# Advanced Red Teaming (ART) Framework

## for the financial sector

DeNederlandscheBank

EUROSYSTEEM

# Inhoud



ART website & documentatie

# 1 Introduction: start with ART

## 1.1 What is ART?

ART is a comprehensive framework that empowers a wide range of financial institutions to conduct advanced ethical red teaming tests driven by high-level threat intelligence. As an evolution of the TIBER framework, ART grants participating entities the freedom to select and customise various modules, ensuring that each test aligns with their specific needs and learning objectives. These modules address different facets of cyber resilience testing, encompassing physical intrusion, incident response, and network and application security.

By enabling entities to choose the modules most pertinent to their cybersecurity posture, maturity level and available resources, ART allows them to optimise their cybersecurity efforts and investments. This approach results in a tailored red teaming engagement that aligns precisely with their unique learning goals. ART's modular nature offers participating entities flexibility and value, reducing documentation requirements while upholding the rigorous cybersecurity testing standards associated with TIBER-EU. Upon successful completion of a test, and if the managing authority (DNB) concurs, it will be officially registered as an ART test.

## 1.2 Who is ART for?

The fundamentals of ethical hacking according to the ART framework are applicable to multiple institutions and sectors. The core principles, such as testing on live systems and secrecy of the test, are the same across all sectors. Detailed implementation, however, may vary depending on the nature of the sector. This specific ART framework is focused on the financial sector. Within this sector, four groups of potential participants can be identified:

1. Entities that have been actively enhancing their cybersecurity posture for several years and are committed to further improvement, but are not yet ready for a full TIBER test or are not subject to DORA/TLPT.
2. Entities that already perform a TIBER/TLPT test every two to three years, but are looking for a more frequent testing framework to bridge the gaps between TIBER/TLPT tests.
3. Entities that are ready for a TIBER/TLPT test, but that are not classified as "vital".
4. Non-financial entities that provide services, software and systems that are critical to the functioning and stability of the financial system.

## 1.3 Main differences with TIBER/TLTP

ART is an evolution of TIBER. Therefore, these two frameworks share a number of fundamental principles and a common objective: to enhance cyber resilience by learning from realistic hacking tests. Nonetheless, there are also notable differences between TIBER and ART. This chapter offers a high-level overview of the main similarities and differences between these two frameworks.

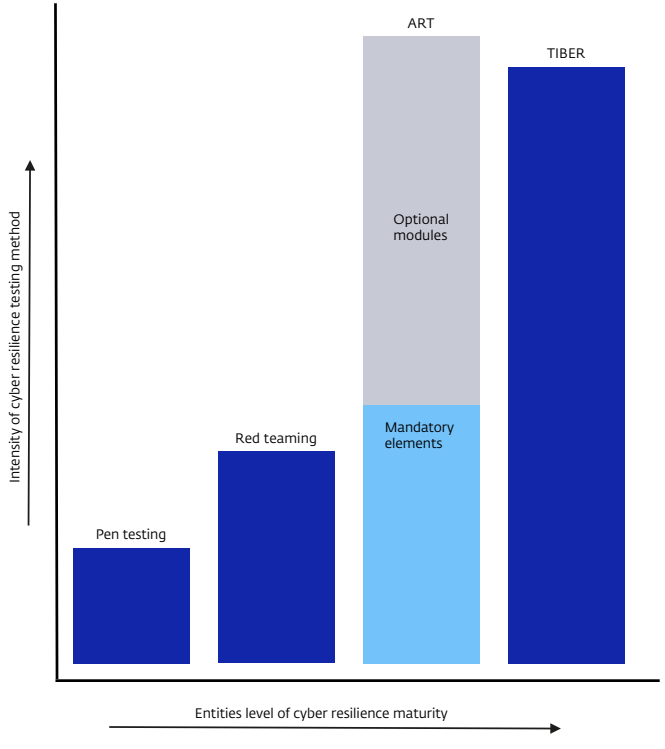

| Subject | TIBER/TLPT | ART |
|---|---|---|
| **General** | | |
| Participation obligation | Mandatory | Voluntary |
| Main participants | Critical FIs and third parties | Important FIs and third parties |
| **Duration (indicative)** | | |
| Total test duration | 9-12 months | 6-9 months |
| Threat intelligence | 6-8 weeks | 2-8 weeks (modular) |
| Red teaming | 10-12 weeks | 6-12 weeks (modular) |
| **Extent of the modules** | | |
| Threat intelligence | Full | Modular |
| Number of scenarios | Two plus a Scenario X | A minimum of one scenario |
| RT test | In-through-out fully | Assumed compromise possible |
| Testing on live systems | Yes | Yes |
| Purple teaming | Full | Standard + additional module |
| Gold teaming | Not yet mandatory | Optional (modular) |
| **Involved parties** | | |
| Test Cyber Team | Mandatory | Mandatory |
| Control team | Mandatory | Mandatory |
| Board engagement | Mandatory | Mandatory |
| Red team provider | Mandatory | Mandatory |
| Threat intelligence provider | Mandatory | Optional |
| Gold team provider | To be determined | Optional |

More information on the differences and similarities between TIBER and ART is provided in Annex 2.

## 1.4 Disclaimer and legal

An official ART test for the financial sector can only be conducted by financial entities and their providers under the guidance of the DNB TCT. The financial entity and the DNB TCT draw up a contract which specifics, among others, risks, responsivities and reimbursement for the duration of the test. A multi-test contract is possible.

# 2 ART overview

The objective of this chapter is to provide an overview of (1) the key players involved in the test, (2) the key steps and milestones in the ART process and (3) the different modules available. A more detailed description of all phases can be found in Chapters 3, 4 and 5.

## 2.1 Key players

An ART test involves many different parties, each with its own role, task and responsibility. The key participants and their roles are set out below. A more detailed overview of the participants and their responsibilities during an ART test can be found in found in the ART guides.

## 2.2 Key steps and milestones

The ART test comprises three distinct phases: preparation, testing and closure. Some are characterised by specific prerequisites that must be met before progressing to the subsequent phase. This chapter offers an overview of the various steps involved, the key milestones to be achieved and the average time allocation for each phase.

| Name | Participation | Role |
|---|---|---|
| Control team (CT) | Mandatory | Test owner. Has the final say in key decisions, for instance regarding the scope of the test, go/no go decisions and learning goals. Also responsible for various practical matters, such as meetings and planning. |
| Red team (RT) | Mandatory | Responsible for the actual ethical hacking based on threat scenarios developed earlier in the test. |
| Blue team (BT) | Mandatory | The team responsible for the entity's cyber defence. It should not be aware of the test until it has been completed. |
| Gold team provider (GTP) | Optional | Responsible for developing, facilitating and evaluating the gold teaming session, based on the results of the red team test. The GT provider can be the same as the RT provider. |
| Threat intelligence provider (TIP) | Optional | Responsible for providing threat intelligence and custom scenarios preceding the red teaming phase. The TI provider can be the same as the RT provider. |
| Board of directors (BoD) | Mandatory | Sponsor of the test. One of the board members is part of the CT and has to formally approve certain milestones during the test. |
| Test Cyber Team (TCT) | Mandatory | Ensures that the test is performed in a uniform and controlled manner, in accordance with the requirements of the ART framework. |
| GT/TIP | Optional | Responsible for the planning and execution of the gold team exercise, if this module is chosen |

