Red Team: Adversarial Attack Simulation Exercises 2.0

for the Financial Industry in Singapore

The Association of Banks in Singapore

September 2024

Table of Contents

1 Executive Summary

Cyber security attacks against organisations such as financial institutions (FIs) are evolving rapidly in scope, complexity, and sophistication. To address this risk, FIs deploy layers of defensive measures, solutions, and controls to reduce their exposure to attacks and improve their response readiness. Offensive simulation exercises complement the defensive layers to assess the effectiveness of defences and improve the security team's preparedness to detect and respond to incidents.

Adversarial Attack Simulation Exercises (AASE), often referred to as Red Team (RT) exercises, are sanctioned, planned, risk-managed and objective-driven cyber security assessments that simulate highly sophisticated targeted attacks against an organisation.

The objectives of AASE are to assess and enhance the resilience of FIs against sophisticated attacks. To efficiently allocate their resources to the unique threats they are facing, FIs are encouraged to create scenarios for their attack simulation by identifying the most likely adversaries and the attack vectors through attack modelling. The goal of these exercises is to assess the capability of an FI to prevent, detect and respond to cyber- attacks that may impact Critical Functions or business continuity. To achieve this, these exercises simulate a full end-to-end cycle of a cyber security attack, replicating actions and procedures utilised by real world adversaries with a high level of intent, sophistication, and capability.

This document provides guidelines and best practices to help organisations in the financial industry in Singapore plan, execute and report of such exercises.

2 Introduction

AASEs are designed to challenge an FI's cyber security defences by modelling and then executing attacks based on real-world adversaries' Tactics, Techniques and Procedures (TTP). Scenarios are designed to be as realistic as possible, and may target the FI's People, Processes and Technology with the intent to compromise organisation's Critical Functions (CF). The primary goal of the exercise is to assess the organisation's ability to prevent, detect and respond to cyber-attacks and discover potential weaknesses that may not be identified through standard vulnerability and penetration testing methodologies.

This guideline is intended to support FIs within Singapore but may also be used by other organisations to plan and conduct such exercises. The document provides guidance on best practices and recommendations on how to adopt them partially or in full and depending on the maturity of the FI's capability in conducting these exercises.

This document aids in planning and executing such exercises but should not be relied on solely to achieve compliance with regulations.

It is expected that the objectives, test scenarios and the report structure will be tailored according to the FI's scale, operations, external threat landscape and risk appetite.

3 Exercise Benefits

An adversarial attack simulation methodology provides a more authentic and holistic view of an FI's resilience.

By simulating realistic attacks during the exercise and taking into consideration the relevant threat landscape and potential adversaries, the following benefits can be achieved:

- An assessment of the organisational resilience against adversarial attack TTPs.
- Identification of weaknesses in security controls and associated risks not detected by standard vulnerability and security testing methodologies.
- An assessment of the FI's security incident management and/or crisis management response and processes.
- A safe, controlled opportunity to identify and enhance the security posture of an FI, reducing risk of cyber compromise.
- An opportunity for the defensive teams, such as the security monitoring or incident response team to gain experience and be more proficient in detecting and responding to incidents.

4 Definitions

The section below outlines key definitions and terms used when conducting adversarial attack simulations.

Term	Definition
Attack Modelling	Attack Modelling is the process where the outputs of the Threat Intelligence are used to generate likely attack scenarios that can accurately simulate real-world cyber security threats under operational conditions.
Attacker (sometimes referred to as Red Team)	An Attacker is an individual or a team who is employed or contracted by an organisation to simulate the attack TTP of a real-world adversary based on intelligence about prevailing and/or probable cyber threats and incidents to stress and provide guidance with regards to enhancing organisational resilience, utilising goals set in the scope of the exercise. For an effective exercise, ideally, the skills and capability of the Attackers should be matched to those expected of real-world adversaries as closely as possible. The provider selection section provides some guidance on selecting a provider with adequate capabilities.
Automated Attack Simulation	Automated attack simulation refers to the use of automation – typically Breach and Attack Simulation (BAS) tools – to remove the manual effort in running and re-running the simulated attacks.
Concessions	Concessions are a method of deliberately altering the course of the exercise by providing explicit, agreed assistance to the attacking team to achieve the goals of the simulation.
Continuous Adversarial Attack Simulation	A continuous adversarial attack simulation is a form of AASE where the Attackers are given a timeline for the attack execution which may stretch for an extended period. This extended timeline allows attackers to pivot to more opportunistic attacks as the threat landscape changes, adjusting their attacks based on the updated threat intelligence which reflects shifting motives or tooling of real-world adversaries.

Critical Functions (CF)	Critical Functions are business functions or services that if compromised would significantly impact business continuity, affect the reputation of the FI, or cause financial loss. These Critical Functions generally represent the greatest opportunity for real-world adversaries who are motivated by financial gain, information, or intellectual property theft, and/or a desire to inflict business disruption.
Cyber Range	NIST defines Cyber ranges as interactive, simulated representation of an organization's local network, system, tools, and applications that are connected to a simulated internet level environment. Ranges can also be virtual environments designed with pre-defined objectives for cybersecurity training, both for defensive and offensive teams. These ranges provide a safe, legal environment to gain hands-on cyber skills and a secure environment for product development and security posture testing.
Defender (often referred to as Blue Team)	A Defender is an individual or a team who is employed or contracted by an organisation to detect and/or prevent a cyber-attack and respond to one when it happens. This virtual team would typically include all resources in the FI's Security Operations Centre, incident response teams, and other technology infrastructure support functions.
Exercise Director	The Exercise Director is employed or contracted by an FI to oversee the development, execution, review and/or approval of the exercise. The role has primary accountability in managing the delivery of the exercise, and operational risks arising from an adversarial attack simulation exercise being conducted on the production environment. The Director should be able to understand the technical details associated with the attacks and its potential risks and consequences. More broadly, the Director is expected to have intimate knowledge of the FI's infrastructure and applications.

Exercise Working Group (Attackers + Exercise Director)	The Exercise Working Group is a team that comprises the Attackers, and the Exercise Director. As the Exercise Working Group will be heavily involved throughout the planning and execution phase, they must ensure that all information regarding the AASE remains strictly confidential to avoid "tipping off" the Defenders. The Exercise Working Group may, however, ensure that senior stakeholders are aware of the exercise and are positioned to provide required authorisation/approval for the attack teams' activities during the exercise.
Letter of Engagement	A Letter of Engagement is to document the exercise that is commissioned, sanctioned, and authorised by the FI. It should be signed by the sponsoring executive or senior management, and it should briefly describe the intent of the test, the extent of involvement of the testers, and instructions on how to authenticate them for their authorisation during a potential investigation. This can be useful in situations where the testers need to prove their intent if they are apprehended by security staff.
Organisational Escalation Paths	The Organisational Escalation Paths represent the different chains of control for any operational issues encountered, as part of the FI's standard business practices, where the FI's staff would escalate problems up the chain of management to the organisation's senior management. An example of this would be suspicious activity being detected by Defenders and escalated up the chain of command as part of incident response.
Target Organisation (also referred to as Organisation, or FI)	Target Organisation or "organisation" or "FI", in this document, refers to the target of the AASE and the potential target of real-world attackers.
Threat Intelligence	Threat Intelligence is the product of the process whereby analysis is performed against real-world threats that would potentially apply as a threat to the FI. Threat Intelligence is supported by a variety of sources including internal intelligence gathering, public and proprietary information feeds, and security intelligence sharing platforms.

Trusted Agents	Trusted agents are employees of the FI in key positions that can assist the execution of the exercise by providing support to the simulated Attacker/Red Team to enable the exercise on request from the Exercise Director. This support is provided by way of pre-approved and pre-defined concessions such as approval to execute a specific command or to produce and release other supporting documentation/resources. Trusted agents may be personnel from the business functions or IT support, and while they are aware of their obligations to assist with an exercise, they are not aware of the details of the exercise, overall objectives, or the progress of an active simulation. Given the potential added privileges granted to the Exercise Director, all actions must be thoroughly documented, and pre-approval should be sought from senior management during the preparation phase.
Tactics, Techniques and Procedures (TTPs)	Tactics, techniques, and procedures are used to describe the behaviour of an actor. A tactic is the highest-level description of this behaviour; techniques provide a more detailed description of the behaviour in the context of a tactic; and procedures provide a lower-level, highly detailed description of the behaviour in the context of a technique.

5 Getting Started with Adversarial Attack Simulation Exercise

Adversarial Attack Simulation Exercise serve to complement other forms of security testing (e.g. code review, vulnerability assessment, penetration testing) and should be incorporated into the security testing program of an FI. The FI may conduct AASE in parallel with other additional exercises to assess its organisational resilience against cyber threats.

It is important to be aware of differences between AASE and other types of security testing. FIs should take note of their current cyber resilience maturity when planning exercises. Furthermore, it is important to note that operational and legal considerations may place some limitations on the ability for an exercise to fully recreate a real-world attack.

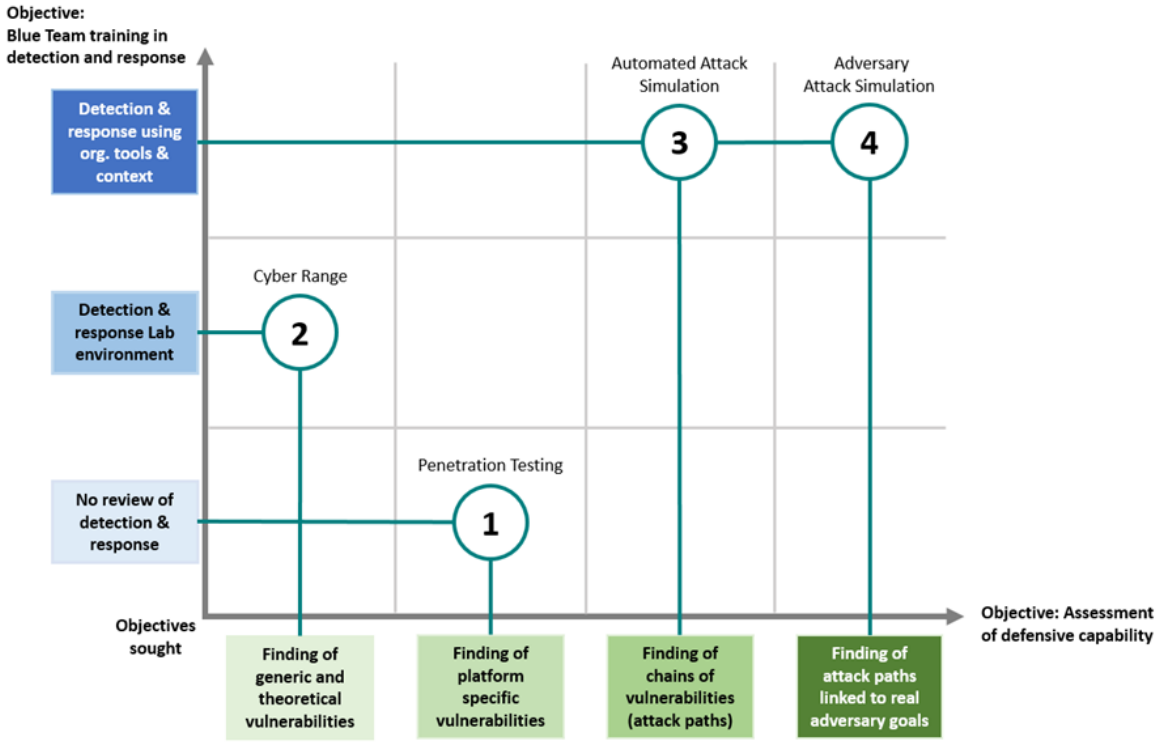

5.1 Adversarial Attack Simulation Exercises Value Realization Approach

The guidance set out in this section aims to provide a list of areas for FIs of different maturity levels to consider when charting their roadmap and gain the utmost value from an AASE.

