ART Red Team Guide

for the financial sector

DeNederlandscheBank

EUROSYSTEEM

Contents

1 About this guide

This red team guide provides guidance to the control team (CT) and control team lead (CTL) on the steps and deliverables in the red team (RT) phase and the procured red team provider (RTP) for drafting the core deliverables, being the Red Team Test Plan (RTTP) and the Red Team Test Report (RTTR). For a list of abbreviations, see annex A.

1.1 Purpose of this guide

The purpose of this guide is to provide the relevant stakeholders with information on the requirements[1] for the content and format of the main deliverables for this testing phase, being a RTTP and RTTR. It also aims at providing guidance on important aspects to be considered while drafting those documents.

This guide is part of the ART framework as published by De Nederlandsche Bank (DNB) on ART-NL | De Nederlandsche Bank | De Nederlandsche Bank (dnb.nl). For enquiries about ART, please contact the DNB Test Cyber Team (TCT) at tct@dnb.nl.

1.2 Target audience

This guidance document is primarily intended for the CT(L) and the RTP conducting the RT phase. In addition to these primary users, it may also provide the other stakeholders of an ART test with useful information.

1.3 Legal disclaimer

This document is intended for institutions within the scope of an ART test. Nothing in this guide should be construed as legal or professional advice.

This guide is an underlying document of the ART-framework. For information on copyrights and creative commons, please refer to section 1.3 of the ART framework.

1.4 Role of the TCT, minimum requirements and attestation

The RT phase is an essential and mandatory part of each ART test. It is strongly connected to the threat intelligence (TI) module. The Test Cyber Team (TCT) consists of a lead test manager (TM), a backup test manager and TI resources. When the institution selected the gold team module, the TCT will be complemented with GT resources.

1 This document also includes operational ART guidance based on best practices, knowledge and experience from numerous previous tests.

To ensure the quality of the test meets the ART standards, the TM will be present throughout the RT phase (and other modules) to ensure the RT is prepared, conducted and evaluated following the requirements as presented in this guide. At critical moments in the module, the TM provides approval on certain deliverables, such as the RTTP, the out plan, and monitors and controls each step during RT together with the CTL and RTP to conduct the test in a controlled way. The TM will work in close collaboration and has short communication lines with the CT and RTP.

Next to the quality assurance (QA) role, the TM is a neutral sparring and guiding partner for the CTL who holds the ultimate responsibility for the ART test within the institution, and for the RTP.

If the test has been carried out in accordance with the requirements of the ART framework, the TM will provide the institution on behalf of the TCT with a DNB attestation document concluding the test.

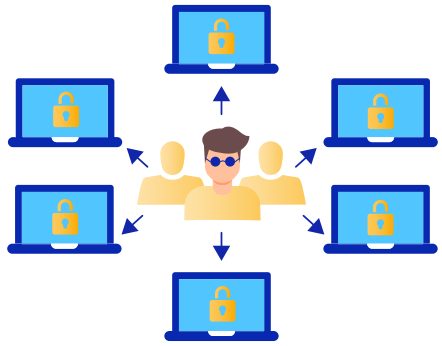

2 Introduction

2.1 What is the RT module in ART?

The RT module is a mandatory component of the ART framework.
It represents the active attack simulation phase, where a Red Team mimics real-world cyber adversaries to test an organisation's resilience under realistic conditions.

2.2 What is the purpose of RT?

RT is a security test designed to simulate real-world cyber-attacks on organisation's critical assets to identify vulnerabilities and improve defences. Its main purpose is to go beyond traditional security tests by thinking and acting like an actual and relevant threat actor. RT strengthens the organisation's overall cyber resilience in the end.

2.3 What are the goals of RT?

- **Gaining insights into cyber readiness**
  RT testing provides insight into the extent to which the organisation is prepared for a real cyber-attack. It provides measurable improvements in the security posture.

- **Identifying weaknesses and gaps within the security organisation**
  RT aims to uncover vulnerabilities and toxic combinations across IT (production) systems, networks, (business) applications and processes, physical security, and human factors. Observations, classified risks and recommendations help organisations to prioritise remediation efforts.

- **Validate Security Controls and Assumptions**
  It challenges assumptions about security posture and validates whether implemented controls actually work under realistic attack scenarios.

- **Testing detection and response capabilities**
  RT measures how effectively the BT can detect and respond to advanced attacks under realistic conditions. This includes stealth operations and multi-vector attacks.

- **Enhancing Crisis Management**
  RT observations and threat actor behaviour can be input for a cyber crisis management exercise. By conducting the optional Gold Teaming module, observations and threat actor behaviour validates realistic strategic decision-making and communication during a simulated cyber crisis, involving senior leadership and cross-functional teams.

2.4 Who is RT for?

Red teaming is generally intended for organisations that want to test and strengthen their security posture against realistic threats. Organisations who typically benefit from it:

- **Large enterprises and corporations**
  They have complex infrastructures and high-value assets, making them prime targets for advanced attacks.

- **Third Party Service Providers for the financial sector**
  These organisations face nation-state level threats and require assurance that their systems can withstand advanced cyber attacks.

- **Critical infrastructure operators**
  Sectors like energy, transportation, healthcare, and finance rely on uninterrupted operations.