## ART phases and modules

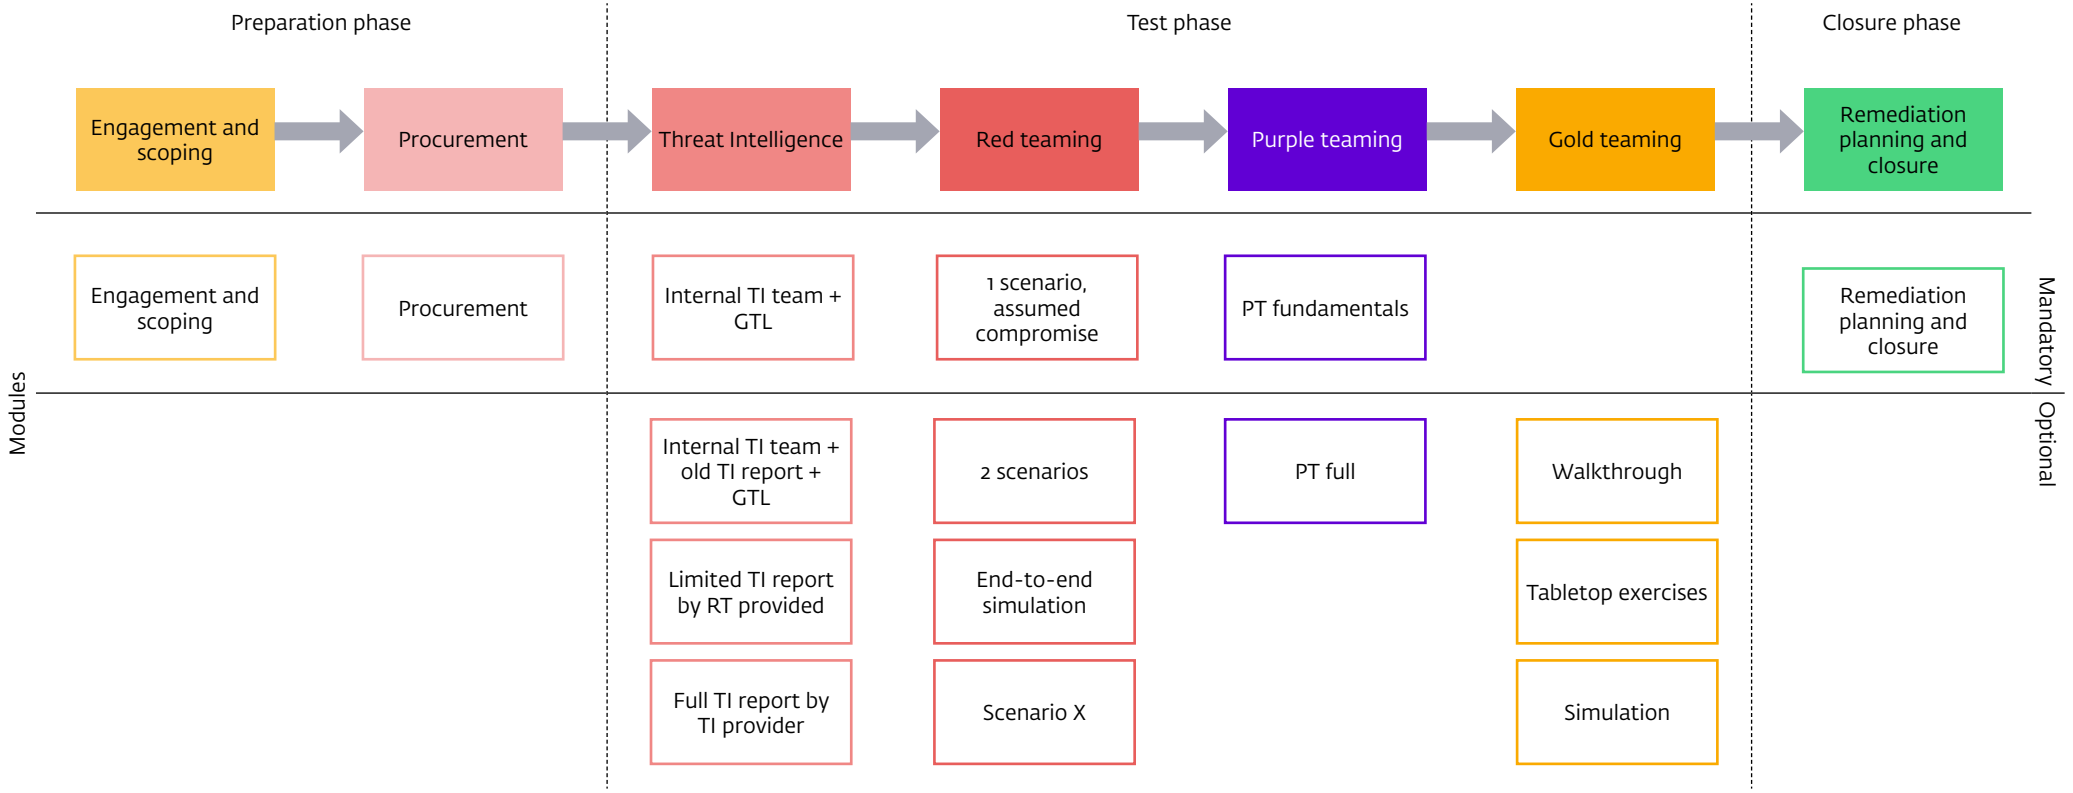

## Engagement and scoping

**What is it?**

In the engagement and scoping phase, the TCT and the Control Team Lead (CTL) collaborate to define the parameters of the ART test. This involves identifying the entity's learning objectives, selecting the appropriate modules, determining the composition of the control team, scheduling the test, assessing the current level of cyber resilience and ascertaining whether the entity has a comprehensive overview of its critical systems and processes.

**Milestones**

- Signed contract between the entity and DNB for ART guidance
- A completed ART checklist defining the test's scope and timeframe
- Board-level commitment
- A filled in and approved scoping document

Average duration: 4-6 weeks

## Procurement

**What is it?**

In the procurement phase, the CTL reaches out to several TI/RT/GT providers to request quotations for the ART test as defined during the engagement and scoping phase. The ART procurement guide serves as a valuable resource to assist the CTL in this process. The duration of this phase is subject to significant variation, largely contingent on the complexity of the entity's procurement procedures. Although the scenario is not yet clear, the FI starts working on a leg up inventory and prepares leg ups.

**Milestones**

- Successful tendering procedure based on the test requirements
- Signed contract between the entity and the TI/RT/GT provider
- Leg up inventory and preparation

Average duration: 6-8 weeks

## Threat intelligence

**What is it?**

In this phase, the TI provider, which can either be a separate entity or the RTP, formulates one or more threat-based scenarios that will form the basis for the RT plan and execution. The required timeframe, scope and depth of the research, as well as the resulting TI report, depend on the number of scenarios selected and the specific TI module chosen.

**Milestones**

- Successful go meeting, formalising the start of the ART test
- A business meeting where the entity provides information about its critical IT and business processes, customers and other relevant developments to the TI/RT provider
- Successful creation (and approval) of one or more TI-based scenarios that serve as the foundation for the RT plan
- Threat Intelligence report

Average duration: Between 2 and 8 weeks

## Red teaming

**What is it?**

In this phase, the RT provider translates the threat intelligence (TI) scenarios into a practical RT plan, structured according to the MITRE ATT&CK framework. Once the board-level sponsor and the CTL have given their final approval to the RT plan, the RT provider proceeds to execute the actual ethical hacking component of the test. The duration of this phase varies based on the specific modules selected for the red teaming test.

**Milestones**

- A draft version of the RT plan
- An actionable and formalised RT plan
- A go/no go meeting where all parties involved vote on whether the RT plan can be executed
- Reaching the RT's predetermined flags

Average duration: Between 6 and 12 weeks

## Purple teaming

**What is it?**

During purple teaming, the red team and blue team set up a collaborative workshop where they discuss the executed scenarios step by step. The two teams work together to share insights, weaknesses and attack paths that were used or discovered during the test. The goal of this collaboration is to enhance the organisation's security posture. The duration of the purple teaming exercise depends on the chosen purple teaming module. It is important to note that the sequence of purple teaming and gold teaming is determined by the characteristics of the ART test.

**Milestones**

- The RT report in which the RT describes in detail which actions it has taken during the ethical hacking phase
- A purple teaming plan created by the RT
- Hosting the actual purple teaming session with the blue team and red team
- A purple teaming report

Average duration: Around 2 weeks

## Gold teaming

**What is it?**

Gold teaming is an optional crisis management module in which the GT provider continues from where the RT phase concludes. In this phase, the simulated consequences of the cyberattack scenario are elevated to CMT level to assess how the organisation's crisis management structure responds to a cyber crisis. The inclusion of a GT module allows the entity to evaluate not only its digital resilience against a cyberattack but also its organisational resilience in addressing the aftermath of a cyber crisis. The duration of the gold teaming exercise is determined by the specific module chosen. It is important to note that the sequence of purple teaming and gold teaming is determined by the characteristics of the ART test.

**Milestones**

- A GT plan based on the chosen module and executed RT scenarios
- Successful go meeting, formalising the start of the GT phase
- Execution of the GT session
- Gold teaming report

Average duration: Between 4 and 10 weeks

## Remediation and closure

**What is it?**

The remediation and closure phase is the final phase of the ART test, in which the entity begins implementing a plan to address and resolve the vulnerabilities, weaknesses and issues identified during the ART assessment. In addition, the ART test and participants are evaluated during a 360 feedback session. All the relevant documentation is formalised, and if the test has been completed in line with the ART standards the TCT signs off on the attestation document.

**Milestones**

- Delivery of the RT summary
- A 360 feedback session
- A 360 feedback report
- Filled in attestation document
- Remediation plan

Average duration: Around 2 weeks

## 2.3 Key mandatory and optional modules

ART offers multiple mandatory and optional modules. This section gives a high-level overview of these modules. For a more extensive explanation of each module, please consult Annex 4.