The guidance should not be inferred as a mandatory pre-requisite or a replacement for AASE, or any other purpose than the intended aim stated above.

The diagram below provides an overview of the suggested standalone activities for FIs to consider prior to conducting an AASE and gain more value out of the exercise. Each standalone activity is categorized into recommended and additional activities to cater to different FI maturity to chart their roadmap.

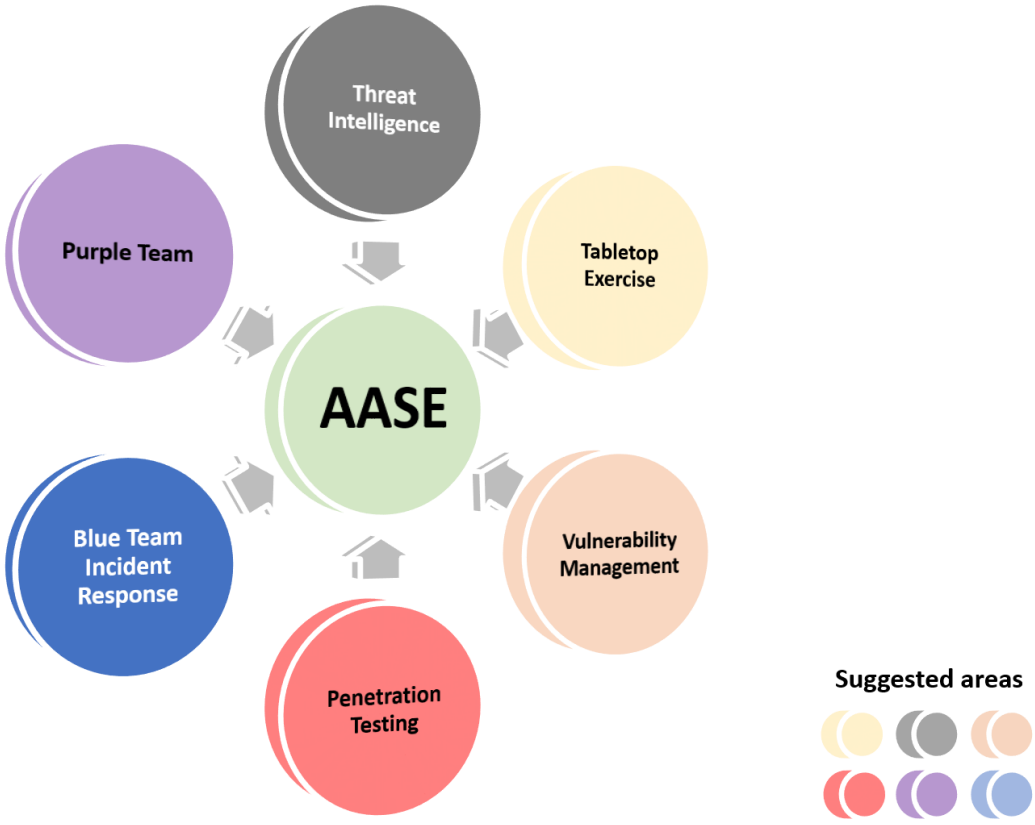

5.1.1 Threat Intelligence

Threat Intelligence involves the gathering, analysis and dissemination of information related to potential cybersecurity threats. For FIs without an internal threat intelligence function, it is

recommended to start with Threat Intelligence from external providers to obtain Generic Threat Intelligence and - if possible - Targeted Threat Intelligence in preparation for AASE.

Generic Threat Intelligence refers to threat that are not specific to any FIs but are threats generic to an industry or region. Targeted Threat Intelligence refers to threats that are specific to the FI's environment stemming from specific Threat Actors.

Recommended Baseline Activities

- Generic Threat Intelligence from Public Sources
- Generic Threat Intelligence from External Provider
- Targeted Threat Intelligence from External Provider

Additional Activities

- Targeted Threat Intelligence from FI's internal team (combining internal knowledge with TI from External Provider)

Section 7.2.2.1 and 7.2.2.2 elaborates further on Generic and Targeted Threat Intelligence for the AASE Attack Preparation phase.

Outcome

The trends and TTPs extracted from a threat intelligence report would help to prioritize TTPs, scope an AASE, and build a realistic attack campaign for AASE. In the ideal state of activities, combining targeted threat intelligence from FI's internal threat intelligence team knowledge with threat intelligence from external provider would provide an accurate and comprehensive threat intelligence coverage.

5.1.2 Cybersecurity Tabletop Exercise

A cybersecurity Tabletop Exercise (TTX) is a structured and interactive discussion-based activity designed to assess an FI process of detection, response, and recovery protocols. The stakeholders of the tabletop exercise should not be limited to security teams (such as SOC, Threat Hunting, and IR) but also include stakeholders (such as application, network, and infrastructure) whose systems are in scope. Depending on the type of TTX, legal, communications, business continuity or regulatory teams may also be included.

Broadly, there are two types of TTX. The first type is the technical team TTX where process and protocols are discussed at a low level, down to the technical level such as what tools are being used to respond or remediate a situation. The second type is the senior management level TTX where policies and processes at a high level are discussed. It is recommended to start with technical teams before moving on to include senior management team.

Recommended Baseline Activities

- Cybersecurity TTX with technical teams

Additional Activities

- Cybersecurity TTX with senior management

Outcome

TTX would result in an improved incident response plan, policies, and procedures. This, in turn, would enhance the FI's incident response preparedness.

5.1.3 Vulnerability Management

Vulnerability Management is a systematic and proactive approach to identify, evaluate, mitigate, and monitor security vulnerabilities in an organization's IT infrastructure, including hardware, software, networks, and systems.

Recommended Baseline Activities

- Asset Management
- Patch Management
- Periodic external perimeter vulnerability assessment

Additional Activities

- External Attack Surface Management
- Internal vulnerability assessment
- Authenticated vulnerability assessment

Outcome

Asset and patch management along with external perimeters scan are recommended programs to be established before carrying out an AASE. This would help to address basic gaps in configuration and patching which could be exploited by an attacker. Knowing the assets in an organization would also help in realistic scenario planning.

5.1.4 Penetration Testing

Penetration Testing (PT) is a proactive and controlled point-in-time security assessment used to evaluate the security posture of computer systems, networks, or applications. Unlike vulnerability assessments which primarily focus on identifying known vulnerabilities, penetration testing goes a step further by identifying and attempting to exploit unknown vulnerabilities. During exploitation attempts, preventive controls could be tested as part of the scope.

Recommended Baseline Activities

- External Infrastructure Penetration Testing
- Web and Mobile application testing (with and without credentials, and full access to documentation) of Internet Facing services

Additional Activities

- Security Controls Validation
- Internal Infrastructure and Applications Penetration Testing with full access to documentation

Outcome

Penetration Testing helps to identify more complex weaknesses and vulnerabilities in infrastructure and applications that attacker could take advantage of.

More information on conducting Penetration Testing can be found in the ABS Penetration Testing Guidelines for the Financial Industry in Singapore.

5.1.5 Blue Team

The defender or Blue Team is composed of resources from the FI's Security Operation Centre (SOC), Incident Response (IR) teams, and other supporting teams that provide defensive functions. The team is employed or contracted by the organization to detect and/or prevent a cyberattack and respond to one when it happens. As AASE assess the FI's security incident management and/or crisis management response and processes, the blue team would need to have an incident response playbook to respond to attacks and remediate incidents.

Recommended Baseline Activities

- Establish Security Operation Centre (Internal or External)
- Establish Incident Response Plan
- Establish Incident Response Team (Internal or External)

Additional Activities

- Establish Threat hunting, Incident Response and Digital Forensics Capabilities

Outcome

It is crucial to establish Blue Team capabilities to detect and respond to cyberattacks. The AASE is a very effective way of assessing and improving the capabilities of the blue team.

5.1.6 Purple Team

A Purple Team Exercise is a collaborative and structured cybersecurity assessment that brings together both the defensive ("blue team") and offensive ("red team") elements to work closely together. The purpose of a Purple Team is to improve the FI's overall cybersecurity posture through testing of detection capabilities. The Purple Team could consist of internal or external red and blue teams. The red team would simulate the attacker TTPs while the blue team would focus on detect and improving its capabilities in respond to the TTPs ran by the red team.

Outcome

The benefits of Purple Team include addressing of existing security gaps, providing the blue team an opportunity to better understand the attack lifecycle, learn from the attacks and foster collaboration between blue and red teams.

5.2 Differences with Penetration Testing

Some of the key differences between Penetration Testing and AASEs are:

Penetration Testing	Adversarial Attack Simulation Exercises
• Primary objective is to identify as many vulnerabilities as possible, in a limited scope	• Primary objective is to stress and enhance organisational ability to detect and respond to adversaries
• Limited scope, asset-based technical assessment	• Objective-based, open-scoped, designed to demonstrate critical impact to a business or organisation. Targets people, process, and technology
• Made known to all the stakeholders	• Covert. Only the Exercise Working Group is aware of the exercise
• Social engineering is not used	• Social engineering may be used
• Physical security is not tested	• Physical security may be tested
• Execution aligned to industry-recognised technical methodologies	• Execution aligned to mimicking Tactics, Techniques and Procedures of real-world adversaries

5.3 Limitations of Adversarial Attack Simulation Exercises

It is important to note some key differences between real-world attacks and Adversarial Attack Simulation Exercises:

Real-world Attacks	Adversarial Attack Simulation Exercises
• Usually not time-bound	• Usually time-bound due to resource limitation granted in a project (but could be open-ended)
• Not bound by Law or Ethics	• Bound by Law and Ethics
• Uncontrolled by the target organisation	• Controlled by the target organisation, risk controls can be applied
• May use physical/psychological violence and extensive coercion	• Respects the integrity and well-being of employees and partners

6 Guiding Principles

The success of these exercises is guided by the following principles:

6.1 Exercise Goals

Exercise simulations should be based on the likely goals of real-world sophisticated adversaries.