- **Organisations with mature security programs**
  RT is most effective when basic security hygiene is already in place. It's not a starting point. It is for those who want to go beyond compliance and traditional security testing.

- **Any organisation concerned about advanced threat actors**
  If the risk profile includes targeted attacks, espionage, or insider threats, RT provides realistic insights.

3 Core principles of red teaming

This chapter describes the core principles of red teaming in an ART test. These core principles provide the reader with a clear understanding of the key notions, ideas and concepts are essential for conducting the RT module.

3.1 Safety and a safe learning environment

For all ART modules, the most important rule is that safety and a safe learning environment come first. Depending on the selected variant, the RT module may involve elements of stress and rapid-fire decision-making, which can challenge the resilience of your staff and team. The CT and TCT must have sufficient confidence in safety, capacity and expertise throughout the RT module. When in doubt, the CT and the TM will discuss how to enable this safe and capable learning environment.

3.2 RT is a strategic matter

RT tests and trains the BT function of the institution that is undergoing an ART test. RT is inherently a strategic matter and should therefore at least involve a C-level staff member of the institution, with an optional back-up. Back-up is advised to maintain awareness at board level should one of the C-level staff members involved leave the organisation.

In all variants of the RT module, tactical and/or operation teams such as (major) incident management teams, Computer Emergency Response Teams (CERTs) and forensic research teams may also be involved to train and exercise cooperation between teams from various levels of the (external) organisation.

3.3 Threat intelligence-based

Except for scenario X, attack scenario simulations are always based on threat intelligence. Realistic threat actor details and behaviour must also be considered and included in the development of the optional crisis management scenario for the GT.

3.4 Executed by professionals with verifiable experience

Developing TI-based attack scenarios followed by executing and evaluating selected scenarios by a RT requires a combination of specific knowledge and skills. In an ART exercise, TI must be executed by internal or external professionals, and RT by external red teams. These professionals should have verifiable experience and must be capable of working closely together with the CT an TCT. Procurement guidelines support the CT to select a security provider.

3.5 Mandate

The CTL engages the control team during the preparation phase. C-level participation and mandate (awareness and approval) is required for conducting an ART test. This applies to testing on live production systems as well as to the preparation and closure phases. Observations and recommendations resulting from the RT require C-level mandate to be mitigated or resolved. The TM is mandated to approve moving from one phase of testing to the next. If the TM is of the opinion that the complete test was conducted in accordance with the ART quality requirements, the tested institution receives a DNB attestation.

3.6 Based on learning goals and sufficient planning

The learning goals defined in the preparation phase run as a common thread throughout the RT phase. Learning goals depend on an institution's current level of maturity and resilience, and the extent to which the CT can challenge it. Just as for the TI, the RT is deliberately planned: the more complex the RT variant, the more thorough the RT plan and planning. For example, if the assumed breach variant is selected and the exercise set-up contains an element of surprise (unannounced exercise) for participants, sufficient containment measures should be in place.

3.7 Confidentiality

Confidentiality is of vital importance during all ART modules, including RT. The BT must be tested in a realistic way without being aware of the planned attack scenario(s). All ART test activities are highly confidential and only known by the CT. The ART test may not be disclosed by an unplanned or uncontrolled action caused by the CT, RT or TCT. The code name of the ART test also applies to the RT module and must remain confidential for the sake of realism.

4 Red Team variants

Three RT variants are available and can be combined to define the full size of the RT module. The CT can choose to only execute one scenario, but it is also possible to execute multiple scenarios. However, when selecting only one scenario to be executed, the assumed compromise variant is the minimum variant.

4.1 TI-based scenario, assumed compromise

The active RT phase of an ART test consists of at least one TI-based scenario, starting with an assumed compromise. In this case the 'in phase' is represented by a TI based leg-up which should be provided by the CT. This decision can be based on learning goals, on time constraints or business knowledge constraints.

As will be explained in section 7.2 on leg-ups, it must be considered that because of the secrecy of the test, the CTL may need more time than expected for arranging the assumed compromise and provisioning of the leg-up.

The RT can be asked to demonstrate that the provided leg-up could have been executed retroactively, based on the knowledge gained during the through and out phase and/or if there was no time limit. This can either bedemonstrated during the RT phase (in this case the order would be through, out, in) or during the PT phase.

4.2 TI-based scenario, end-to-end simulation

To make a scenario as realistic as possible, an institution can choose to incorporate all the steps of the TI-based scenario (including the 'in phase'). This means that threat actor behaviour is fully simulated during the scenario, following the same sequence of actions that would occur in a real-life situation. This method of RT generally requires more time and resources but provides the most comprehensive view of the institution's cyber resilience, covering all relevant areas, such as people, processes and technology or physical properties.

4.3 Scenario X

A scenario X variant can only be included in addition to the other variants. It cannot be used as a substitute for an assumed compromise or end-to-end scenario. The goal of the scenario X variant is to simulate attacks that may be expected in the future or that are solely based on specific learning goals of the institution. This scenario can focus on innovative techniques and emerging tactics.

When executing scenario X, the RTP can use observations from scenarios executed earlier in the test, meaning it can be defined during the RT phase. In that case, a finalised scenario must receive timely approval during the RT phase, depending on the planning of the overall RT-module. However, it can also be defined upfront together with the definition of the TI based scenarios, when based on a specific learning goal of the institution. The goal of scenario X is to target a CIF, often by using a highly creative approach.