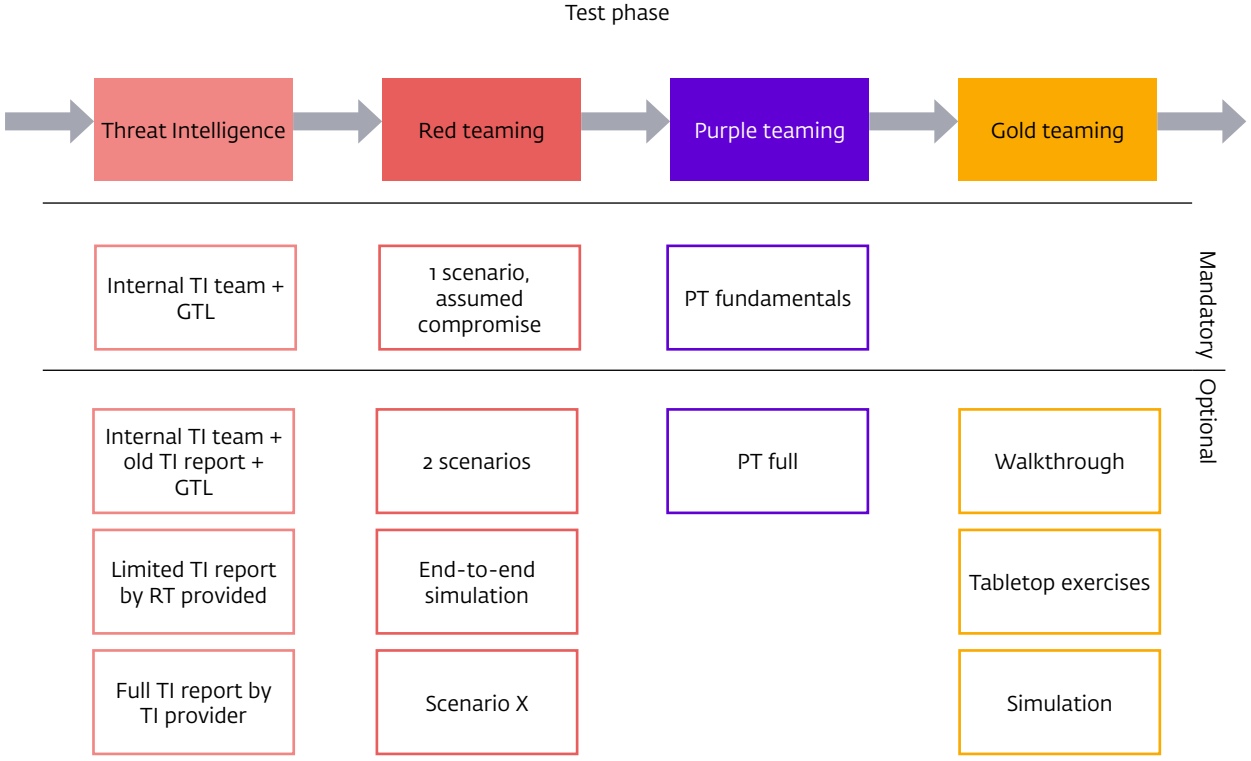

**Why are there mandatory and optional modules?**
Financial entities differ tremendously in terms of the products they offer, the systems and suppliers they use and their current level of cyber maturity. ART acknowledges this diversity by using a modular build. Some ART test components are mandatory, to ensure that the test meets the minimum standards. Besides these mandatory components, entities get to choose from several optional modules to make the ART test fit their budget, learning goals and security posture. This way, the entity gets maximum learning value for its investment.

**Who chooses the modules?**
In the earliest stages of an ART test, before procurement, the CTL and the TCT meet up to discuss a number of things (see Chapter 3). An important topic is the scope of the ART test. Which optional modules should be included depends on a number of factors, as mentioned in the previous paragraph. Although the decision as to which modules to include ultimately lies with the entity itself, the road leading up to this decision is a collaboration between the CTL and the TCT. The following factors play a role in this process:

- the entity's size
- the entity's characteristics
- How this test relates to earlier and planned other tests
- the entity's budget
- the entity's previous experience with threat-led penetration testing
- the entity's learning objectives
- the entity's ambitions
- the frequency of the test

**Which modules are there to choose from?**
The optional modules are all part of the test phase. They include the following steps: threat intelligence, red teaming, purple teaming and gold teaming. In this section, each of these steps is described, along with the available modules and their main objectives. A more detailed overview of the different modules can be found in Appendix 4.

## Threat intelligence

Internal TI team + GTL

Every year, DNB produces the Generic Threat Landscape (GTL), which identifies key threats, actors and scenarios in the Dutch financial sector. This report is shared with every entity and provider that conducts ART tests. Based on this GTL, the entity can produce its own basic ART TI report with the help of internal TI experts. This option can only be chosen if the entity has a suitable internal TI expert. This is the minimum TI option that must be included in the ART test. The advantage of this option is that it requires relatively little time and effort. The disadvantage is that, based solely on the GTL, only relatively general scenarios can be created, and there will likely be a lack of true entity-specific depth. If the entity cannot or does not want to create scenarios on its own, the TI will be carried out by a TI or RT provider.

Internal TI team + old TI report + GTL

Some entities could have, in the twentyfour months prior to the ART test, commissioned a TI report from an external provider. The findings from this report can, under certain conditions, be incorporated into the ART TI phase. Based on this previous TI report and the mandatory GTL, the internal TI experts(s) of the entity can create an ART TI report ART scenarios. Whether the entity itself is able to create ART scenarios depends on its capacity

and maturity level. This determination is made in collaboration between the CTL and the TCT. If the entity cannot or does not want to create scenarios on its own, the TI will be carried out by a TI or RT provider (see next two options). The suitability of the previous TI report for the TI phase of ART is assessed in consultation between the CTL and the TCT.

Limited TI report by RT provider

In this option, the entity procures limited TI services from the RT provider. It is also possible to obtain a limited TI report from a specialised TI provider. The scenarios in this limited report are more comprehensive and in-depth than those in GTL and Previous TI report, but they may still lack the technical depth and exploration that a full TI report typically provides.

Full TI report by TI provider

For a full TI report, the entity must hire a provider that specialises in TI reporting. Sometimes, the RT provider will also offer this service. A complete TI report always provides multiple scenarios for the entity to choose from, as well as targeted threat intelligence (TTI), which can be a significant source of information for the RT as it creates its RT plan. This is the most comprehensive and costly TI module, but it also provides the greatest added value for an ART test.

## Red teaming

1 scenario, assumed compromise

The ethical hacking phase of an ART test includes at least one mandatory scenario. If the available time and resources are limited, or if the specified learning experience calls for it, it is possible to skip one of the classic in-through-out phases with proper justification. This makes it possible to focus the available time and resources on the part where the organisation can learn the most. In some cases, the skipped phase can still be simulated later during the purple teaming phase. For example: it is possible to efficiently simulate the in-phase after through and out with knowledge gained during those latter two phases.

2 scenarios

Often, a TI report will indicate that more than one threat scenario is realistic. Therefore, it is possible for the entity and RT provider to execute more than one scenario during the red teaming phase. Especially when there is enough diversity in actors, tactics, techniques and procedures (TTPs), and objectives in the different scenarios, valuable additional findings can emerge from a second scenario.

End-to-end scenario

To make an ART test as realistic as possible, an entity can choose to incorporate all the steps of the TI scenario into the actual ethical hacking. This means that the in, through and out phases, if possible, are all fully simulated during the RT phase. This method of ethical hacking generally requires more time and resources, but it provides the most comprehensive view of an organisation's cyber resilience across all aspects.

Scenario X

A Scenario X can be included in addition to the planned scenarios. The goal of a Scenario X is to emulate attacks that may be expected in the near future. This scenario can focus on innovative techniques and emerging tactics. Scenario X utilises findings from the earlier scenarios and is developed during the RT phase. The ultimate goal of a Scenario X is to target a critical function, often using a highly creative approach.

## Purple teaming

PT fundamentals

The mandatory option for the PT phase is a one-day purple teaming exercise. During PT fundamentals, the RT and BT share intelligence, review the simulated attacks and analyse the findings. They also propose ways to improve the entity's defences. A PT fundamentals option is suitable for ART tests with a relatively compact RT phase. Towards the end of the RT phase, the TCT will agree with the CT whether PT fundamentals is adequate or not.

PT extended

The extended option for the PT phase is a purple teaming exercise of more than one day. During PT, the RT and BT share more intelligence and review the simulated attacks and analyse the findings in greater depth, before proposing ways to improve the entity's defences.

## Gold teaming

Walkthrough session

This is the most low-key and accessible GT variant in ART. It can be used by entities with no or very limited experience in crisis management. A walkthrough session can also be a good choice for entities that have seen significant changes in their crisis management structure and personnel. A walkthrough session is a discussion-based meeting aimed at validating plans, processes and procedures.