AASEs are goals driven. The intent of the exercise is not to identify and report vulnerabilities an FI may have. The goals represent what real-world adversaries may aim to achieve by compromising an FI.

For example, an FI could be targeted by an adversary for financial gain and the objective could be to demonstrate an ability to perform a large unauthorised funds transfer. The scenarios will therefore be developed to include all the steps to reach that goal.

It is not necessary to rotate or vary goals for every exercise as threat intelligence or other information may indicate that real-world adversaries could still have the same motivation.

Scenarios and goals need only be changed if Threat Intelligence indicates that the threat landscape has changed, or controls around Critical Functions vary.

6.2 Exercise Secrecy

The scope, nature and timing of the exercise should be kept secret to adequately assess the effectiveness of security defences and response to cyber-attack scenarios.

The information related to the exercise should only be circulated and discussed within the Exercise Working Group for the duration of the exercise. The Defenders are expected to react to the simulated attack in the same manner that they would do when they detect real attacks. Keeping the scope, nature and timing of the exercise secret would therefore provide a more accurate assessment of the capability of the Defenders to prevent, detect or respond to real-world or simulated cyber-attacks.

6.3 Targeting of Live Critical Functions

Exercise scenarios should be designed to target Critical Functions of the organisation and in a manner that is aligned with the motives of expected adversaries (in the production environment).

Real-world adversaries are motivated to compromise Critical Functions to achieve maximum

impact. Motives of adversaries can include business disruption, financial gain, and/or information theft. As such, the same functions should therefore be targeted by the Attackers in the exercise.

Given the purpose of the exercise is to stress and enhance the FI's current resilience to a real-world adversary, the exercise should be conducted against the live environment. Assessments in staged or otherwise non-live environments may influence the outcome of the exercise and would not be representative of organisational cyber resilience against real- world threat, unless the specific target is a development environment, or the team is testing for weaknesses in the protections between a test and production environment within the organization.

Assessment in a live environment would include, for example, the exploitation of staff cognitive biases, identification and exploitation of gaps within processes utilised by an organisation, and subversion of existing business process or technologies within an environment.

6.4 Exercise Frequency

Organization should employ a risk-based approach to determine the appropriate frequency for AASE.

Factors such as the prevailing threat landscape, recent changes to FI's cybersecurity operations as well as results and lessons learned from past AASE may be considered in arriving at the suitable cadence.

This guide recognizes that different FIs may be at different stages in their AASE journey and therefore recommends that AASE be performed once every 24 months minimally, but FIs have the flexibility to increase the frequency to every 12 months should this be deemed necessary.[1]

With each AASE iteration, the scenarios and sophistication of the TTPs may be adjusted in line with the improvements made in the organization's cybersecurity posture and changes in the threat landscape.

[1] Unlike a Tabletop exercise, an AASE deployment typically faces greater risks and uncertainties throughout its lifecycle and thus the success of an AASE is not always guaranteed.

A planned attack may be rendered invalid at the last minute due to unforeseen circumstances such as technical changes to the target infrastructure (e.g.: patched server vulnerabilities).

Additionally smaller FIs without in-house AASE capabilities may engage external service provider and this would incur additional time and effort in performing the onboarding due diligence such as background checks before the external capabilities can be operational.

With these considerations and in alignment to relevant industry practice it is determined that it is in the best interest of the industry that FIs have the flexibility in deciding their AASE frequency according to their risk assessments, budget and resources to allow focus on delivering quality AASE.

6.5 Exercise Duration

Exercise duration will be determined by the complexity of the attack scenarios, and the scale of the organization being examined.

The duration of the exercise is considered during the planning phases and is based upon the perceived persistence, and operating procedure, of the adversary that the FI is looking to simulate. Typically, a large financial institution may operate a cycle of between four to six months for complex scenarios (inclusive of Attack Modelling).

Some FIs may choose to test its defences through continuous adversarial attack simulations whereby the Attacker would be able to adjust their attack depending on the changes to the threat landscape. Exercises may also be open-ended (no defined timing but only closed upon achieving the goals).

At the end of each AASE, there should be a comprehensive evaluation from the Red Team's perspective of the FI's current prevention, detection and response capabilities, and a roadmap of on-going improvements over varying periods of time.

7 Methodology

Overview

The typical methodology of executing AASE involves five distinct phases - the planning phase, attack preparation, attack execution, exercise closure and post-exercises actions.

The AASE starts with a “planning phase” where the scope of the assessment is defined and described, service providers are sought, and a budget is set aside. An Exercise Working Group is formed, and communication protocols are defined.

The second phase, “attack preparation”, involves the creation of scenarios based on the defined Critical Functions and exercise threat model.

The third phase, “attack execution”, is usually comprised of several sub-phases which are covered extensively in section 7.3 ‘Attack Execution Phase’ below.

The fourth phase, “exercise closure”, describes reporting, cleaning-up, transferring knowledge, remediating, and communicating outcomes at the conclusion of the exercise.

The final phase, “post-exercise actions’, covers formalization of action plans on the AASE recommendations.

7.1 Planning Phase

Prior to the start of the AASE, careful planning is required to ensure that the exercise can be conducted effectively and with minimal risks to the FI.

It is recommended that the planning stage includes the following steps:

- **Determine the scope and duration of the exercise.** This will affect the number of goals that could potentially be achieved during exercise period. Given the focus of the exercise is to stress organisational resilience and controls, it is recommended that defined goals and objectives, where relevant, should cover both technical and business controls by the FI. In addition, Critical Functions, which are important to the organisation should also be covered and tested.
- **Define the key exercise parameters** including desired outcomes. This will affect the degree of secrecy required, the composition of the Exercise Director team, the escalation and cautions required, the data management, and lastly the selection of providers. Those parameters will impact the course and results of the exercise.
- **Determine the participation model.** For instance, whether the exercise can be performed by the FI's own experienced staff, by external service providers or a combination of both. The exercise requires Threat Intelligence analysts to perform threat intelligence gathering and qualified Attackers to execute the attack (although both roles can be performed by the same, appropriately qualified person). It is recommended that these capabilities are provided by impartial parties, ensuring the integrity and realism of the exercise.
  - The planning stage should include the selection and procurement of the external services, if applicable.
  - The FI should only engage service providers that are reputable and have experienced persons to perform such exercises. Recommendations in selecting the right provider are detailed in Section 7.1.4 "Provider Selection".
- **Creation of the Exercise Working Group** that will oversee and direct the entire exercise. Members of the group should maintain secrecy throughout the entire exercise and not disclose the scope and plans of the exercise to persons outside of the group. To avoid any security alerts generated by the exercise from being escalated beyond what is necessary for the exercise (i.e. to law enforcement or government authorities), the Exercise Working Group should consist of members who are part of the Organisational Escalation Path, who would be informed of such escalations and are able to intervene where needed.
- **Determine the escalation process** and key persons for the escalation of any emergency issues encountered or caused during the exercise. If the escalation process and the key persons cannot be determined at the planning stage, it is acceptable to do so at the later stage after the Critical Functions that are in scope have been determined.

The decisions made during the planning phase are usually captured in the Exercise Preparation Report.

7.1.1 Goals and Scope

Scoping of the exercise should be performed by the Exercise Working Group, which typically involves both business and IT leaders, and should be led by the project sponsor, which will typically be the FI's CISO or its delegate. Scoping should include identification of Critical Functions within the FI that are likely to be targeted by real-world adversaries.

7.1.2 Defining Exercise Parameters

Defining parameters must be carefully agreed during the planning of an adversarial attack simulation exercise, as they will directly influence the realism of the simulated attack, and therefore the benefit of the overall exercise.

Determine the Degree of Secrecy of the Exercise

The first defining parameter is the degree of secrecy desired prior to and during the exercise. Higher levels of secrecy allow the exercise to demonstrate and validate the FI's actual and unfettered response to a cyber-attack, but such an approach can introduce additional operational risk e.g., high risk recovery actions or responses initiated in the production environment could result in prolonged outages to services and business disruption. Lower degree of secrecy tends to diminish the execution of detection and response capabilities and processes. This will influence or otherwise undermine exercise outcomes. FIs should determine the ideal degree of secrecy to allow adequate testing of their ability to detect and respond to threats, while minimizing risks of disruption to the FI's business. FIs should also plan recurrent exercises with higher levels of secrecy in subsequent iterations.

Select the Exercise Director and Exercise Working Group

Typically, the Exercise Working Group supplements and assists the Director in conducting the exercise, by observing, opining, and intervening where necessary. Regular meetings should be established within the Exercise Working Group to oversee the progress of the exercise.

The Exercise Working Group would usually consist of staff in key functions who are able to influence incident responses in cases where there is a need to alter the outcomes of the exercise, and other processes. The participants may include members of:

- Senior Management
- Risk Units (Business Risk, Technology Risk, Operational Risk, Cybersecurity Risk)
- Business Continuity and Crisis Management

- Incident Management
- Information Security
- Relevant Business and System owners
- Legal and Compliance

The Exercise Director should be able to decide, at every stage of the exercise, whether to proceed with an action, pause, or terminate the exercise. The Exercise Director should be qualified and sufficiently aware of both technology and business to make such decisions, and these decisions can be supported by the Exercise Working Group.

Define Exercise Data Handling Protocols

Data handling and management protocols should be established prior to the exercise commencement for the Exercise Working Group. These protocols must define how data exposed or acquired during the exercise should be handled (e.g. usage, data protection and retention).

Given the nature of the simulation, the Attacker may potentially gain access to the FI systems and data.

The protocols should limit access to this data by the Exercise Working Group, as its members may not be authorised to handle this data by their individual scope of roles. Access to acquired data should be managed, and the Attacker, should only present to the Exercise Working Group a subset of data sufficient to validate achieved goals. This subset can be produced from smaller samples of acquired data that has been redacted to remove sensitivity. Ideally, all data verification should be performed by the business owner of that data set, and whose scope of role includes authorisation to access the data.