5 Description of the red teaming steps

Following the approval of the TI report by the CT and the TCT, the RT activities move into focus. During the RT phase, the RTP plans and executes the RT phase based on the selected scenario(s) for the target systems and services that underpin the selected CIFs in scope.

The RT phase consists of three separate process steps:

- Step 1: Creating the Red Team Test Plan (RTTP)
- Step 2: Red team testing
- Step 3: Compiling the Red Team Test Report (RTTR)

The TI report forms the basis for the RTTP. When the RTTP is in its final stage, the stakeholders (CT, RT and TCT) hold a 'go/no go RTTP meeting' to approve the plan. After approval of the RTTP, the active testing process step starts. During active testing, the RTP aims to reach all the flags, as defined in the RTTP.

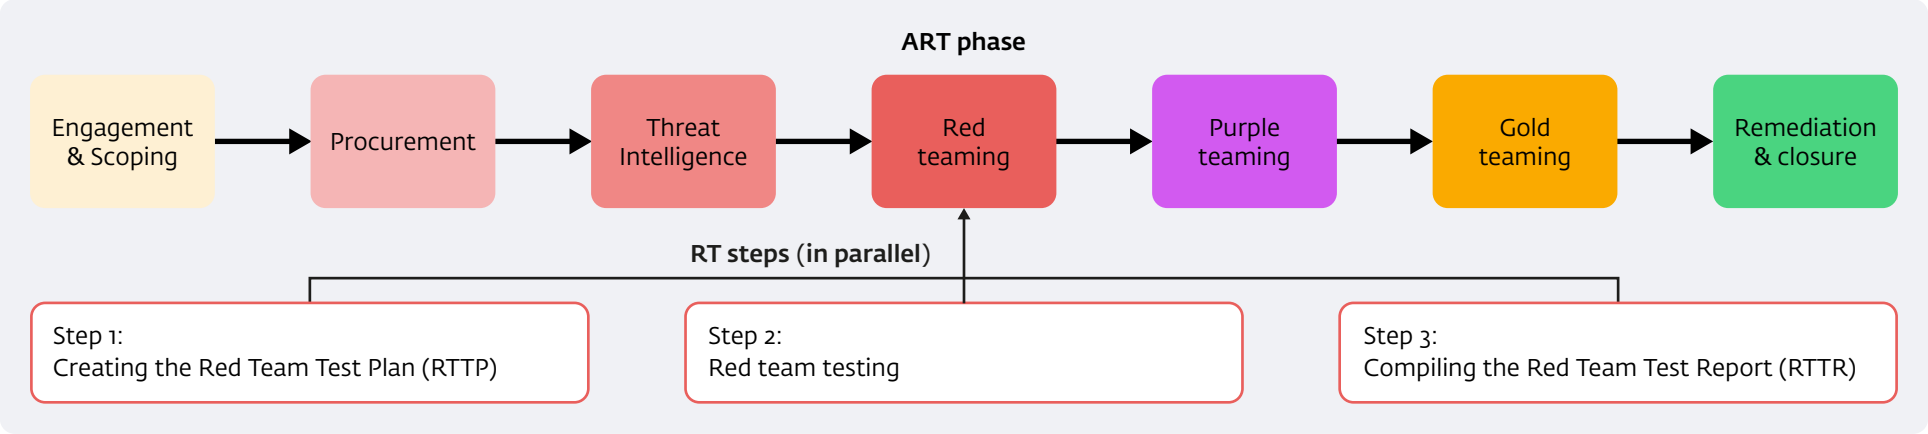

Step 1: Creating the Red Team Test Plan (RTTP)

The RTTP must be drafted before RT testing starts. In this process step, the RTP develops and integrates the attack scenarios into a RTTP, leveraging on the selected scenario(s) from the TI report.

The RTTP must include a description of the out phase. The out plan is not a formal part of the ART framework but can be added as an annex to the current RTTP, to be set up and finalised during RT testing. This must be done after the through phase but before starting the out phase.

Required content and considerations for writing the RTTP and finalising the out plan will be explained in Chapters 6 and 7.

Step 2: Red team testing

Once the RTTP is approved by the CT and TM, the RTP should initiate the active execution of the test. Any changes to the RTTP after its approval must be approved by the CT and TM. The attack scenarios are not a prescriptive playbook which must be followed precisely during the test. If any obstacles occur, the RTP should show its creativity (as a selected threat actor would do) and develop alternative ways to reach the test objective or flag.

A minimum of 6-12 weeks must be allocated to active RT testing, to allow the RT to conduct a realistic and comprehensive test in which all attack phases are executed, and the flags can be reached. Within this time frame, if there is more than one attack scenario, the attack scenarios can be executed in parallel or in sequence. If scenarios are executed in parallel, the RTP still needs to complete the scenarios' in, through and out phases, depending on the RT variant or combination of RT variants selected (see Chapter 6).

The RTTP must include a comprehensive description of the out phase. Before the start of this phase, the RTP must determine if the out phase – as described in the RTTP – is still aligned with the current planned execution of the scenario. If not, the RTP must specify how it will approach the new out phase. This does not have to be recorded in a formal document, but the RTP must prove that it is in control during the out phase and the CTL must have a record of the discussion and decisions taken in the discussion of the out actions.