Tabletop exercise

This GT variant is an accessible discussion-based exercise and a good choice for entities with a crisis management team that already has some experience in crisis management, but that do not want to subject their team to a full simulation. The goal of a tabletop exercise is to train crisis management capabilities (based on learning goals) in a low-stress environment.

Full simulation

A simulation is the most elaborate and challenging GT variant in ART. It is intended for experienced crisis management teams that want to step up their game. The goal of a simulation is to test and train crisis management capabilities (based on learning goals) under stress, by confronting team members with a realistically simulated scenario unfolding in real time.

## 2.4 Key documentation

The following frameworks, guides, and formats are essential for organizing and understanding an ART test. The ART framework is leading in this regard. The processes and steps outlined in it will be elaborated upon in the underlying guides.

ART framework
ART Quality Assurance Format
ART Procurement guide
ART Control team guide
ART Threat Intelligence guide
ART Red team guide
ART Purple team guide
ART Gold team guide
ART Scoping Format
ART 360 Feedback Format
ART Summary Format

These documents can be found at https://www.dnb.nl/voor-de-sector/betalingsverkeer/art/

# 3 Organising an ART test

This chapter provides an overview of the key elements that need to be addressed, prepared and organised before a financial entity begins an actual ART test. It includes insights related to risk management, project management, reporting and responsibilities.

## 3.1 Key players in an ART test

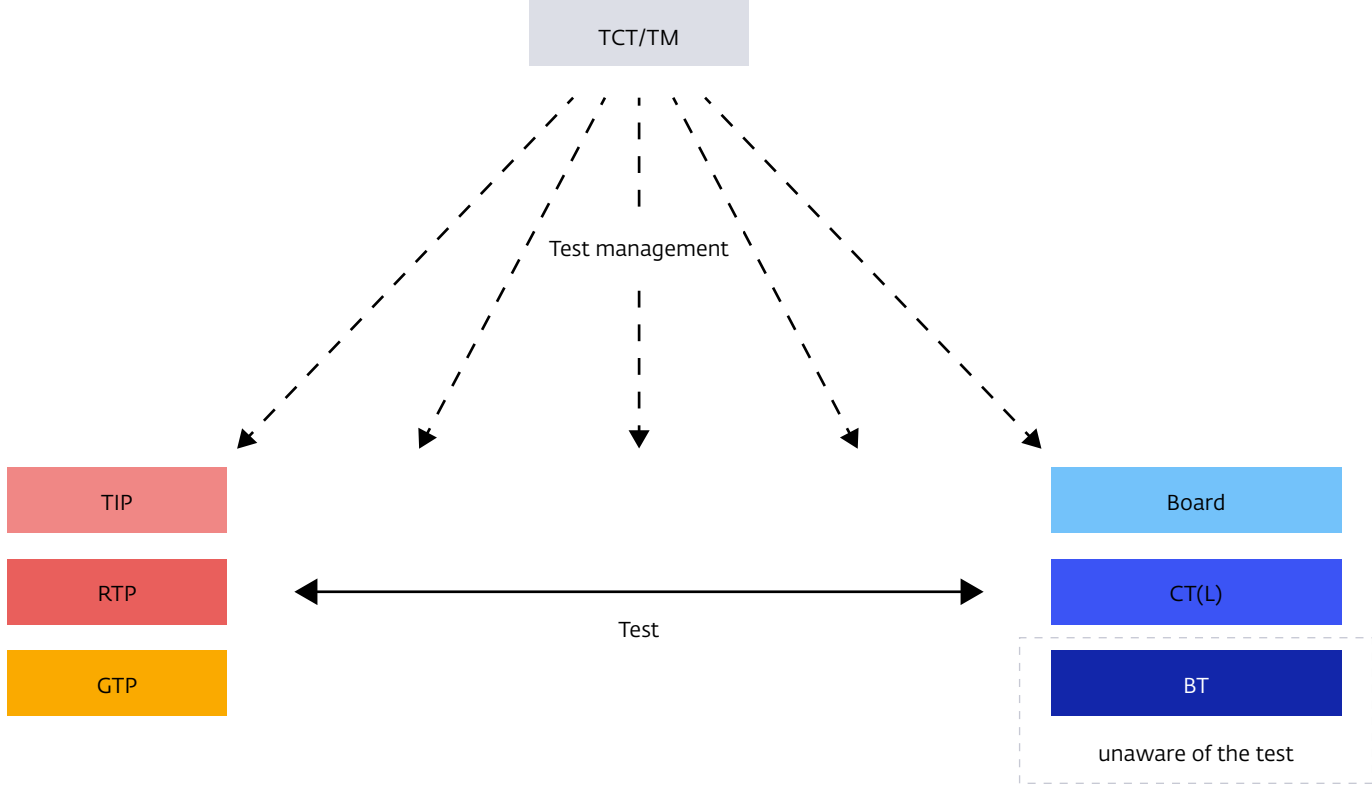


The most important stakeholders in a test are:

- Entity's board of directors
- Control team and control team lead (CT and CTL)
- Entity's blue team (BT)
- Red team provider (RTP)
- Optional: threat intelligence provider (TIP)
- Optional: gold team provider (GTP)
- Test Cyber Team (TCT) and the test manager (TM)

**Entity's board of directors**

The board of directors is an important stakeholder throughout the test, in various ways. One of the board members is part of the control team and has to formally give the go-ahead at the start of the test. This "C-level sponsor" will be actively aware of the test and what is happening. If necessary, this person can make decisions with regard to certain events during the test. It is the responsibility of the CTL to keep the board member involved and

up to date during the test. The other board members are not aware of the test and thus only involved during the closure and learning phase. This can either be during the purple teaming or gold teaming sessions, or when the test is finished. After each test, the CT and the board must allocate time to allow the CT to present the findings and proposed improvements.

**Control team**
The control team (CT) is the team that manages the entity's involvement in the test. The CT members are the only staff within the entity who are fully aware of the test. The CT consists of a control team lead (CTL) and their mandatory substitute, subject matter experts and, if necessary, third-party members. If possible and desired, the CTL can involve an internal TI expert in the TI phase of the test to improve the TI scenario(s). A sponsoring board member is also part of the CT, but does not receive daily updates. This person is kept in the loop on important developments by the CTL.

For more information on the control team, please consult:
1. the control team guide. For more information.

**Blue team**
The blue team (BT) is the entity's defensive team. This is usually a security operations centre (SOC), but it can also be another department. The BT should not be aware of the test until it is finished. However, a situation may arise where, due to chance or as a result of good work, the BT finds out about the test (or parts of the test) before it has been completed. After the test phase is over, the BT can be made fully aware of the test. Together with the RT, it will evaluate the findings and create its learning experience during the purple teaming session. Besides technical personnel, such as security operations centre staff and IT administrators, the BT consists of everyone who is not part of the CT and has therefore not been informed about the ongoing test. This ranges from staff who receive phishing emails to personnel whose accounts might be compromised during the test.

**Red team provider**
The red team provider (RTP) is responsible for carrying out the scenario-based ethical hacking part of the ART test, for which it should provide a team of technical experts. There must be a red team lead and a number of other members who specialise in various fields of red teaming. The main products delivered by the RTP are the red team attack plan and the red team test report. The red team is also responsible for organising and running the purple teaming sessions.

For more information on the red team provider, please consult:
1. the ART procurement guide. For more information.
2. the red team guide. For more information.
3. the red team test report format. For more information.

**Threat intelligence provider**
The threat intelligence provider (TIP) is responsible for providing (targeted) threat intelligence during the test phase and, if necessary, provides additional intelligence during the RT phase. The TIP team should consist of a threat intelligence lead and one or more analysts. The main product delivered by the TIP is the TTI report, which contains a company overview, a threat landscape for the entity, targeted threat intelligence and possible scenarios.

For more information on the threat intelligence provider, please consult:
1. the ART procurement guide. For more information.
2. the threat intelligence guide. For more information.

**Gold team provider**
The gold team provider (GTP) is responsible for organising and executing the selected gold teaming exercise, based on the scenario used during the red teaming phase of the ART test. The GTP can be the same as the RTP, if the RTP has the required qualifications to organise this component of an ART test. The main products delivered by the GTP are the gold team plan, scenario, exercise materials and the gold team report (for the tabletop and simulation variant).

For more information on the gold team provider, please consult:
1. the ART procurement guideline. For more information.
2. the gold teaming guide. For more information.

**Test Cyber Team**

The role of the Test Cyber Team (TCT) is to make sure that entities are tested in a uniform and controlled manner, in accordance with the requirements of the ART framework. The TCT writes and provides the GTL. It also appoints a test manager (TM) for each test, who works closely together with the entity's CT throughout the entire ART process. The TM guides the CT through the ART phases, but can in no way be held accountable for the CT's actions or any consequences of the ART test. The TM has a close relationship with the CT but is not formally part of the team. They have the right to escalate major deviations from the test scope or scenario to the TCT programme manager, who they directly report to.