There should be an agreed protocol to adequately protect the data which the Attackers may have access to. The protocol should guarantee at minimum the same level of protection adopted by the FI, ensuring data is only available to the members of the Attacker group that are directly involved with the simulation, and that data is destroyed after the exercise. There should also be an agreed process to dispose data used during the AASE in a secure manner.

Assess the Impact to Production Systems and Operations

While creating the scope of the exercise, impact to the environment should be carefully considered with a thorough risk assessment. If the potential operational impact to specific components in the live environment during an exercise is deemed to be unacceptable, the FI may consider performing parts of the exercise in a non-production environment that is a close replica of the actual instance. However, the FI should be aware that performing the exercise in a non-production environment may limit the realism of the exercise, its results and reduce the overall

benefits of the exercise. Hence, conducting the exercise in the non-production environment should be an exception rather than a norm as the results may not reflect the true security state of the FI.

Any parts of the exercise that are not performed on the live environment should be clearly documented with reasons in the report. For example, certain functions, assets, or areas/zones can be designated as "no-fly zones" due to excessive risk involved with testing.

Operational impacts should also be considered arising from recovery actions or response taken during the exercise such as wiping endpoints or resetting passwords.

Major changes on systems being attacked during the exercise may increase the risk of unforeseen operational impact or reduce the effectiveness of the simulation. In this case, FIs should consider reviewing change plans with their IT Project Management Office or Change Management teams during the planning phase to be aware of any planned changes during the planned exercise timeframe.

Establish the Budget

When budgeting for such exercises, the FI should consider the intended frequency, complexity, size of scope and duration of the exercises, as well as the cost of potential service providers.

Consider Engaging External Providers

While larger FIs tend to have their own internal specialised teams, external providers may be engaged to offer independent views on conducting the Attack Modelling and the simulated attack. Maintaining the same provider across exercises provides benefits, such as building deep understanding of strengths and weaknesses of an organisation akin to real-world adversaries, an organisation may be breached multiple times over the years to develop understandings and situational awareness – this can be simulated by utilising the same provider.

Engaging different providers may on the contrary offer different insights, perspective and TTP.

The FI should perform their own assessment to determine whether keeping or changing providers would enable it to achieve the intended outcome.

7.1.3 Risk Management

AASE should always be conducted against live production to assess real-world cyber resilience.

Depending on the approach, certain threat scenarios could involve high-risk activities, and once carried out on a production environment, could lead to uncontained impact resulting in

damages greater than the benefit of the adversary simulation itself. Therefore, FIs should perform a risk analysis during the scoping phase.

If the assessment indicates that risks arising from testing in the real-world environment are deemed excessively high and could result in operational failure, certain elements of testing could be performed in non-production environments to minimise impact. This approach however does not allow for an accurate reflection of the FI's security state and should only be used in circumstances of very high uncertainty of operational risk.

If such a simulated approach is chosen, high-risk activities can be conducted through coordinated actions on both the production and a replicated environment, i.e., obtaining access on the production and then, using the same level of access on the replicated environment, demonstrate the ability of altering critical business data. For example, when targeting a payment processing infrastructure, the Attacker could try to obtain access to the production environment but will not make changes such as submitting payment messages or creating new user accounts; instead, these activities will be carried out on a replicated environment with the same configurations and level of access that they had previously obtained in the production environment. While potentially close to the real-life environment, this arrangement should be clearly documented in the reports. This method allows for the execution of a goal when only technical hurdles or challenges remain and cannot be replaced by subversion of business controls (such as legitimate maker/checker processes). Proper oversight and escalation paths should be established before the start of the test to avoid any outages of critical infrastructure.

A number of key risks should be carefully considered during the planning phase of the exercise.

Business Continuity and Operational Risk

Due to the level of sophistication of typical attacks, there is an inherent accentuated risk of operational failures of infrastructure being assessed during the exercise. Hence, it is usually prudent to have members of Business Continuity or Crisis Management teams intimately familiar with the exercise, unless an objective of the exercise is to invoke and assess those specific processes.

The Exercise Director and Exercise Working Group should be sufficiently informed and qualified to make decisions to drive the direction of the exercise should such events occur.

Legal and Regulatory Risk

Legal risks to be considered include legal liability as an outcome of adverse reactions to testing, such as failure to continue business operations.

FIs using systems or functions hosted by an external provider (i.e. Cloud service or external business function) should review the contractual terms of the service, as some providers do not allow attack simulations or require prior notice. The FI should consider these contractual obligations in the selection of Critical Functions to be tested.

Data Risk

The risks around data include data leakage, data destruction and alteration, and data unavailability.

Should the risks be deemed acceptable by the FI, the Attacker could be requested by the Exercise Working Group to provide access to exercise execution logs. However, since execution logs contain information which may allow replay of the attacks, proper controls would need to be implemented to ensure the security of this information.

Third-party Risk

Due to the nature of external engagements to perform AASE, third-party risk is one of the highest risk categories. Third-party risk may involve either malicious intent or reputational risk. It is important to mitigate third party risk by conducting adequate due diligence on service providers.

- **Malicious Intent**

There is a potential risk of the hired third party acting with malicious intent. Provisions should be made in the contractual phase that covers legal liability and indemnity in case of a potential rogue activity (e.g. theft of sensitive information, installation of malicious software like backdoors).

- **Insufficient Expertise of the Attackers**

There is a risk of insufficient expertise of the provider, which could in turn cause events that would expose the FI to operational risks due to execution errors or negligence. The provider selection process is an important step, and such risks should be duly considered.

Reputational Risk

All the risks as mentioned could potentially damage the reputation of the FI which can be mitigated with careful planning and execution.

7.1.4 Provider Selection

Due to the sensitivity of the systems potentially accessed, the provider selection process should at minimum seek assurances on appropriate testing methodology, data handling procedures, criminal background checks for its consultants, and adequate insurance and indemnity to cover legal liabilities.

When choosing a provider, its reputation, experience, and references should be considered.

Reputation can be established in industry collaboration forums, and references can be obtained from the potential provider.

7.1.4.1 Key Requirements of the Provider Organisation

Areas to consider when selecting third party AASE provider:

Methodology and Process

The provider should be able to curate threat intelligence and demonstrate its relevance to the FI as part of threat modelling. The attack scenario created in tandem with the FI should be realistic and align to the real-world adversary's tooling and motivation to target the FI's critical functions.

While executing the attack, the provider should constantly be in communication with the FI to update on the progress made, as well as to request for consent prior to execution of any actions that may have a high potential risk and consequence.

During the exercise, any indications of a concurrent real-world attack should be reported back to the FI. The provider should also be cognizant of what tooling and commands fall under this category and able to demonstrate capability to document actions done to assist in clean-up and reporting at the end of the exercise.

Any loopholes and vulnerabilities found as part of the exercise should be reported to the FI.

Reputation, Experience and Track Record

Providers should demonstrate strong capability by way of possessing recognised accreditations or having proven track record in analysing and replicating adversarial techniques from an offensive and defensive perspective, as well as providing evidence that they have helped to enhance their clients' organisational resilience and defensive capabilities. Any experience in managing incidents or identify threats in the environment would be a plus. A sound reputation in the industry across multiple geographies and jurisdictions is usually associated with a well-rounded provider.

Depending on their technology stack, FIs may want to consider the following non-exhaustive technical competencies:

- Experience running complex red team operations.
- Experience deriving realistic scenarios from Threat Intelligence.
- Understanding of and capabilities in cloud technologies.
- Understanding of and capabilities in DevOps and Continuous Integration / Continuous Delivery.

- Understanding of and capabilities in Internet of Things and/or Operational Technology devices.
- Expertise in defence evasion.

It may also be of interest for FIs to evaluate the track record of the actual personnel involved in the exercise.

Data Protection

Providers should have robust data management procedures that cover the protection of data at rest and data in motion, data retention and disposal procedures.

Background Checks

Due to the high sensitivity of the work performed, each provider should have standard practices of criminal record and other background checks for all consultants hired by them and trained in applicable regulatory compliance obligations.

7.1.4.2 Provider Obligations

Providers should be able to demonstrate a real world understanding of sophisticated adversarial TTPs, which can be substantiated through publicised research, experience, and recognised capability in responding to sophisticated cyber incidents. Providers should also be able to continuously communicate with FIs across the whole AASE process, and guiding FIs along who may not be familiar with the processes around this type of engagement.

In addition, it is important that the selected provider has considered the obligations below.

Compliance

The provider must comply with all relevant laws and regulations during the entire course of the exercise.

FIs should ensure that the provider is kept legally liable for its activities by stating such requirements in the Scope of Work, as well as through signing a Non-Disclosure Agreement.

FIs should ensure there are enough legal policies and procedures to be legally compliant, insured and indemnified for the type and nature of work. Specifically, considering the possibility of unintended damage to the systems being attacked, the FI should pay particular attention to mutual damages and liabilities, and to the provider's financial capacity to withstand such an eventuality.

Evidence Keeping

The provider should be required to keep only enough information as is required to provide evidence that specific exercise goals have been achieved and to reproduce the methods used to

achieve the goal. They should not be allowed to manage and keep data belonging to the FI after the completion of the exercise, unless necessary for future exercises.

When third parties are employed to conduct the attack, it is recommended that a secured storage space for data is set-up to securely store the data obtained from the FI's system. Such practice could facilitate the removal of data after the exercise.

Data Retention and Destruction

The provider typically only needs to retain a small number of artefacts it obtained during the exercise to produce the exercise a report, such as screenshots or small pieces of data. The contractual obligations may be specified to require the provider to only keep such reports for a limited amount of time. All other data acquired by the provider through the course of the exercise should be destroyed upon the exercise completion.

The provider should reduce the amount of real, sensitive, production data (e.g. classified business information) sent outside of the FI's network while still achieving the exercise objectives. This might be achieved by the provider exfiltrating representative data only, thereby allowing it to prove the attack path but limiting exposure of the sensitive data.

Physical Supervision

Depending on the exercise mode of operation and goals desired, certain elements of the exercise execution may benefit from real-time physical supervision by the Exercise Director, such as final execution of an action upon a goal, or physical access to premises. This should be considered in the planning stages.

7.1.4.3 Key Requirements of the Attacker (Threat Intelligence and Attack Teams)

Ideal characteristics include:

- Demonstrable technical skills, expertise, industry experience and competency relevant for the type of work provided, considering the industry and specific requirements, related to AASE.
- Demonstrable understanding of real-world adversarial activities.
- Ability to communicate well with senior management as well as technical experts.
- Education on the legal and compliance aspects to enable management of legal and regulatory frameworks applicable to clients and their relevant geographies and industries.