Regardless of whether it is aligned with the attack, the out phase must be discussed with the TM and the CT before it is executed.

Limited purple teaming (LPT)

During active RT testing and before PT, there might be circumstances that require disclosing (part of) the test activities to the BT. In that case, the limited purple teaming (LPT) is started, which means involving the BT in the RT phase. LPT is a measure to ensure that the red team testing still provides as much added value as possible, but it can never be carried out instead of the PT module. The CT, TCT and RTP decide whether to start LPT. The CT will then inform the institution's BT about a part of the test, without disclosing the specifics of that scenario or the execution of other scenario(s).

The planned attack scenario(s) move on in a safe or meaningful way that has been aligned with the BT. The BT can also be asked to share their testing needs. It is important to stress that the BT should continue operations as normal. They should be instructed not to put additional efforts in trying to find suspicious activities that could be performed by the RTP. The BT can possibly even cooperate with the RTP. For example, when detecting a security event or incident for example, the BT could check whether the indicator of compromise (IoC) is from the RTP or from a real threat actor. Sometimes this is called 'catch & release'.

Back to overview red teaming steps

After LPT and the closure of the active red teaming, the formal PT phase will start anyway to give the BT full insight into the activities conducted during RT testing.

Requirements for LPT

It is not possible to provide an exhaustive list of circumstances that could result in LPT. However, one of the main criteria should be that an event – outside the control of the CT, RTP or TM – prevents the continuation of the execution of a scenario in the RT phase without compromising the secrecy and/or the security of the scenario.

After consulting with the TM and the RTP, the CT can decide to start LPT:

- after detection or other circumstances forcing the test to be disclosed
- due to other circumstances e.g. to avoid:
  - risk of impact on data
  - risk of impact on contractual agreements with suppliers or third parties
  - damage to assets
  - disruption to critical important functions (CIFs), services or operations of the financial institution itself, its third-party service providers, or ICT intragroup services providers
  - disruptions to its counterparts or to the financial sector.
- if the continuation of the test is not otherwise possible.

Starting LPT during the RT phase requires the following steps:

- the CT consults with the RTP and TM to formally start LPT
- the CT details the specific scope and objectives in close cooperation with the RTP, TM and BT
- the CT asks the BT about their needs and the scenarios to be considered before starting LPT
- the CT consults with the TCT, RTP and TIP where necessary to adapt existing scenarios or implement alternative (TI-based) scenarios to maximise the learning outcomes for the institution
- a communication channel must be set up with the BT, to ensure short lines between BT, CT and RTP.

Circumstances leading to LPT

Before proposing to start LPT, the CT must carefully consider alternative options for moving forward after disclosing (a part of) the test. Advised by the TM, the CTL, RTP and involved stakeholders must then decide whether starting LPT is the most effective way to maximise the learning outcomes.

Examples of potential circumstances that may lead to LPT are described below.

- If the BT detects the RTP and the secrecy of the scenario or complete test is permanently compromised. Note that it is possible that the BT may detect some RTP actions during a test. However, this alone does not necessarily mean that LPT is the right way forward. It may be possible to continue the test in its original manner using a cover story (e.g. a local penetration test) to explain certain detections to the BT or to only introduce LPT for the disclosed attack scenario(s). In addition, it may be possible that the BT detects part of a test. The CT can then instruct a freeze on RTP activities to allow the elevated threat level to subside. In such cases, it is crucial to have alternative approaches and techniques at hand, as it is a common mistake to pause only to reuse the same attack vectors that have already been detected.

- LPT can also be started as a result of unforeseen situations, where there is a high degree of risk that the simulated attack could lead to business impact or disruption to CIFs. In such cases, it is advisable to discontinue testing and to start LPT for these systems instead. This would enable the BT, once informed, to take timely action to prevent and minimise any impact, or advise on safe, alternative TTPs.

- If a real cyberattack occurs during an ART test, and the BT must fully shift its focus to prevention and containment of disruptions. This may result in the ART test being disclosed to ensure the BT can distinguish between the test activities and the malicious attack. Among other possibilities, the test may be postponed to a later date.

Back to overview red teaming steps

- If it is likely or evident that the uninformed BT's response to contain the detected emulated attack will have a disruptive impact on CIFs. This potential overreaction might be appropriate in the event of a real attack, but not in the context of an ART test, given that the ethical RTP will never deliberately cause real harm. If the BT is not aware of the ART test, they have no way of knowing if their response is adequate.

- To prevent situations that can lead to the BT deviating from normal response procedures when suspecting the detected attack in a test. This would both reduce the realism of the test and hamper its learning outcomes.

- When the CT is unable to stop escalation by the BT, and the BT is about to or even already has involved external parties such as the police, intelligence services, government authorities, industry bodies or financial institutions. In real life, this could be the appropriate course of action due to the perceived severity of the incident. However, involving these parties will put an unnecessary strain on those authorities and could have a severe impact on current and future testing activities. That is not the intention of performing an ART test. So, in that case, the test should be halted immediately and (partially) disclosed, and LPT should be considered.

Step 3: Compiling the Red Team Test Report (RTTR)

The CT, RTP and TM should agree on when the active RT phase is to end. Following the end of the active RT phase, the CTL will inform the BT that a test was conducted. The RTP will start writing the RTTR and deliver it within four weeks after completion of the RT phase. Required content and considerations for writing the RTTR will be explained in Chapters 8 and 9.