For more information on the Test Cyber Team, please consult:

1. the ART procurement guide. For more information.
2. the control teaming guide. For more information.

## 3.2 The control team's project management responsibilities

The CTL is responsible for managing the project and the risks of the ART test. This means that it is, amongst others, responsible for planning the mandatory meetings, agreeing on ways of communication, password policies, keeping track of risks and drafting a high-level overall schedule for the entire test. Project management also involves making sure that internal stakeholders, such as the board, are included in the test in a timely manner, and that the external parties deliver according to schedule or that the schedule is adjusted in the event of changes. The schedule must be created and shared with every party involved.

## 3.3 Overview of mandatory reporting

One of the objectives of ART is to limit mandatory documentation where possible. A lower documentation load prevents a unnecessary burden on CTs and RTs. Nevertheless, a certain degree of documentation is essential. Not least because several (essential) reports and logs form the foundation for improvements within the tested organisation. The following documentation needs to be produced/maintained during or after the completion of an ART test:

| Name | Author | Goal |
|---|---|---|
| ART agreement | TCT and CT | The ART agreement outlines the legal foundation, the scope of the ART test, the statement of work and other procedural agreements that underlie the test. |
| Scoping | CT | A document that identifies vital targets, systems and assets that will be included in the test. The scoping document ensures that the CT, TCT, RTP and TIP have a clear understanding of the areas, systems and processes of interest, and that they stay within the predefined scope. |
| TI report (depending on module) | TIP/RTP/CT | A document that outlines the financial entity's threat landscape. It also identifies risks and crown jewels, and offers a business overview and attack scenarios. Based on this information, the RTP makes its RT plan. |
| RT test plan (depending on module) | RTP | A document in which the TI scenarios are translated into a technical attack plan, complete with TTPs and MOs of the selected threat actors (using MITRE). Additionally, the document should include descriptions of potential leg-ups, risk management and the expected timeline for execution. |
| RT report | RTP | An RT report is a comprehensive document that sets out the findings, observations and insights from the RT phase. |
| GT plan | GTP | A GT plan describes the plan of approach for the preparation and execution of the GT. It includes the scope, learning goals, high level scenario and risk management of the GT. |
| GT report | GTP | A GT report is a comprehensive document that sets out the findings, observations and insights from the GT phase. |
| Test summary | RTP/CT | The test summary is a concise document that provides an overview of the key findings, outcomes and insights from the test. It serves as a high-level summary for stakeholders who may not require detailed technical information but need a clear understanding of the test's results and implications. |
| Attestation document | TCT | A document certifying that the ART test has been performed in accordance with ART standards and that the test has been formally completed. An attestation document is not an indication of the quality of the entity's defences. |
| 360 feedback report | TCT | A summary of the key findings from the evaluation of the ART test. |

## 3.4 Important meetings

This is a non-exhaustive overview of the key meetings during an ART test. For a complete list of meetings, please refer to Chapters 4, 5 and 6.

| Name | Authors | Goal |
|---|---|---|
| Pre-launch meeting | TCT and CT | Meeting before the formal start of the ART test between the CT and the TCT in which the scope, the modules and other fundamental prerequisites are discussed. |
| Scoping meeting | TCT, CT, a board member and, if procured, the TIP and RTP | During the (final) scoping meeting, the scoping document is agreed by the TCT and the entity's C-level sponsor. |
| Launch meeting | TCT, CT, (TIP) and RTP | Formal launch of the ART test. During the launch meeting, the following topics are discussed:<br>■ the ART process and documentation<br>■ other TCT members involved<br>■ stakeholders, roles and responsibilities<br>■ project planning<br>The end of this meeting marks the formal start of the ART test. |
| Business overview workshop | TCT, CT, (TIP) and RTP | Workshop given by the entity's business expert to support the RTP/TIP in its understanding of the entity. |
| Weekly update meetings during TI and RT phase | TCT, CT, (TIP) and RTP | During the test phase, there are weekly update meetings where the TIP/RTP gives an update on the progress made in the preceding week. The activities for the upcoming week are also discussed. |
| Go/no go TI report | TCT, CT and TIP or RTP | After the RTP/TIP/entity delivers the TTI report, a meeting is held in which the report is formally approved. |
| Go/no go RT attack plan | TCT, CT and RTP | After the RTP has created the attack plan, a meeting is held to formally approve the attack plan and start the RT phase of the test. |
| Purple teaming | CT, RTP<br>Optional: TCT | Purple teaming kicks off with a replay session, during which a chronological summary of the test is created. After that, the BT and RTP focus on the areas with the most learning opportunities. |
| Optional:<br>gold teaming kick-off | TCT, CT, RTP and GTP | During the kick-off, the scope, learning goals and set-up of the gold teaming phase are defined. It also marks the formal start of the gold teaming phase. |
| Optional:<br>go/no go GT plan | TCT, CT and GTP | After the GTP has created the GT Plan, a meeting is held to formally approve the GT plan and start the GT phase of the test. This can either be a go or no go meeting. |
| Optional:<br>gold teaming dry run | CT and GTP<br>Optional: RTP and TIP | Depending on the chosen variant, a gold teaming dry run is conducted to check and verify that everything is in place for successful execution. |
| Board meeting | TCT, CT and board<br>Optional: TIP, RTP and GTP | After the PT and GT sessions and finalisation of the RT report, a board meeting is held to communicate the results and the impact of the test. |
| 360 feedback session | TCT, CT, (TIP), RTP (and GTP) | During the 360 feedback session, all parties that were actively involved evaluate the test. The subject of the evaluation is the ART-NL process, not the results of the test. |

## 3.5 Risk management

An ART test always involves potential risks. This is due to the critical role of the targeted systems, people and processes.

**Mapping and reducing risks**

Before an entity engages in an ART test, it should conduct thorough due diligence of any systems that might fall within the scope of the test to ensure that backups are in place and any damage can be restored. Furthermore, the entity should conduct an assessment of the risks involved in an ART test, take these into consideration and put in place effective mitigation measures. Such a risk assessment should at least consider the following risks:

- risks related to entering into the contractual relationship with (a) provider(s) and the confidentiality of the information that becomes accessible to that provider;
- 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 red teaming;
- risks related to operational defence;
- risks related to clean-up after completion of the test.

When hiring a provider (RTP, TIP or GTP), the entity makes sure that there is mutual agreement on at least the following aspects: the scope of the test, boundaries, timing and availability of the providers, contracts, actions to be taken and liability (including insurance, where applicable). In addition, the TM's involvement in the ART test ensures that the test proceeds according to the agreed test scope, scenario, planning and process, as described in the jointly developed framework documents. The minimum requirements for cybersecurity service providers are set out in the ART services procurement guidelines.

Risks are also mitigated by sound planning, informing only a select group of people in higher management about the test and its scope, keeping track of an up to date risk register during the testing process and a clear definition of the scope and predefined escalation procedures. It is important to note that the entity remains in control of, and responsible for, the test. At any time, the CT can (temporary) 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 help prevent miscommunication and knowledge about the ART test leaking out.

**Ethical boundaries**

An ART test should mimic the (seen, current and potential future) actions of a real threat actor. Criminals usually do not stick to ethical rules, and an ART test should use the same kind of "creative thinking" criminals would use – up to a point – to make the test as realistic as possible. Despite this objective, 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 names, logos or otherwise identifiable information of real people or companies.

**Code names**

To prevent the leakage of sensitive information, code names must be used. These code names should be used throughout all documents related to the ART test as best as possible, but at least in document titles and throughout 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 entity. The TCT will assign a code name to each unique test. This code name will be used in all communication and documentation between the parties involved in the test. Besides this code name, providers and/or the

entity are free to use their own code names for internal communication.

**Escalation and stopping 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 paused so that measures can be taken to prevent these authorities from getting involved.

**Personal identifiable information**

It is up to the entity to set up contractual agreements with the RTP regarding, for instance, the inviolability of their employees' privacy. Under no circumstances may privacy-related information be included in test reports.

## 3.6 Stopping the test and/or removing the ART label

As the TCT is not involved in the commercial relationship between the RTP and the entity, it cannot stop the test. However, it does have the power to remove/deny the ART label, which means the test will not recognised as an official ART test. For multi-sector tests, this also means that the test will not be recognised as an ART test in other sectors. The TCT must therefore exercise restraint in deciding to remove the ART label, giving due consideration to the quality and safety of the exercise. Any decision to remove the label must always be made in consultation with the CTL, unless the situation does not permit this.

The TCT can remove the ART label in at least the following situations:

- either the TIP or the RTP has (repeatedly) shown that it cannot live up to the standards set out in the ART framework and/or has lost the confidence of the TCT and/or CT it can perform its duties in a controlled manner appropriate to the delicate nature of the covert test;
- the test has been compromised by the RTP, TIP or the entity, either intentionally or as a result of (gross) negligence;
- if there is foul play by the CT or BT;
- other situations that compromise the quality, safety or secrecy of the test.