The Threat Intelligence and Attack teams should be able to demonstrate expertise in selecting and using TTPs for the exercise, for example:

- Malware samples from known threat groups/actors (e.g. banking trojans, custom built ATM malware targeting specific protocols, command and control RAT, etc.).
- Command and Control ("C2") traffic and protocols (e.g. ICMP, DNS, HTTP/S, TCP, etc.).
- Type of different infrastructure that can be used (e.g. Service provider, Known C2 hostnames and IP, etc.).
- Vulnerabilities that were exploited by known threat group/actors.
- Known lateral movement and persistence techniques.
- Operational hours.

Providers having a dedicated research and development team that can create custom tooling for the exercises based on such TTPs would be a plus.

Threat Intelligence Team Specialist Requirements

The Threat Intelligence specialist should be able to perform with high level of confidence in all areas of threat intelligence, operational security, data collection, data analysis, intelligence production, scenario development and enrichment.

The threat intelligence specialist should be able to provide insights to the following questions:

- Who may want to attack the organisation?
- Why would they want to attack the organisation? (What are the motives for the attack)
- How would they attack the organisation? (Tactics, techniques and procedures of the adversary)

Attack Team Specialist Requirements

Specifically, the Attacker team should consist of specialists in Adversarial Attack Simulation with proven track record of:

- Technical expertise in skills such as reconnaissance, perimeter breach, persistence, lateral movement, and privilege escalation.
- Understanding and translating of sophisticated adversarial TTPs.
- A strong understanding of a wide breadth of typical preventative and investigative security technologies.
- Executing actions in both realistic test environments as well as real environments in a safe and legal manner.
- Strategic understanding of capabilities required to increase organisational capability

to detect and respond to TTPs executed.

- Producing written documents and delivering presentations to all levels of business.

7.1.4.4 Qualifications (including Certification and Accreditation)

The provider and all its practitioners should hold accreditations or certifications that require practical demonstration of an advanced understanding of offensive and/or defensive security skills and the tactics, techniques and procedures utilised by real world sophisticated adversaries. These accreditations should also require that providers, and practitioners, demonstrate an understanding of legal and operational security frameworks that apply to exercises of this nature.

There are accreditation bodies that provide certification and accreditation for this specific type of testing, and therefore such qualifications should always be sought. Upon request by the FI, the provider should supply the list of consultants that are assigned to the project, together with their biographies.

It should also be noted that due to the regional reach of accreditation bodies, there will be tendency for a provider from specific regions to hold specific certifications and accreditations, and FIs should not discount providers solely due to the fact that they do not hold a particular accreditation, or its practitioners do not hold a particular certification.

7.1.5 Concessions

Concessions that may be ceded during an exercise should be planned. These concessions can include software, hardware, people, processes, etc. Some examples could be:

- A dedicated Identity & Access Management ("IAM") administrator instructed to issue administrator or plain accounts on demand to the Attacker.
- A local IT person instructed to issue laptops or tokens on demand.
- Information about assets, processes, information systems, or key personnel attached to a function.
- Access control bypass to a network outlet that would otherwise be physically secured.
- A trusted agent to unzip and runs malware infected file to provide initial entry point for the attacker.

Some concessions can be provided upfront (e.g., assume an endpoint compromise will be successful and provide a laptop to be used for the exercise), and some can be issued during the exercise, at points where the Exercise Director establishes that they would be beneficial for the outcome (e.g. deliberately distracting incident responders to prolong the attack vector's duration).

In the event where all concessions are exhausted (including providing internal foothold), FIs may consider allowing the attacker to continue the exercise with the blue team passively monitoring.

All concessions used, whether upfront or during the exercise, should be recorded, their impact analysed and reported in the Exercise Report.

7.1.6 Communications

7.1.6.1 General Awareness

An AASE's efficacy is highest when absolute secrecy is maintained in the FI. The major benefit of secrecy is that the entire response to the attack from all parts of the FI is genuine and organic, and therefore truly reflects the shape and form of the incident response. However, in AASE where secrecy was not preserved, the information leaked may negatively impact the behaviour of the incident response teams substantially and may compromise the benefits of the exercise. (Refer to the section 6.2 Exercise Secrecy)

7.1.6.2 Use of Code Names

To protect the secrecy of the exercise, the Exercise Working Group should use code names to refer to the exercise utilising terms that do not reveal the existence or nature of the exercise.

7.1.6.3 Stakeholder Management

As the AASE is a holistic approach to test the cyber resilience of an entire FI, the organisation's Board of Directors should be fully supportive of the principles, goals, and implementation of the exercise. These duties are typically delegated to the CISO, CIO, or a senior executive on behalf of the Board. The Board, or sponsoring management, should be informed of the intent of the exercise and the final outcomes. The senior management should review and approve the final report and commit to remediation efforts via a management action plan.

7.1.6.4 Regulatory Disclosure

Financial Institution regulators in certain countries require some form of AASE to be executed periodically, while others suggest it or remain silent. It may be prudent to advise the regulators of the intent to execute such exercises upfront. It should also be noted that in certain cases some adverse reactions may occur which would impact the FI's ability to operate as required and may even require actions such as regulatory notification.

7.2 Attack Preparation Phase

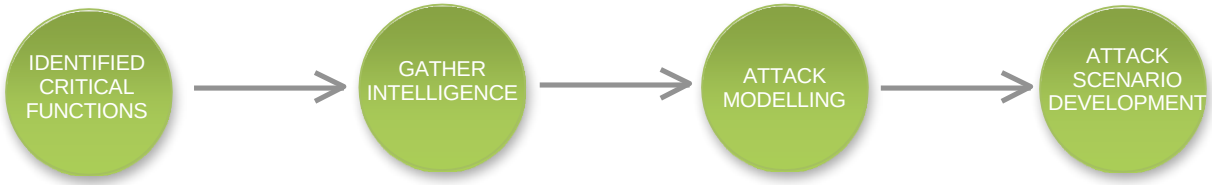

During the attack preparation phase, the Exercise Working Group will develop attack scenarios that would impact Critical Functions or business continuity.

Firstly, Threat Intelligence is gathered to determine the tooling and known capability of relevant threat actors of concern. Information is gathered, and the range of possibilities are expanded. With that information available, an attack model is created that helps define possible ways to attack the organization based on known TTPs, tooling and threat actor motivations. Finally, the possible attack paths from the attack model into attack scenarios that broadly define how the exercise will be executed.

7.2.1 Identified Critical Functions

Critical Functions are business functions or services that if compromised, would significantly impact business continuity, affect the reputation of the FI, or cause financial loss. Identification of Critical Functions should ideally be done on a regular basis. The Critical Functions identified during the Planning phase will be used to formulate the end goal of the Attack Scenario.

7.2.2 Gather Intelligence

In this step, threat intelligence and information about the target organization are collected. In gathering threat intelligence, the FI aims to determine probable Threat Actors who will target the organization and the TTP that may be employed. This threat intelligence gathering can be performed by threat intelligence analysts under the organization's employ, attained from external service providers, or a mixture of both.

For the purposes of the AASE, the threat intelligence that needs to be gathered relates to possible threats that may impact the FI from Threat Actors known and unknown. This can be conducted in the following two ways:

7.2.2.1 Generic Threat Intelligence

Generic Threat Intelligence refers to threat intelligence information not specific to any particular FI, but are threats generic to the industry or a specific geography. As examples, this could refer to

reports of incidents affecting another institution in the Financial Services sector, or another organization within the same geography.

Intelligence sources could be publicly available or paid intelligence reports, information shared on online sources such as social media, chat groups or news sites, information gathered from the Dark Web, or even data from the MITRE ATT&CK matrix; This is but an unexhaustive list for consideration.

An approach could be to consolidate the available threat intelligence to identify vulnerabilities exploited, tooling applied, and Tactics, Techniques and Procedures (TTPs) utilized during these incidents.

7.2.2.2 Targeted Threat Intelligence

While generic threat intelligence focuses on incidents affecting the broader sector or geography, targeted threat intelligence looks at threats that are specifically relevant to the organization coming from specific Threat Actors. This requires a deeper understanding of the possible motivations of Threat Actors to understand their possible goals and targeting. This may also take into consideration any geopolitical issues brewing in the recent timeframe.

By looking from the perspective of known, named Threat Actors, FIs should look to identify Threat Actors that would plausibly pose a high risk to the organization, and identify the plausible end goals of each of these high-risk Threat Actors in targeting the firm.

An approach could be to identify known tooling and TTPs exhibited by these Threat Actors, to identify plausible attack vectors that may be utilized by them in the event of a targeted attack on the organization. It is also important to try to identify possible goals of such attacks.

7.2.2.3 Target Reconnaissance

On top of understanding the threats facing the organization, the Attacker should also have a good understanding of the targeted organization FI to be able to formulate a good Attack Model. At this point in time, it would mostly involve passive reconnaissance and open-source intelligence gathered off the open web and public domain. A profile could be built for the organization, encompassing information that could be utilized during the execution to make the attack more convincing. For example, this could be email subjects used for targeted phishing campaigns, or specific individuals of interest that would be targeted.

7.2.3 Attack Modelling

In Attack Modelling, the Exercise Working Group taps into all the threat intelligence gathered, and looks to integrate the various information points. Identified Threat Actor motivations and tooling are combined with TTPs across different campaigns, to paint potential attack pathways. Where there are gaps in the intelligence, the Exercise Working Group should also look to patch them based on current research or commonly used attack techniques.

The Attack Modelling step aims to refine the threat intelligence gathered into actionable plans to attack the organization.

7.2.4 Create Attack Scenario

In the last step of the preparation phase, the Exercise Working Group develops attack scenario(s) based on the inputs from the earlier steps.

From the Attack Model created, specific Attack Scenarios are created based on the probable attack pathways. The convergence of attack pathways into a smaller subset of attack scenarios will consider various factors; For example, the infrastructure used by the organization based on target reconnaissance, the availability of the technological capability of the Attackers to emulate the Threat Actor, or current priorities of the Exercise Working Group relating to the scope of the attack and the risk of the attack resulting in early conclusion of the exercise (eg. Detection by the SOC).

Attack scenarios should also include the success criteria in compromising a critical function. For example, where the target is a payment processing gateway, the FI may deem that the attack is successful based on the Attacker's ability to gain access to payment users' credentials and their access to the payment systems, instead of requiring the Attacker to fully demonstrate an actual fraudulent transfer. In other cases, the end-to-end demonstration may be required to test the effectiveness of all controls including the payment authorization, the reconciliation with the logs as well as the various limits and alerts that may have been put in place. The Exercise Director should be consulted at the point in time before such attacks are executed, to determine the risk and impact of such an attack and decide what would be the ideal action to take.