Back to overview red teaming steps

6 Required content of the RTTP

The RTTP for the active RT phase describes the attack steps to be performed, the techniques to be used, when and how to provide leg-ups if milestones cannot be reached, as well as what flags to aim for at the target systems.

The RTTP may be drafted in any preferred format, provided that all required information is included. Drafting a RTTP is required for each RT variant selected from the ART framework.

The RTTP must clearly define which RT variant is selected to determine the composition of the RT module. Selecting scenario X alone is not allowed, because that would enable a test that is not TI-based. When the scenarios include a sufficiently wide range of actors, TTPs and objectives, executing a second or third scenario can yield valuable additional learnings.

On top of that, the RTTP must at least include information on:

- The teams involved in the test (which could be the CT, TCT, TI, RT and GT) specifying the team members and their roles in the test.
- a description of the scenarios selected, including:
- the simulated threat actor
- their intent, motivation and goals
  - the targeted CIF(s) and the supporting systems and services and potential flags
  - the targeted confidentiality, integrity and availability aspects
  - the tactics, techniques and procedures (TTP) to be used for the (in), through and out phase, based on the MITRE ATT&CK framework.
- potential leg-ups to be provided by the CT, including deadlines for their provision and potential usage
- the scheduling of RT activities, including a time planning for the execution of each scenario, divided into the in, through and out phases
- a description of the activities for the out phase. Potential additional detailing must be prepared before initiating these out actions on the objectives. This detailing can be delivered in the RTTP or after completion of the in and through phase
- risk management measures to be followed up on by the CT and RT. These measures can be included in a separate risk register to be updated on a regular basis when new risks arise which must be treated in a certain way
- communication channels and procedures to be used
- if applicable, ethical boundaries for social engineering, and how the privacy of involved parties is being safeguarded
- particularities of the financial entities' infrastructure to be considered during testing
- if any, additional information or other resources necessary to the red team for executing the scenarios.

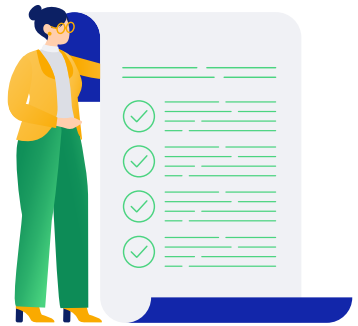

7 Considerations when drafting the RTTP

7.1 Scenario description

The planned attack scenarios based on the selected threat scenarios from the TI phase form the core element of the RTTP. Based on the finalised flags, the RTP should develop appropriate attack scenarios for each flag. The scenarios are written from the attacker's point of view in a creative process. Each scenario should be written as a film script and set out how specific targets can be reached (i.e. the flags to be captured). Scenarios can be executed in sequence or in parallel but must be independent from each other. The RTP should list creative options for each attack phase, considering various TTPs used by advanced attackers to anticipate changing circumstances, or if the initial option does not work. Except for scenario X, TTPs must be TI-based and in line with the MITRE ATT&CK framework.

The TTPs do not simply reflect real-life attacks seen in the past but combine the techniques of the various relevant threat actors. The RTP may include more than one threat actor in some scenarios, since some of them share similar motivations and/or objectives. This makes it possible to include additional test activities and gives the RTP more flexibility in selecting the attack methods based on a broader set of TTPs. Threat actors must be aligned with the TI report, except for the scenario X variant.

For each RT variant, the RTP should clearly describe the following:

- **Simulated threat actors**
  Based on the collected threat intelligence in the TI report, elaborate in further detail and more precisely what the motives, intent and goals of the threat actors are, how they would seek to target the specified CIFs, which of the CIA triad (Confidentiality, Integrity, Availability) they would seek to compromise and how they would focus their efforts on capturing the final flags.

- **Targeted critical functions, underlying systems and services, and flags**
  The relevance of the CIFs being tested, and how the flag/objective relates to the critical or important function and its underlying systems and services.

- **Tactics, techniques and procedures (TTPs)**
  What tactics, techniques and procedures the specific threat actor(s) would use to capture the specific flags. These TTPs should be set out in line with the MITRE ATT&CK Framework (https://attack.mitre.org/).
  In some cases, the implementation of the framework may also include using TTPs trying to breach the physical security of the institution to gain access to the network or plant a device.

7.2 Potential leg-ups

If the RTP is unable to progress to the next phase of the agreed attack scenario – whether due to time limitations, insufficient knowledge of business processes, or because the institution has effectively defended itself – a leg-up may be required. Although the RTP simulates the behavior of determined threat actors, it must always operate in an ethical and responsible manner. When these boundaries prevent the scenario from advancing, the RTP can request a leg-up, provided that the realism of the test is maintained. A leg-up is not intended as a shortcut; rather, it compensates for the advantages a real attacker would typically possess, such as greater time, resources, or fewer operational constraints.

Leg-ups generally fall into one of three categories:

- Information that, under normal circumstances, the RTP could have discovered with additional time;
- Support with a task that the RTP is unable to perform within the (ethical) constraints of the exercise;
- Access that a real threat actor could realistically have obtained.

Any leg-up should always be limited to the minimum assistance required for the scenario to continue.