Should the TCT decide to remove the ART label, the entity can choose to continue the test for learning purposes, or it can consult the TCT about the steps that would have to be taken to secure ART recognition.

# 4 Building an ART test

## the procurement phase

"Building an ART test" refers to the preparation phase of the ART test, which includes the engagement, scoping and procurement of external parties. This chapter will cover these three aspects of the preparation phase.

### 4.1 Engagement

The main goal of the engagement phase is to define the entity's learning objectives and the scope of the test, and to get commitment from all parties involved. During the phase, the TCT and the CTL also determine which modules are going to be included in the test. Another objective of this early phase is that the entity makes sure that all relevant internal stakeholders in the ART test are involved and aligned. The entity also ensures that the TCT is engaged. With the guidance of the TCT, the entity can then begin setting up the ART test. During this phase, the TCT and the CTL can use the ART checklist to determine if all the steps required to formally start an ART test have been completed. All of the above is discussed during a pre-launch meeting (or in multiple pre-launch meetings), which is the first official meeting of the ART test. During this meeting, the TCT asks the CTL to establish a CT, comprised of a select number of senior staff who have the required expertise and/or are part of the security incident escalation chain. The CTL makes sure that they are aware of the ART test, the need for secrecy and the process the team should follow if the BT detects and escalates an ART-related incident.

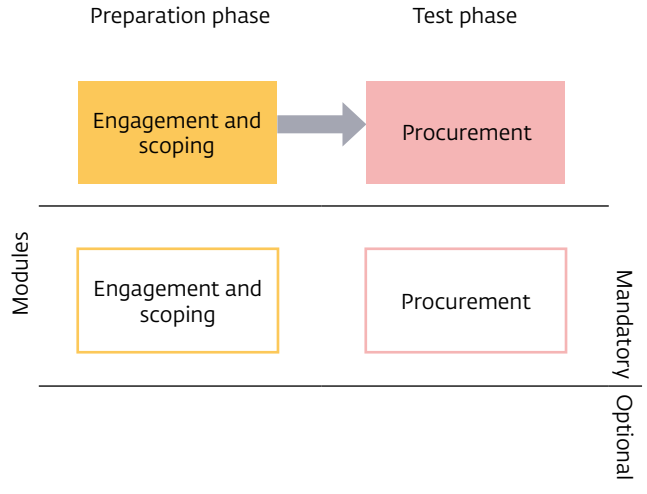


**Choosing the right modules for the right test**

As mentioned above, during the preparation and launch phases, the CTL and the TCT meet to discuss the scope of the ART test and to decide which modules to include. The final decision depends on a number of factors, including:

- the entity's size
- the entity's characteristics
- the entity's budget
- the entity's level of experience with threat-led penetration testing (TLPT)
- the entity's learning objectives
- the entity's ambitions
- the frequency of the test

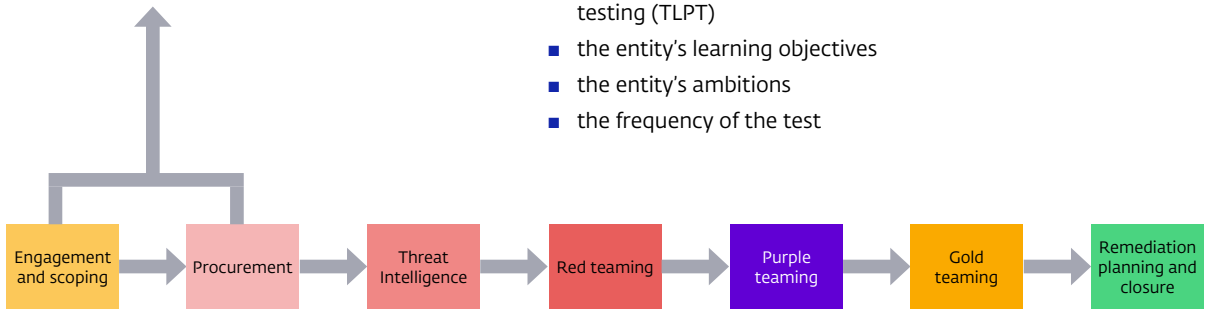

Although the decision as to which modules to include in the ART test ultimately lies with the entity itself, the road leading up to this decision is a collaboration between the CTL and the TCT.

**Meetings**

- Pre-launch meeting

**Milestones**

- Signed contract between DNB and the entity
- Filled in ART checklist that includes the modules chosen for the test
- Establishment of a control team
- Agreement on code names and communication channels.
- The ART scope is approved by a C-level executive of the entity

## 4.2 Scoping: critical functions and systems

Critical functions (CFs) are defined as the people, processes and technologies required to deliver a core service which, if disrupted, could have an impact on the Netherlands' financial stability, or on the entity's safety and soundness, customer base or market conduct.

Entities across the financial sector support and deliver these functions in different ways through their own internal processes, which are underpinned by critical systems. It is these critical systems, processes and the people involved in them that are the focus of ART threat intelligence and red teaming. Flags are placed on the critical systems in the ART scope specification. These flags will later serve as goals in the test scenarios, which are based on relevant threat intelligence. The TCT involves supervision and/or oversight during the scoping phase to verify that the scope is a realistic representation of the entity.

During the scoping process, the entity must complete the ART scope specification. In addition to defining the scope, the ART scope specification lists the key systems and services that underpin each CF. This information helps the CT place the "flags" to be captured, which are essentially the targets and objectives the RTP must strive to achieve during the test.

The CT should discuss the flags with the TM, who must approve them. Although the flags are set during the scoping process, they may be changed in some cases, based on threat intelligence and as the test evolves.

**Meetings**

- Scoping meeting(s). Preferably face-to-face

**Milestones**

- Filled in scoping template
- Scoping document approved by C-level sponsor and TCT

## 4.3 Procurement

Based on the agreed scope and modules, the CTL starts the procurement of an RT provider and/or TI provider and/or GT provider. The ART procurement guide can assist the CT with this task. With regard to contractual considerations, smooth delivery of an ART test requires that the process is transparent and that appropriate information and documentation flows freely between the relevant parties. To facilitate the free flow of information, non-disclosure agreements (NDAs) can be used. The RFP (request for proposal) used to procure an RTP (and TIP and GTP) is shared with the TCT. The TCT then makes sure that the RFP contains all the necessary elements listed in the ART services procurement guidelines. The procurement process is started after the pre-launch meeting. During this process, the entity shares a shortlist of potential providers with the TCT. The TCT can arrange contact between the entity and other TIBER/ART entities to request references for the provider. During the procurement phase, the CT must complete the ART test project plan, including the schedule of meetings to be held between the entity, RTP and TCT (and TIP and GTP). Apart from the mandatory meetings, the TCT and CT should have regular meetings to discuss the progress made. The TCT can, if needed, support the CT in the procurement process or participate in workshops to create a scoping document. After procurement has been completed, the launch meeting can be held to align all stakeholders in the test. During this meeting, a number of practical agreements are made regarding the frequency of meetings during the TI and RT phases, communication channels, documentation and responsibilities.

**Meetings**

- Launch meeting

**Milestones**

- Signed contract with RT provider and/or TI provider and/or GT provider

# 5 Running an ART test

## the test phase

"Running an ART test" refers to the test phase of an ART test. This is the part where threat intelligence is gathered, scenarios are developed and the actual ethical hacking is performed. In this phase, purple teaming and (optionally) gold teaming also take place.

### 5.1 Threat intelligence

The threat intelligence phase is the first active phase in an ART test. It involves gathering information about the entity and identifying potential threats. This information is used to simulate real-world attack scenarios in the later stages of the test. During the ART test, one or more of four TI modules are selected by the entity, each of which fits different learning goals. Despite the differences between the TI modules, the procedures and mandatory steps remain the same.

One of the first steps (in module 3 and 4) of the TI phase is a business overview meeting organised by the CT. It is of the utmost importance that the RTP and the TIP not only understand the technical components of the financial entity, but also the business processes.