The end goal of the Create Attack Scenario step is to create realistic attack scenarios for the Attacker to simulate. This scenario should include the initial access vector, information or artifacts that the Attacker is looking for once in the environment, and the end goal of the Attacker. Scenarios should broadly paint the entry vector of the attack and specify a success criterion as an end goal based on a Threat Actor's motivation, without being too prescriptive in directing how the attack will unfold.

For example, a scenario could be “The Attacker utilizes phishing to get a foothold in the organization by targeting privileged users. Subsequently, it aims to locate and pivot to SWIFT servers to transfer $100,000 in funds via fraudulent transactions.” Another example could be “The Attacker targets the software supply chain to gain a foothold in the organization, heavily utilizing PowerShell to move laterally and find sensitive data for exfiltration. The exfiltrated data is then used to extort the organization. This data should be either sensitive customer PII or important server configurations or private code.”

At this stage, possible Concession points could also be planned out, for example to allow the Attackers to skip over an undefeated control or to simulate a threat vector such as a zero-day vulnerability or a supply chain compromise. Such concessions aim to maximize the value of the engagement but should only be applied when the Exercise Director determines that the Attacker is unable to proceed, or has tested a control without success for an extended period of time, or if the scenario portrays the use of a TTP that is a plausible risk factor but not technically feasible in a simulation (e.g., A 0-day vulnerability). The use of concessions should be recorded in the final Exercise Report, and risks from subsequent findings should be factored for accordingly.

7.3 Attack Execution Phase

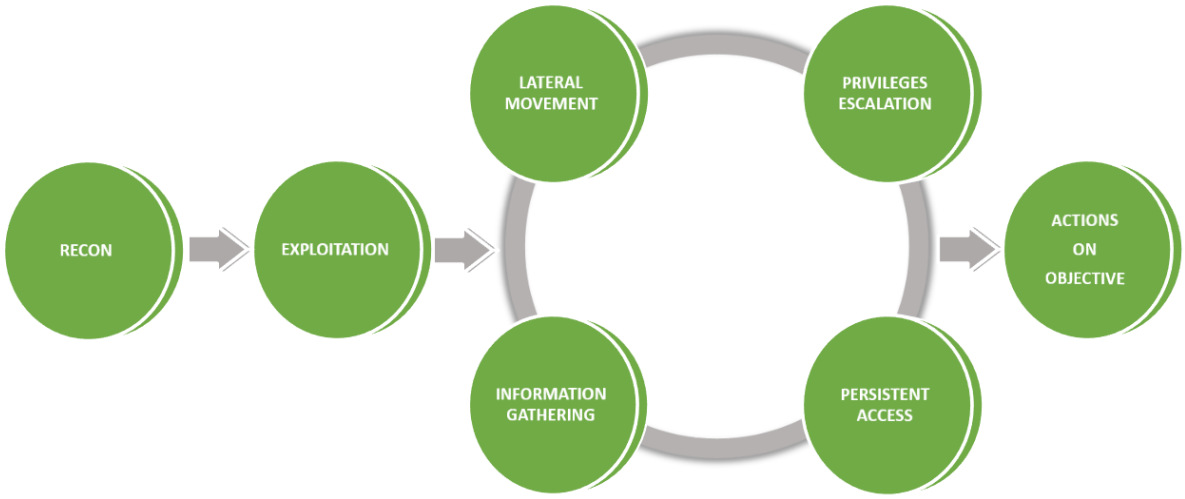

The Attack Execution phase involves the execution of the attack on the identified Critical Function based on the attack plan and scenarios that are formulated in the previous phase. During the attack execution phase, the Exercise Director should closely monitor the progress of the exercise and supervise the Attackers. While it is recommended for the Exercise Director to have full supervision at all stages, his involvement at certain stages where there is no impact such as reconnaissance can be reduced to a "keep informed" basis.

All steps taken by the Attacker and their observations should be continuously captured in the Exercise Log Report with the level of details determined in the requirements during the preparation phase. The Execution Log Report will also contain changes of plan, use of concessions and relevant approvals provided by the Exercise Director. As mentioned in the Risk Management Section, sharing the details of the attack with the Exercise Working Group carries a risk should the content be exposed unwittingly to either the Defenders or a real- world adversary.

To reduce risks during the attack execution phase, the "test/halt engagement model" should be applied.

7.3.1 Test/Halt Engagement Model

AASEs are a risk-based approach to holistic systemic testing. Simulated attacks are planned and analysed before they are executed, therefore the fallout and consequences can be predicted and minimised. These exercises typically consist of high- r i s k activities, so the attack plans should be carefully analysed before approving them.

Once the exercise has begun, the Exercise Director should be in supervision of the efforts and is able to decide whether a particular attack path should be pursued, diverted, or if the entire exercise should be paused or terminated, depending on desired outcomes and potential adverse reactions.

Publicly exposed vulnerabilities

In the event where high-risk security vulnerabilities that pose immediate threats to the FI are identified by the Attacker during the exercise, the Attacker should inform the Exercise Working Group immediately. Examples of such high-risk security issues may include vulnerabilities allowing arbitrary code execution in internet-facing systems. The Exercise Working Group should determine whether such a finding would need to be escalated for remediation immediately, as this kind of escalation might compromise secrecy of the ongoing exercise. If such a situation occurs, access to the system is given to the Attacker to assume that the vulnerability would have been successfully exploited.

Evidence of prior compromise

If the Attacker discovers evidence or indication of prior compromise during the exercise, they should immediately escalate such findings to the Exercise Working Group for investigation and deliberation on the future course of the exercise. The FI should seek to complete the incident response with regards to the said findings immediately, before continuing to pursue goals of the exercise, even if this requires halting or pausing of the exercise.

Attribution of TTPs during the exercise

If the exercise is designed such that the Exercise Director has insight into the incident response processes of the FI, then the Exercise Director should be able to identify incident notifications that are related to the exercise. The service provider should establish a direct hotline that the Director can use to validate and attribute TTPs observed in incident notifications that point to attacks made by the Attacker. This can be used by the Exercise Working Group to potentially intervene, either to diffuse or to delay/alter the incident response, should this mode of exercise be chosen.

A typical attack is not necessarily linear but will often involves repeating some of the following stages based on the position of the Attacker.

7.3.2 External Reconnaissance and Perimeter Breach

A reconnaissance is performed by the Attacker from outside of the perimeter, to find entry points into the target FI. This consists of discovery exercises, to identify potential weaknesses within external infrastructure and identify legitimate services provided by the FI.

Attackers will also review legitimate sources of information which may have been inadvertently or intentionally disclosed (job descriptions as an example). This information is used to inform further stages of the exercise.

Entry points typically include (but are not limited to):

- Physical breaches
- Infrastructure breaches
- Social engineering (phishing, vishing, etc.)
- Supply chain attacks

The Attacker may choose to adopt techniques to bypass security controls and obtain access to corporate portals used by employees that may be externally accessible.

7.3.3 Lateral Movement

Attackers will typically look to move between, and gain access to, various systems within an environment. Lateral movement typically involves the compromise of additional systems or processes until an adversary has access to systems deemed and believed to be necessary for the success of the simulated adversary's objectives. Access to further systems can assist attackers in:

- Increase level of access
- Evade detection
- Obtain further situational awareness and information
- Gain further perspectives on the network

Systems typically targeted for lateral movement can include:

- Asset management systems
- Anti-virus management systems
- Software delivery mechanisms

The Attacker may choose to adopt some of the above techniques during the exercise.

7.3.4 Internal Information Gathering

Attackers will perform internal reconnaissance to achieve situational awareness and inform further stages of the exercise.

7.3.5 Privileges Escalation

Attackers will typically attempt multiple escalating privileges methods to other systems or users on roughly the same level (horizontal privilege escalation) and higher privilege users or system (vertical privilege escalation). This should be performed in a way that mimics real- world advanced targeted breaches.

7.3.6 Persistent Access

Attackers may try to maintain persistent access upon breaching a network. The type of persistence techniques that may be used depends on the environment, considering relevant controls.

7.3.7 Action on Objectives

Attackers will proceed with the attack after he is satisfied that he has the resources and the right condition to perform the attack.

7.3.8 Directing the Attack (Escalation Path - Chain of Control)

The Exercise Director has the oversight of course of actions of the Attackers and can alter the course of actions of the Attackers or even alter the environment and situation if deemed beneficial for the outcomes of the exercise, or if the potential risks and consequence is high should there be a breach.

7.3.9 Use of Concessions

Concessions can be used by the Exercise Director to alter the course of the exercise. When used, they should be documented and reported.

7.4 Exercise Closure Phase

The exercise should be called to a close when either all outlined objectives have been met, or when the planned timeline has ended.

7.4.1 Clean-up and Tactical Vulnerability Containment

The primary objectives of clean-up activities and tactical vulnerability containment, conducted based on the information provided in the final report, are to remediate immediate issues found during the exercise, as well as eradicate any left-over attack tools and artefacts.

A clean-up report should be prepared by the attacker as soon as the attack report is published. Internal parties are typically best placed to work with the attacker to execute the clean-up, due to knowledge of internal processes, legitimate credentials and access held and knowledge of controls or intricacies that the attacker may not have visibility of. The deliverable is outlined in section 8.3.2.

7.4.2 Reconciliation

The reconciliation phase should be used to identify security controls, either missing or in need of improvement, that would otherwise have prevented or detected the attack. This phase should identify all aspects of the exercise where the attack was detected, observed, reacted upon, tracked, contained, and eradicated, or lost sight of. The deliverable is a defence report outlined in section 8.3.3. The defence report is optional if the exercise report could explain a balanced view from the defenders.

7.4.3 Attack/Defence Joint Replay

A joint post-attack exercise can be organised to enable a step-by-step replay of the attack, either as a table-top or lab exercise for the Defence team's learning benefit, or in the actual live environment, to demonstrate failed controls in real-time.

In addition, the replay could include some hypothetical attack simulation for the benefit of the Defenders. The Attackers would then explain other steps they could have taken following another path and the steps they would have taken if the Attackers had more time or resources.

Lastly, the replay offers an opportunity for the Attackers and the Defenders to provide an opinion on the various activities performed throughout the exercise as potential improvement points. The Attackers could also offer an opinion on the performance of the Defenders, as a comparative analysis with similar organisations the Attackers worked with, should the Attackers be a third-party provider.