For each potential leg-up, the RTP must clearly define its scope, specify who is authorized to approve it, outline the deadline for approval, and describe the procedure for submitting the request.

Potential leg-ups must be documented and incorporated into the RTTP. Therefore, before the test begins, the RTP should discuss the circumstances under which leg-ups may be necessary, as unplanned leg-ups during execution can significantly influence the test timeline. Addressing this early enables a smoother and more efficient testing process. Following these discussions, the RTP must establish a clear request process and identify any scenario-specific leg-ups that may be required. This preparation allows the CT to make the necessary arrangements in advance, ensuring the test can proceed without disruption. Managing and delivering leg-ups during a covert test can be challenging for the CTL, given the need to maintain secrecy within the organisation. Once approval is granted by both the CT and TM, a leg-up may be employed.

Depending on the nature of the test and its progress in real time, the RTP may deviate from the TTPs to remain agile and flexible. Therefore, it is not always possible to document leg-ups and deviations in advance. The RTP should retain a degree of flexibility to allow improvisation during the test, even if certain TTPs are not included in the RTTP. This may reduce the need to use leg-ups, relieve demand on the CTL and potentially improve the learning experience.

The need for and status of leg-ups is usually discussed and decided during the weekly updates.

7.3 Scheduling red-teaming activities

Given the complexity of an ART test, the RTP and CT must have a well-defined plan in place to ensure the test can be completed successfully. The RTTP should be prepared as early as possible and must describe any relationships between the different scenarios and the dependencies on leg-ups to visualise the orchestration of the test. A structured and prescriptive timeline will increase the RTP's opportunities to capture all of the flags. If the RTP is unable to get to the next phase or flag within the test time, they can request a leg-up.

In this section of the RTTP, the RTP should also provide a general timeline for the execution of the scenarios. At a minimum, this timeline should be divided into the 'in, through and out' phases that a tester progresses through during the testing phase. It should include all the key milestones, dates of meetings, activities and deliverables. It can be represented in the form of an illustration to provide a clear and accessible overview. The minimum time span of the active red team test is 6 - 12 weeks. This excludes the preparation of the RTTP, delays resulting from changes due to detection or incidents, freeze periods, further development of scenario X (when learnings from TTPs that have been executed earlier in the test are used), preparation of the out plan, planned pauses due to holidays and preparation of the RTTR.

Quality is more important than speed, and RT members need time and space to stay creative while executing the selected TTPs during the RT phase. Depending on the learning goals (for example if the CT wants to know how long it might take to be detected once the RT has obtained a foothold) the CT may ask the provider to extend the duration of test without the RT increasing the planned work hours.

It is important that the RTP can map timelines to flags and end goals, although the timelines may change during the test due to the inherent unpredictability of an ART test. Setting out more structured and prescriptive timelines ensures that the RTP can attempt to capture all the prescribed flags. As the timelines are likely to change during the test, the CT and RTP

can align expectations and ensure the right amount of time is spent on respective activities, ensuring all important parts of the institution are tested, as is expected and required in an ART test. This more prescriptive timeline will also allow the CT to highlight any issues that may interfere with test activities.

7.4 The out phase

The RTTP must include a (detailed) description of the out phase. After the through phase, the out activities must be detailed, set up and aligned before starting the out phase. This out plan is not a formal part of the framework but can be added as an annex to the current RTTP. The RTP – together with the CT and TM – must determine if the out phase is still aligned with the planned execution of the scenario as described in the RTTP. In any case the RTP must specify how they will approach the out phase. They should prove that the activities are TI-based and that they are in control during each step of the out phase. The activities in the out phase must always be discussed with and agreed by the TM and the CT before they are executed.

The time required for setting up, aligning and approving the out activities should be considered and be part of the time planning. The out plan describes any additional measures to stay in control during the out phase, such as security considerations that are in place during exfiltration of the data from the institution, installing (fake) ransomware or making payments.

A subject matter expert from the institution could join the CT to discuss any unforeseen risks and mitigation measures. A ready-made communication plan for disclosing the test on each level within the organisation can be valuable.

It is advised to ask the RTP to present the out activities to the TM and CT including the C-level member before starting the out phase. This creates awareness at the board level, not only concerning recent achievements in previous phases, but also concerning the actions to be taken regarding targeted objectives (CIFs).

7.5 Risk management

The RT phase always involves risks, due to the critical role of the targeted systems, people and processes and physical properties. The RTTP should therefore clearly describe, either in the risk management section or in the separate risk register, how the RTP and CT will take appropriate measures before, during and after the RT phase. Risk management is an active and continuous process. Risks must be considered before each step during the RT phase. The risk management section or risk register should therefore be updated on a regular base, and potential risks and measures must be discussed during the weekly calls. A secure messaging app (see Chapter 7.6) can be used when an unexpected risk or issue arises. The RTP should liaise with the CT to confirm the intended risk management approach.

The initial risk assessment by the RTP and CT should cover the following risks (non-exhaustive):

- Risks related to reputational damage if the confidentiality of the test is breached or in case of unethical conduct
- Risks related to crisis and incident escalation
- Risks related to operational RT
- Risks related to operational defence
- Risks related to clean-up after completion of the test
- Risks related non-compliance with the framework.