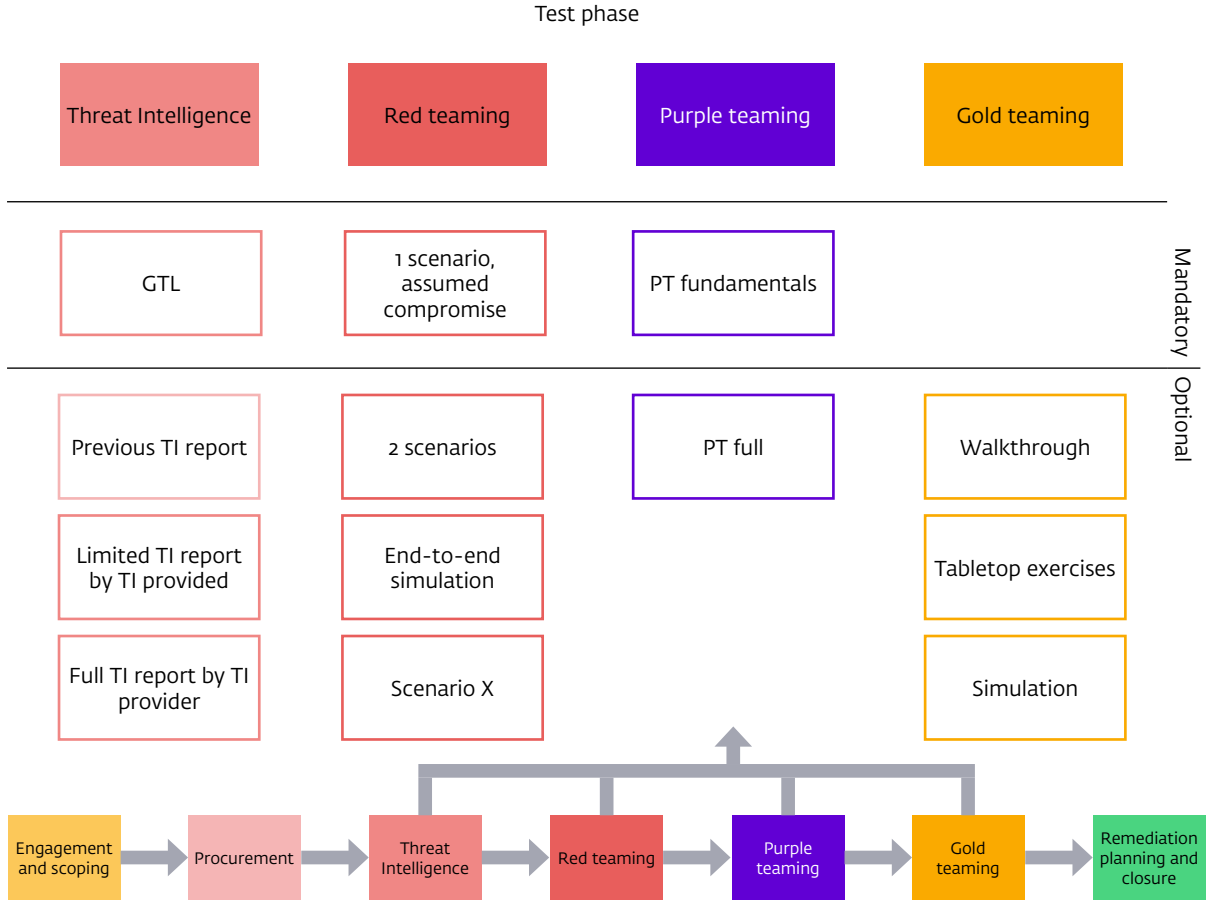

A business overview meeting organised by a CT expert will help the RTP or TIP fully understand the scoping document and make a better assessment as to which threats are applicable to the entity. Based on the scoping document, the business overview meeting and its own intelligence gathering, the RTP/TIP produces a TI report containing one or more realistic threat scenarios.

During this part of the TI phase, the RTP/TIP regularly informs the CT and the TCT of its findings and progress. At an appropriate time, a draft version of the TI report is provided. This ensures that the report continues to meet the requirements set out in the threat intelligence guide. If necessary, the TCT can give feedback on how to align the report with the framework. The CT organises regular meetings to discuss progress, the frequency of which depends on the progress made during the TI phase. The TI phase concludes with a go/no go meeting, attended by the TCT, RTP, TIP and CT, as well as the C-level sponsor. During this meeting, the RTP/TIP presents its findings and the corresponding TI report. The C-level sponsor approves the selected scenario(s) to initiate the preparation of an RT attack plan by the RTP, based on the TI report.

More information on the threat intelligence phase can be found in the threat intelligence guide.

**Meetings**

- Business overview workshop
- Go/no go TI report meeting
- Weekly meetings

**Milestones**

- Finalised version of the TI report
- Go for the creation of the RT attack plan

## 5.2 Red teaming

The RT phase consists of a number of steps. First, the TI report and chosen scenarios are converted into a red team attack plan. Once this plan is approved by the C-level sponsor, TCT and CT, the RT executes it by attacking the entity's live systems. The goal is to identify weaknesses, vulnerabilities and potential gaps in the entity's defences, providing insights it can use to improve its cybersecurity posture and incident response readiness.

**The red team test plan**

In the test plan, the RTP sets out operational attack scenarios for the ART test that:

- incorporate the modules agreed by the CT and the TCT;
- use the TI scenarios drafted in the TI part of the ART engagement to ensure credibility;
- provide background information on the tradecraft of the type of threat actor that is mimicked in the test;
- gather additional OSINT information to help the simulated threat actor achieve its goal;
- provide creative elements using TTPs that have not yet been used in practice but that will likely be used in the future according to the RTP, based on its professional knowledge;
- would, in case of a real attack, have an impact on the Netherlands' financial stability;
- also include some elements that test the entity's response, showing whether the attack would immediately be detected or could have a fair chance of succeeding.

The attack scenarios are written from the threat actor's point of view and are intelligence-led. The RTP presents a number of creative options in each of the test phases based on various TTPs used by advanced threat actors. It does so to anticipate changing circumstances or in case the first option does not work. The RTP should also indicate where a leg-up might be needed if the attack is not successful and what this leg-up will entail. Writing the scenario is a creative process. The TTPs mimic those seen in the past and can combine techniques used by various relevant threat actors to save resources. The RTP should substantiate why it is possible to combine multiple techniques.

**Optional: Scenario X**

In addition to these scenarios, a Scenario X is prepared based on advanced attacks that will likely be used in the near future. This scenario can focus on a specific innovative technique, tactics that are currently being developed (possibly combined with societal developments) or developments in the threat landscape that will

impact the entity in the future. While the ultimate objective of a Scenario X is to compromise a CF, it attempts to do so using a highly creative approach. The CT and TM must agree on the use of a Scenario X, with support from the RTP (and TIP).

**Rules of engagement**

The red team test plan should include rules of engagement, in which the RTP lays down the rules it will follow. The rules of engagement should contain at least the following:

- high-level description of the techniques used during the attack;
- list of excluded techniques;
- detailed description of scenarios used for social engineering;
- how the privacy of both voluntarily and involuntarily participants is being safeguarded in compliance with relevant legislation.

**Approval of the attack plan**

The red team test plan must be approved:

- before the start of the actual test phase (by the CT, TM and RTP);
- optional: when Scenario X is finalised (by the CT, TM, (TIP) and RTP).

**The ethical hacking**

This is where the fun starts. After every party has approved the attack plan, the RTP begins executing the ART test. During this phase, it performs an intelligence-led red teaming test on the target systems. The scenarios are not prescriptive playbooks that must be followed to the letter during the test. If obstacles occur, the RTP should show its creativity (as advanced threat actors would) and develop alternative ways to achieve the test objective. This is always done in close consultation with the CT and TM. All of the RTP's actions are logged so they can be replayed with the BT, as evidence for the RTP report and for future reference. The test objectives are pre-designated "flags", which the RTP must attempt to capture during the test as it progresses through the scenarios. Of course, all captures happen in close cooperation with the CT, and the overall aim is to improve the BT's capabilities. The scenario should be played out from beginning to end. The RTP may need some help to overcome barriers and may be discovered, but the scenario must continue to make full use of the ART exercise within the given timeframe and test all phases (in, through and out). RTPs are constrained by the time and resources available, as well as by moral, ethical and legal boundaries. The RTP may therefore require occasional leg-ups and/or information assistance from the CT to help it progress. If this happens, the assistance must be logged by the RTP. This ensures that all stakeholders derive maximum benefit from a time-limited test.

At all times, the RTP liaises closely with the entity's CT and the TM. The TM is updated at least once a week by the RTP and the CT on the progress. Physical meetings between the CT, TM and RTP during this phase are strongly encouraged, since the discussions this leads to add significantly to the quality of the test. Entities have also had very positive experiences when a member of the CT was onsite with the RTP for some time during the engagement. In case of an in-through-out scenario, the test will have a potential cut-off point. If the RTP has not been able to complete the in phase, it should be given realistic leg-ups so the rest of the scenario can be played out. Alternatively, if the RTP has gained a foothold using another scenario, it can be allowed to use that path for the rest of the scenario where the in phase failed.

**Out phase plan**

The RT attack plan 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 RT attack plan is still aligned with the current planned execution of the scenario. If not, the RTP has to 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 of/during the out phase. Regardless of whether it is aligned with the attack, the out phase has to be discussed with the TM and the CT before it is executed.