7.4.4 Exercise Reports and Recommendations

The Exercise Report, also known as the Attackers' Report, is the final report published by the Attackers after the exercise is concluded. It describes the actual scenario-based attack as it played out, listing the attack elements that were critical to the success of the attack, e.g., weaknesses discovered that enabled the Attackers to progress to the next stage. This report is used as the source of information for remediation and clean-up activity planning. The deliverable is outlined in section 8.3.1.

7.5 Post Exercise Phase

The post exercise phase comprises of action plan formalization and remediation timelining based on the recommendations from the exercise report. Fis may also consider sharing the outcome of the AASE with a wider audience for common benefit.

7.5.1 Strategic Remediation Management Action Plans

The FI should formalise a remediation timeline based on the exercise report, along with the organisation's risk appetite and compensating controls. The strategic remediation may involve items such as process changes within the FI, tightening of security controls, additional investments, end-user security training, architecture redesign, etc. The risk assessment should take into consideration the sensitivity of the data residing in the systems and the location of these systems.

A project manager should be assigned to track these actions to closure, with periodic reporting to the Sponsoring Executive or senior management on the progress. Details of the remediation management action plan is outlined in section 8.3.4.

7.5.2 Sharing with a Wider Audience

The outputs of the exercise can be summarised into a "Lessons Learned" document, redacting details that would otherwise expose sensitive information, that can be shared with other members of the industry or a broader community for the common benefit.

8 Documentation

This section provides indicative samples of deliverables to be tailored to suit individual FIs, depending on level of services provided and outcomes desired. The table below summarizes the stakeholders involved in each of the phases and the suggested priority of the reports. An artifact is a supporting document which contributes to the writing or outcome of a report.

Phases	Reports / Artifacts	Owners	Priority
Attack Preparation	Threat Intelligence Report	FI's Internal TI / TI Provider	Strongly Recommended
	Attack Modelling Artifact	Exercise Working Group / TI Provider	Recommended
	Exercise Preparation Report	Exercise Working Group	Recommended
	Target Reconnaissance Report	Attacker	Optional
Attack Execution	Execution Log Artifact	Attacker	Recommended
Exercise Closure	Exercise Report	Attacker	Strongly Recommended
	Clean-Up Report	Attacker	Strongly Recommended
	Defence Report	Defender	Optional
	Remediation Action Plan	Exercise Working Group	Strongly Recommended

8.1 Attack Preparation Phase

8.1.1 Threat Intelligence Report

The Threat Intelligence Report is produced either by the Attacker or by a separate group (can be a different provider) that specialises in Threat Intelligence. Threat Intelligence analysts will analyse the kinds of threats that are currently prevalent, either non-specific or specific to the FI, and will put this information in a report that is used to create plausible and credible scenarios of attack.

In the Threat Intelligence Report, the intent and capability of the threat actors are also assessed and ranked using a threat matrix and summary table. The FI would then use the information in this report to subsequently plan the exercise. The Threat Intelligence report may consist of either a generic and / or targeted Threat Intelligence type.

Generic Threat Intelligence

Generic Threat Intelligence should contain information on potential real-world attackers' modes of operation and their TTPs. This does not necessarily apply to the FI being targeted but applies to most organisations in general.

Targeted Threat Intelligence

Targeted Threat Intelligence should contain information on threats, attackers and TTPs specific to the FI or its industry type. This may include threats that are specific due to a certain geopolitical condition, FI's involvement in certain projects, competitor landscape, etc.

Report structure sample

1. EXECUTIVE SUMMARY

2. THREAT MATRIX

3. THREAT SUMMARY TABLE

4. INTRODUCTION

5. CYBERCRIMINAL THREATS

Organised Cybercriminal Groups (OCG)
DDoS extortion
Ransomware extortion
Data leak extortion
Malicious Insiders
6. NATION STATE THREATS (NST)

Nation State Attacker 1
Nation State Attacker 2
Nation State Attacker 3
Nation State Attacker 4
7. CYBER ACTIVIST THREATS

Cyber Activists
8. THREAT SCENARIOS

Scenario 1: OCG1 attempts fraudulent transactions via payment and settlement systems
Scenario 2: NST clients to exfiltrate sensitive intellectual property
Scenario 3: OCG2 compromises various components of the retail banking function
ANNEX 1: CYBER THREAT DEFINITIONS

Threat matrix sample

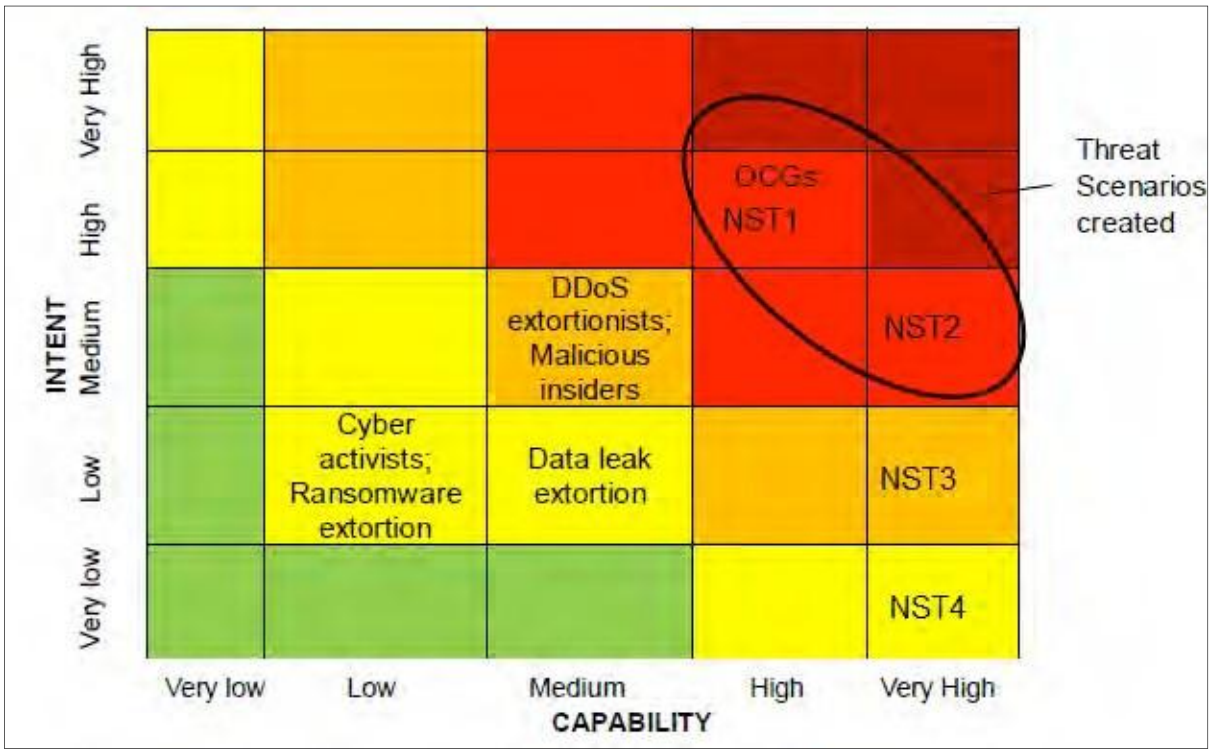

Threat Summary Table Sample

Threat Actor	Intent	Capability	Threat	Summary
Organised cybercriminal groups (OCGs)	High	High	HIGH	OCGs are the most sophisticated of cybercriminal actors and have demonstrated their capability to compromise various types of systems in scope for this engagement. Although they have been more active in financial centres other than Country X, Organisation X will still likely represent an attractive target.
NST1	High	High	HIGH	
NST2	High	Very High	HIGH	

DDoS extortionists	Medium	Medium	MEDIUM	DDoS extortionists have successfully targeted organisations with a similar profile to the organisation in the past and may look to disrupt the organisation's public-facing web portals to extract ransoms.
Malicious insiders	Medium	Medium	MEDIUM	Malicious insiders' privileged access to key systems and information potentially renders them among the most capable actors in this assessment, though we have uncovered no explicit evidence to suggest that insiders are looking to target Organisation X.
NST1	Low	Very High	MEDIUM	
Ransomware extortion	Low	Low	LOW	Organisation X will likely face a high volume of ransomware extortion activity, though the perpetrators will generally lack the intent or capability to specifically target the in-scope systems.
Data leak extortion	Low	Medium	LOW	Data leak extortionists may have the ability to breach and threaten to release sensitive information from Organisation X key systems, though are more likely to pursue smaller and more vulnerable organisations.
Cyber activists	Low	Low	LOW	With the possible exception of launching DDoS attacks against online banking portals, cyber activists lack the capability required to target Organisation X in-scope systems. This will also deter their intent to do so.
NST2	Very Low	Very High	LOW	

8.1.2 Attack Modelling Artifact

The Attack Modelling Artifact is prepared by the exercise working group and / or Threat Intelligence providers to enhance attack scenario creation. The artifact should detail how the table of threat events were derived and included in the Appendix of the Exercise Preparation Report.

Artifact structure sample

Appendix – Attack Modelling Artifact Table of Attack Model Scenarios

Table of Attack Model Sample

TI Threat Scenarios	Point of Entry	Attack Event	MITRE Tactics	MITRE Techniques
OCG1 attempts fraudulent transactions via payment and settlement systems	Enterprise Workstation	OCG1 perform targeted phishing attacks on SWIFT system user and deploy Remote Access Trojan (RAT) to establish a backdoor. OCG1 captures user credentials for access to SWIFT servers. With credentials obtained, OCG1 lateral move to SWIFT server and attempt to transfer $100,000 in funds.	Initial Access	Phishing
			Execution	Command and Scripting Interpreter
				Native API
			Credential Access	Input Capture
				Credentials from Password Store
			Lateral Movement	Remote Services
			Impact	Financial Theft
	Web Server	OCG1 exploited N-Day vulnerability to deploy Remote Access Trojan (RAT) to establish a backdoor. OCG1 discovers SWIFT user group and lateral move to targeted user. OCG1 captures user credentials for access to SWIFT servers. With credentials obtained, OCG1 lateral move to SWIFT server and attempt to transfer $100,000 in funds.	Initial Access	Exploit Public-Facing Application
			Execution	Command and Scripting Interpreter
				Native API
			Lateral Movement	Internal Spearphishing
				Taint Shared Content
				Remote Services
			Credential Access	Input Capture
				Credentials from Password Store
			Impact	Financial Theft

8.1.3 Exercise Preparation Report

The Exercise preparation report is created by the Exercise Working Group during the initial preparation phase. It will contain the details and decisions for the exercise.