If external providers are hired for the test, the institution must ensure that there is mutual agreement on the following aspects (non-exhaustive): the ART service procurement guidelines are met, the scope of the test, boundaries, timing and availability of the providers, contracts, actions to be taken and liability (including insurance, if applicable). Any changes to the RTP composition must be communicated to the TM and CTL in a timely manner and included in an updated version of the RTTP.

The RTP should share all infrastructure, domain names, hashes, emails (i.e. indicators of compromise) used by the RTP with the CT before the RT phase starts – to the extent possible at this stage. This will allow the CT to distinguish between the ART test and actual attacks and take the appropriate steps to manage an actual cyber attack. The RTP should share the details with the CT on the risk management approach to be taken on securing the RT infrastructure set-up before start of the test. This includes guardrails to secure RTP's infrastructure and its location to make sure the RT infrastructure or any other assets used as part of the RT activity are protected. The RTTP should describe ethical boundaries for social engineering, as well as how the privacy of the parties involved is being safeguarded.

In addition, the TM's involvement in the ART test ensures that the test follows all required process steps, resulting in the agreed test scope, scenario, planning and process as described in the ART framework and associated guides. The minimum requirements for cybersecurity service providers are set out in the ART service procurement guide.

Risks are also mitigated by sound planning, informing only a select group of people in higher management about the test and its scope, maintaining an up-to-date risk register during the entire test and a clear definition of the scope and predefined escalation procedures. It is important to note that the institution remains in control of, and responsible for, the test and the risk management. At any time, the CT can (temporarily) suspend the test if concerns are raised about damage (or potential damage) to a system or business process. Trusted contacts within the CT positioned at the top of the security incident escalation chain can help to prevent miscommunication and share knowledge about a possible ART test detection.

Ethical boundaries

The RTTP should mimic the (previously seen, current and potential future) actions of a real threat actor. Criminals do not adhere to ethical rules, and an ART test should use the same kind of "creative thinking" criminals would use – up to an ethically acceptable point – to make the test as realistic as possible. This ethically acceptable point will be different for every institution, and it is up to the CT to define this.

There are certain types of behaviour that are strictly forbidden in ART:

- Unauthorised destruction of equipment
- Unauthorised modification of data/programmes
- Unauthorised jeopardising of continuity of critical services
- Extorting, kidnapping, threatening or bribing employees
- The use of mailing lists, names, logos or otherwise identifiable information of real people or companies, without explicit approval of the holder of those attributes.

Code name

To protect sensitive information, a code name must be used for the test. The CTL may choose a code name for their test. If they prefer, the TCT can provide a code name. This code name should be used in all documentation related to the ART test, in document titles as well as in the documents themselves. Elements where code names cannot be used (such as URLs and screenshots) are exempt and may contain the full name of the institution. The code name will be used in all communications, such as meeting invites and documents exchanged between the parties involved in the test.
In addition to the code name assigned by the TM and/or CT, providers and/or the institution are free to use their own code names for internal communication.

Escalation and suspending the test

The test may reach a level of escalation that causes the BT to inform relevant authorities, such as the police, intelligence agencies or data protection agencies. The CT must always try to prevent this from happening, as external authorities should not be burdened by an ART test. In case the CT is informed of an active escalation to outside authorities, the test must immediately be suspended so that measures can be taken to prevent these authorities from getting involved.

Personal identifiable information

It is up to the institution to set up contractual agreements with the RTP to ensure protection of their employees' privacy. Under no circumstances may privacy-related information be included in test reports.

Artifacts

In the RTTR or in an annex to this report, the RTP must describe any artifacts that could plausibly remain on the institution's systems after testing, such as toolsets that remain on compromised hosts/systems. The institution should remove these artifacts, as they may pose a risk to the systems or interfere with any future incident investigations or security assessments. Typically, these artifacts are described using filenames, paths, hashes, hostnames, IPs, email addresses, email subjects, email domains and web domains.

7.6 Communication channels

The RTP should indicate in the RTTP how the stakeholders (i.e. CT, TMs, TIP and GTP) will be updated during the testing process. All communication must be conducted via secured channels, for example, end-to-end encrypted chat and email. Realtime and active communication between CT, TIP, RTP and TCT can be done via a secure messaging app (e.g. Signal). During active communications in the test, the participants should always refer to the test by its code name rather than the institution's name, to minimise risk in case of information leaks. In addition, all communication between the TIP, RTP, CT, GTP and TM should be safely stored to ensure evidence of approvals and communications conducted via secure channels.

Throughout the active RT phase, the RTP reports on the progress made during status meetings, which are held in-person or online at least on a weekly basis and involve the CT, TM, TIP and the GTP as optional participant. In the weekly status meetings, the RTP should report on planning, progress, activities executed (what was successful and what wasn't), next steps, potential risks, expected leg-ups and the status of planned actions. The RTP should share the documented weekly information just before the weekly update calls with all stakeholders so that they can prepare for the meetings.

Furthermore, during the test, the RTP may need to communicate directly, immediately and urgently with the CT. For example, before moving further in case an issue arises.

8 Required content of the RTTR

The RTTR is created by the RTP after the RT phase and serves (together with the optional BT report) as input for the purple teaming (PT) phase, the Test Summary Report (TSR) and the Remediation Plan (RP).

