a block of concrete and sealed in a lead-lined room with armed guards - and even then, I have my doubts. – **Gene Spafford**

Biographies

Felisha Mouchous completed her MSc in Information Security at Royal Holloway, University of London, graduating in 2019 with Distinction. She currently works as a Security Analyst for the Bank of England in their Technical Vulnerability Management team. She is responsible for scoping, co-ordinating and conducting intelligence led penetration tests for critical systems at the bank. In 2019 she was short-listed in the newcomer category as one of three finalists for the Women in Cybersecurity Awards.

Daniele Sgandurra is a Senior Lecturer in Information Security at Royal Holloway in the Information Security Group (ISG), where he leads the Systems and Software Security Research Lab (S3Lab). His current research interests lie within systems and software security, focusing on malware analysis, sandbox evasion and virtualization security.

Series editor: S.- L. Ng