Report structure sample

1. *Introduction including the team composition (name and role of each team member)*
2. *Exercise Preparation*
2.1 Exercise Description
2.2 Exercise Code Name
3. *Attack Execution*
3.1 Objectives and Scope
3.1.1 Attack Objectives
3.1.2 Attack Scenarios
4. *Communication Management Strategy*
4.1 Contacts
4.1.1 Organisation
4.1.2 Attackers
4.2 Email and File Exchange
5. *Risk Management Strategy*
5.1 Risk Mitigation
5.2 Organisation's Escalation Process
5.3 Attackers' Escalation Process
6. *Governance, Control and Reporting*
6.1 Test Plan and Reporting Progress
6.1.1 Test Plan
6.1.2 Weekly Progress Meetings
6.2 Final Report
7. *Legal and Liability Insurance*
8. *Appendices*
8.1 Appendix – Methodology
8.2 Appendix – Attack Modelling Artifact
8.2.1 Table of Threat Events

8.1.4 Target Reconnaissance Report

The Target Reconnaissance Report prepared by the Attacker contains all the information about the targeted organisation. It will help the Attacker breach the organisation and move closer to the objectives by understanding the various controls in place to bypass as well as the overall attack surface. The Target Reconnaissance Report is optional if the attacker is aware of the attack surface.

Report structure sample

1. *Executive Summary*
2. *Targeting Organisation*
2.1 Phishing Campaigns
2.2 Security Controls (Email filtering, File integrity monitoring, Egress restrictions, Office macros, …)
3. *Key Systems (Retail banking, Commercial banking, Trading and sales, Payment and settlement, …)*
3.1 Key Personnel (Related to the systems identified above)
3.2 Key Suppliers (Related to the systems identified above)
3.3 Technical Details (Related to the systems identified above)
4. *Organisation and People*
4.1 Organisation Structure
4.2 Organisation Policies (Employee benefits, Whistle blowing campaigns)
4.3 Organisation Announcements (Acquisition, Deployment of large solutions)
4.4 Organisation Awards (Best employers, Safest Bank, ...)
4.5 Organisation Events (Sports related, charity, ...)
4.6 Social Media Presence (Facebook, Twitter, Instagram, LinkedIn, ...)
4.7 Office Locations
4.8 Key Customer, Supervisory and Audit Relationships (MAS, ABS, ...)
4.9 Employees
5. *General Networks and Systems*
5.1 Autonomous System Numbers (ASN)
5.2 IP Netblocks (IPv4 and IPv6)
5.3 DNS Domain Names
5.4 Registered X.509 Certificates
5.5 Enumerated Hosts
6. *Email*
6.1 Gateway Infrastructure
6.2 Content
6.3 Address Format
7. *General Technologies*
7.1 Critical Functions
7.2 Cloud
7.3 Security
7.4 Networking
7.5 Server
7.6 Desktop
7.7 Mobile
7.8 Other
8. *Key Suppliers*
9. *Appendix I - Reconnaissance Methods*

8.2 Attack Execution Phase

8.2.1 Execution Log Artifact

The execution log artifact contains information on resources that were utilized to achieve goals and/or scenarios and highlight the attack path taken by the Attackers during the execution phase.

It is a live document that should be maintained by the Attackers throughout the exercise, providing

an overview of the exercise and describing the course of actions, proposed future strategies, decisions made, changes in direction and other such reasoning. The documented artifacts are used to compose the Exercise Report detailed in the next section.

Artifact structure samples

1. *Attack Path*
1.1 Network foothold, sustained access, action on objectives
1.2 Out of scope techniques: Phishing attacks on Organisation's customers, Compromising third party sites as part of a watering hole attack
1.3 Command and Control (such as Cobalt Strike) logs
2. Concessions and specific Information for each stage
2.1 Reconnaissance
2.2 Staging
2.3 Exploitation
2.4 Control and Movement
2.5 Persistence and Egress
2.6 Attack Execution
3. *Attack Attribution (Protocols used to clarify whether an actual unauthorised attacker has gained access to sensitive systems or data, and to prevent unnecessary initiation of expensive incident handling procedures in response. On the other hand, if an attack attributed to the actual unauthorised attacker is detected, then this provides the opportunity to halt the AASE to avoid interfering with or confusing the incident responders.)*

8.3 Exercise Closure Phase

8.3.1 Exercise Report

The Exercise Report should contain security strengths, comprehensive analysis of organizational capability and recommendations on remediation, both tactical and strategic, based on the Attackers' expertise and experience.

The Exercise Report should also include the methodology, evidence of goals achieved, details of the attack path undertaken, and any concessions used. A timeline containing pertinent information about the attack as it unfolded could be included, and should include a balanced view from the defender, to reconcile the activities of both sides.

Report structure sample

1. Executive Overview
1.1 Strategic Recommendations
1.2 Security Benchmark
2. Project Scope and Methodology
3. Findings
4. Recommendations
5. Prevention
6. Detection & Response
7. Attack Details
7.1 Threat Intelligence and Modelling
7.2 Reconnaissance
7.3 Phishing
7.4 Local Privilege Escalation
7.5 Information Gathering
7.6 Use of concessions and deviation from live environment
7.7 Lateral Movement
7.8 Researching Attack Objectives
7.9 Goal 1
7.10 Goal 2
7.11 Goal 3
7.12 Two-Factor Authentication
8. Appendices
8.1 Appendix I - Phishing Campaigns
8.2 Appendix II - Indicators of Compromise
8.3 Appendix III - Compromised Assets
8.4 Appendix IV - Incident Response
8.5 Appendix V - Old Operating Systems
8.6 Appendix VI - Compromises via technique 1
8.7 Appendix VII - Cleaning Up
8.8 Appendix VIII - Approach to Testing

8.3.2 Clean-up Report

The Clean-up Report should detail IOCs and artefacts on any residual data or changes made during the exercise, for removal by an internal system administration team.

Report structure sample

1. Introduction
2. Initial Access
2.1 Phishing Target A
2.1.1 Campaign 1
Date, File / Payload Name, File / Payload Hashes, Target Recipients, and any other details necessary for clean-up should be included
2.1.2 Campaign 2
Date, File / Payload Name, File / Payload Hashes, Target Recipients, and any other details necessary for clean-up should be included
3. Compromised Systems
3.1 System Name 1
Domain username, artifacts hash, persistency deployed (e.g. Scheduled Task, WMI Event Triggered Execution, etc)
3.2 System Name 2
3.3 Domain username, artifacts hash, persistency deployed (e.g. Scheduled Task, WMI Event Triggered Execution, etc)

8.3.3 Defence Report

A Defence Report, or a Blue Team report, should be prepared by the defence teams based on the Exercise Report, to reconcile the attack timeline from the defenders' point of view. It should describe at which stages and at which points the Attackers were uncovered (if any), or their artefacts encountered.

Report structure sample

1. Introduction
2. Reconnaissance (External / Internal)
Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.
3. Initial Compromise
Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.
4. Establish Foothold
Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.
5. Persistence
Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.
6. Privilege Escalation
Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.
7. Lateral Movement
7.1 Timestamp, Hostname, Username, Action taken, Artifact Name (if any), Artifact Hash (if any), Rules, etc.

8.3.4 Remediation Management Action Plan

The action plan should outline the identified gaps through the course of the exercise, and the Management's committed activities to remediate them.

Action plan sample

Identified Gaps	Gaps Description	Recommendation	Team Responsible	Remediation Plan	Timeline	Assessed Risk
PowerShell Execution by non-privileged users	Users from the front office were allowed to execute PowerShell which led to the initial callback from the attacker payload.	Disable PowerShell with Group Policy or limit PowerShell execution to authorized privilege accounts.	XXXX	Jun 2024: Assess the impact of disabling PowerShell in UAT environment. Aug 2024: Roll out the GPO to a small group of production users. Nov 2024: Roll out to all production users.	Dec 2024	High

9 References

9.1 Relevant Frameworks

CBEST Framework

- From the Bank of England and Prudential Regulation Authority
- CBEST implementation guide
- CBEST services assessment guide
- Understanding cyber threat intelligence operations
- Cyber resilience questionnaire

TIBER-EU Framework

Threat Intelligence-based Ethical Red Teaming

- From the European Central Bank

TIBER-NL Framework

Threat Intelligence-based Ethical Red Teaming

- From the Dutch Financial Stability Committee
- Derivative of TIBER-EU

iCAST, part of HKMA's Cyber Resilience Assessment Framework

Intelligence-led Cyber Attack Simulation Testing

- From the Hong Kong Monetary Authority

CORIE Framework

Cyber Operational Resilience Intelligence-led Exercises

- From the Council of Financial Regulators, Australia

F.E.E.R. Framework

Financial Entities Ethical Red Teaming Framework

- From the Saudi Arabia Monetary Authority

G7FE-TLPT

G-7 Fundamental Elements for Threat-Led Penetration Testing

- From the G-7

9.2 Relevant Guides

The Association of Banks in Singapore:

- https://www.abs.org.sg/industry-guidelines/cyber-security
- ABS Penetration Testing Guidelines

Cyber Security Agency of Singapore:

- https://www.csa.gov.sg/legislation/supplementary-references
- Guide to Cyber Threat Modelling

9.3 Additional references

ATT&CK (Adversarial Tactics, Techniques & Common Knowledge):

- https://attack.mitre.org/wiki/Main_Page

10 Glossary

ABS	Association of Banks in Singapore
BAU	Business-As-Usual
C2, C&C	Command and Control infrastructure
CBEST	Cyber resilience program – Bank of England
CF	Critical Function
CISO	Chief Information Security Officer
C-RAF	Cyber Resilience Assessment Framework – HKMA in HK
DNS	Domain Name Service
HTTP	Hypertext Transfer Protocol
IAM	Identity and Access Management
iCAST	Intelligence-led Cyber Attack Simulation Testing
ICMP	Internet Control Message Protocol
IT	Information Technology
RT	Red Team
SCCS	Standing Committee for Cyber Security
TIBER	Threat Intelligence Based Ethical Red-teaming – DNB in NL or ECB in EU
TCP	Transmission Control Protocol
TTP	Tactics, Techniques and Procedures used by attackers

11 Acknowledgements

- ABS Standing Committee on Cyber Security (SCCS)
  - Chairmanship: Standard Chartered Bank

- SCCS Working Group (AASE)
  - Project Lead: Citibank, UOB
  - Members
    - BAML
    - BNP Paribas
    - Citibank
    - Credit Suisse
    - DBS
    - HSBC
    - JP Morgan Chase
    - MAS
    - Morgan Stanley
    - OCBC
    - UOB