The RTTR may be drafted in any preferred format, provided that all required information is included. Delivering a RTTR is required for each ART test.
It summarises the test and its results on a detailed technical level.
It includes all the information about all attack actions performed by the RTP, weaknesses identified and detections by the blue team (BT) as well as specific analyses of root cause indicators and recommendations for remediation and improvement. It therefore serves as a basis for all the activities conducted during the closure phase. The RTTR should also clearly set out how, when and under what circumstances the various stakeholders communicated with each other.

The RTTR should include information on at least all of the following:

- Information on the performed attack, including:
  - the targeted CIF and identified systems and services, processes and technologies supporting the CIF, as identified in the RTTP
  - a summary of each scenario (at least the minimum variant)
  - flags captured and not captured
  - attack paths (TTPs) followed successfully and unsuccessfully
  - TTPs used successfully and unsuccessfully
  - deviations from the RTTP, if any
  - leg-ups used, if any.
- All actions known to testers that were carried out by the BT to reconstruct the attack and to mitigate its effects.
- Discovered vulnerabilities and other observations, including:
  - vulnerability and other observation description including their criticality translated into classified risks
  - provisional root cause analysis of successful attacks
  - recommendations for remediation including indication of the remediation priority.
- The RTTR should clearly set out how, when and under what circumstances the various stakeholders communicated with each other.
- A summary of observations and recommendations for remediations which could be supplemented and refined after the PT sessions and used in the TSR.

9 Considerations when drafting the RTTR

9.1 Management summary

The RTP should draft a short summary in non-technical language for senior management and higher-level governance bodies (such as a Board of Directors), to:

- explain what critical or important functions and underlying systems were tested
- provide a high-level timeline of the test and an overview of the scenarios tested, including references to mimicked threat actors from the Threat Intelligence Report, and context of the successful and unsuccessful attack methods employed, including the use of leg-ups
- highlight the main observations (based on criticality), including strengths and challenges, and possible root cause indicators based on the attack methods used
- give insight into the main categories of recommendations to address the observations and possible root cause indicators
- note any significant observations and exceptions in the test.

9.2 Storyline

The RTP should draft an end-to-end storyline of the test for the selected scenarios, where relevant divided into in, through and out phases, outlining:

- The CIFs and underlying systems that were targeted, including the identified ICT systems, processes and technologies supporting the CIF as identified in the RTTP
- A summary of the TTPs that were utilised (both successful and unsuccessful) in line with the MITRE ATT&CK framework, including the flags and objectives reached and not reached (linked to the relevant aspect(s) of the CIA triad)
- An annex including a timeline with relevant logs and details for the BT to create an optional BT report can be added.

9.3 Observations

The RTP must document the observations identified in the testing process. Each observation must be categorised by criticality and complexity, and each observation must contain a clear description on how the CIF was compromised and the impact of the compromise (including the potential real impact if there were no limitations of a test). The observations should describe both technical and non-technical elements, if applicable.

9.4 Provisional root cause analysis

The RTP must use their experience and expert judgement to determine whether they can draw conclusions on root cause indicators of the observations outlined above. In order to do this the RTP needs to consider people, processes and technology holistically and not limit their view on the technological aspect alone.

Using these root cause indicators, the RTP must also extrapolate from their actions what more could have been done to advance the attack on the institution and what the possible impact could have been.

It should be noted that the provisional root cause analysis will be judgement-based and only preliminary. The aim of such a preliminary analysis is to provide the BT with a basis to reflect and to facilitate a robust discussion in the replay exercise. This section of the RTTR should be more analytical in nature and aims to facilitate the replay exercise being more forward thinking, rather than solely technical and retrospective.

9.5 Recommendations for remediations

The RTP should develop clear conclusions and identify specific recommendations for remediations, which can lead to future action. The RTP must extensively document the recommendations for remediations on the observations in the following manner:

- the prioritisation of recommendations for remediations must be commensurate to the observation it aims to address, and
- the recommendations for remediations must be adequately described to determine the institution's objective and to be able to implement the actions under each recommendation.
- the recommendation for remediations on the provisional root cause indicators should be drafted at a level that provides options and guidance to the institution undergoing the test, given that there may be alternative ways to address the provisional root cause indicator.

Annex: List of abbreviations

ART	advanced red teaming
BOD	board of directors, also referred to as executive board
BT	blue team
CIFs	critical or important functions
CMT	crisis management team
CT	control team
CTL	control team lead
GT	gold teaming
GTL	generic threat landscape
GTP	gold team provider
GTTP	gold team test plan
LPT	limited purple teaming
NDA	non-disclosure agreement
OSINT	open source intelligence
PT	purple teaming
RFP	request for proposal
RT	red teaming
RTP	red team provider
RTTP	red team test plan
RTTR	red team test report
SME	subject matter expert
SOC	security operations centre
SSD	scope specification document
TCT	test cyber team
TI	threat intelligence
TIBER	threat intelligence based ethical red teaming
TIP	threat intelligence provider
TIR	threat intelligence report
TM	test manager
TPSP	third party service provider
TTP	tactics, techniques and procedures

De Nederlandsche Bank N.V.
PO Box 98, 1000 AB Amsterdam
+31 (0) 20 524 91 11
dnb.nl/en

Follow us on:
Instagram
LinkedIn
X

DeNederlandscheBank

EUROSYSTEEM