**Completing the RT phase**

The output of the ethical hacking phase is a red team test report produced by the RTP and delivered to the entity. The draft report must be shared within two weeks of the test's completion. It must give an overview of the entire ART process, including the CFs, the scope, the threat intelligence base of the test, the planned

scenarios, the executed scenarios, the test findings and the RTP's advice to the entity, and should be written using the RT test report format. At this point, key members of the entity's BT are informed of the test. If desired, they can write their own report ahead of the purple teaming session. If the BT is not able to write a full report as a result of findings or omissions in the monitoring, the RT report can be shared with the BT.

**Meetings**

- Red team test plan go/no go meeting
- Out plan go/no go meeting
- Weekly updates
- Scenario X meeting

**Milestones**

- Finalised version of the RT test plan
- Go for the execution of the RT attack plan
- Reaching pre-designated flags during the ethical hacking
- Filled in red team test report
- Filled in blue team report (if necessary)

## 5.3 Purple teaming

After the RTP delivers its report, the entity organises a purple teaming workshop. Often, the PT phase is perceived as the most educational, leading participants to spend more time on this part of the process. The goal of this workshop is to enhance the learning experience. This workshop can last either one day (PT fundamentals) or two days (PT extended), depending on the scope and duration of the test. Towards the end of the RT phase, the CT and TM will agree which PT module fits best for each test, although a preliminary assessment can be made at the start of the ART engagement.

Purple teaming in ART is an expansion of the replay and enhances the learning experience for both the BT and the RTP. During the purple teaming workshop, the RTP and entity should replay the attack and work together to enhance the entity's defensive capabilities. This will also improve the attacking capabilities of the RTP. The TM is present during parts of this meeting. PT is described in more detail in the ART purple teaming guide.

**Meetings**

- PT planning session

**Milestones**

- Execution of the purple teaming

## 5.4 Gold teaming

Gold teaming (sometimes publicly known as a "tabletop" or "crisis management simulation") is a collaborative session with the crisis management team (CMT) of the financial entity. A gold teaming allows an entity to validate, train and exercise crisis management in a controlled environment. Gold teaming is also a perfect opportunity to get (additional) C-level engagement. During a GT exercise, participants from various departments within the entity gather to discuss and respond to the crisis caused by the completed red teaming scenario. Within the ART framework, the optional gold teaming (GT) module allows the entity's CMT to validate and test crisis management structures, plans and procedures, and to practise managing strategic impact, following a scenario as played in the RT phase. This chapter will give an overview of the different GT variants, the GT planning and the rules of engagements.

**GT modules**

The GT can be developed, organised and facilitated by the tested entity or an external GT provider (GTP). Three GT variants can be considered:

1. **Walkthrough session** – This is the most low-key and accessible GT variant. It can be used for entities with no or very limited experience in crisis management. A walkthrough session could also be a good fit for entities that have seen significant changes in their crisis management structure and personnel. During the walkthrough session, all steps in the crisis management process, from detection to closure, are discussed and completed in detail. A walkthrough session addresses the roles and expectations of CMT members, as well as actions and measures to take in specific crisis scenarios. It is a discussion-based session with the purpose of validating the crisis management processes and increasing the CMT's knowledge of how to act if a specific crisis unfolds. The walkthrough is facilitated by a process supervisor or facilitator.

2. **Tabletop exercise** – This GT variant is an accessible discussion-based exercise and a good fit for entities with a CMT that already has some experience in crisis management but that do not want to submit their CMT to a full simulation. The goal of a tabletop exercise is to practise crisis management in a low-stress environment. The CMT members gain knowledge and skills on an individual level, but the exercise also trains their ability to collectively respond to challenges and work effectively as a team. During tabletop exercises, participants practise specific crisis management capabilities, such as gaining situational awareness, communicating actions and statements, information management and decision-making (depending on the exercise goals). At the start of the tabletop exercise, all participants receive the same scenario information. After this the exercise starts and more role-specific information can be shared with individual participants. A tabletop exercise is always facilitated by a facilitator, trainer and/or observer for training and evaluation purposes.

3. **Full simulation** – The simulation is the most elaborate and challenging GT variant. It is intended for experienced crisis teams that want to step up their game. In an interactive crisis simulation, the entity's crisis management team experiences what it is like to be confronted with a real crisis. The goal of a simulation is to practise and train crisis management capabilities (based on learning goals) under stress, by confronting team members with a realistic crisis scenario unfolding in real time. Under time pressure, the CMT members must decide what to do to mitigate the impact of the crisis. Scenario information and events can be inserted in multiple ways. There are two subvariants:

a. **Simulation without counterplay** – in this subvariant, static or pre-defined scenario injects are sent to the exercise participants through various (fictitious) channels from the exercise control cell;
b. **Simulation using counterplay** – this subvariant uses dynamic scenario injects that are sent to the participants from the exercise control cell. Counterplay events are based on actions taken by the CMT, in order to make the exercise more realistic.

For both subvariants, the simulation set-up includes an exercise bubble containing the exercise participants, facilitator, trainer and observer. The scenario injects and/or responses are sent to the exercise bubble from the exercise control cell, which is located in a separate room or location. Please note that a successful simulation requires profound and meticulous preparation (including a dry run), execution and evaluation.

**GT planning**

GT can follow either the RT phase or the PT phase of an ART test. Which one should come first is to be discussed by the CT and TCT. Regardless of the order of the exercises, there are a number of factors that should be taken into consideration to create a realistic, evidence-based GT. These factors should all be described in a GT plan.

There are multiple GT variants that can be included in an ART test. All have different objectives and levels of complexity. The CT and the TCT decide which one suits the organisation best. During a GT, only fictitious decision-making should take place in a controlled environment. GT is intended for the CMT members of the tested entity. Depending on the scenario and the entity's learning objectives and other response teams can be involved in the GT.

Every GT begins with a kick-off meeting and the creation of a GT plan. Ideally, this process should start well ahead of the GT exercise, preferably during the early stages of the RT phase. The high-level scenario that is described in the GT translates the technical implications of the RT phase to a strategic level, focusing on the organisational (reputational, operational, safety and security) impact of the findings of the ART test. For the development of the scenario, it is important to involve the IT department of the tested entity.

The plan should incorporate a number of elements focused on risk management in the GT exercise. This is to prevent the "contained" exercise from accidentally leaking out to people or personnel that are not part of the test, possibly creating a situation where test injects or information "escape" into the wild. Thorough risk management prevents such escalations. As a result, GT enables entities to practise their crisis management response in a controlled and safe setting. It is important that the GT takes place not too long after the RT and/or PT phase. The effectiveness of

the GT exercise will be increased if the "pain" from the red teaming phase is still felt. For example, if the RT provider simulated a ransomware scenario during the red teaming phase, the gold teaming phase will start as soon as possible after the ransomware has been deployed on the targeted systems. Even if the RTP did not manage to reach the flags in the scenario, the gold team scenario will assume that it did. The impact that the RT had on the organisation should be a factor in the timing of the GT. For example, in some cases the start of the GT should commence directly after the GT. While in other instances it is advisable to delay the start of the GT with a couple of weeks to achieve optimal learning. The ideal starting point of a GT will be determined in consultation with the TCT and CT.

**Rules of engagement**

The gold team plan should include rules of engagement, in which the GTP lays down the rules it will follow. The rules of engagement should contain at least the following:

- Stipulation that exercise participants cannot be involved in planning and development
- Confirmation that the GT is always a safe learning environment for all participants
- The GT scenario builds upon the scenario as played in the RT phase

**Approval of the GT plan**

Before the actual gold teaming exercise begins, both the TCT and the C-level sponsor of the entity's CT have to approve the GT plan. This happens in a GT go/no go meeting.

**Meetings**

- GT kick-off meeting
- GT plan go/no go meeting
- Dry run (depending on the module)

**Milestones**

- Finalised version of the GT plan
- Go for the execution of the GT plan
- Execution of the GT exercise

## 5.5 Removing the ART label from a test

As the TCT is not involved in the commercial relationship between the RTP, TIP and/or GTP and the entity, it cannot stop the test. However, it does have the power to remove/deny the ART label. More on how and when the ART label can be removed/denied, and the consequences of this decision can be found in chapter 3.6.

# 6 Learning from an ART test

## closure phase

The financial entity learns a lot about its own level of cyber resilience during the threat intelligence, red teaming, gold teaming and purple teaming phases. However, there is also much to be gained from the testing experiences of other entities. DNB aims to facilitate a mutual exchange between entities by encouraging information sharing through community building. This topic will be explained in this chapter.

### 6.1 Test summary

The ART test summary summarises the ART process and should draw on the delivered documentation, such as the RT and BT reports, the targeted threat intelligence and (if available) the remediation plan(s). The entity should use the test summary format for this. The gathered intelligence and lessons learned from the test will serve as input for the Generic Threat Landscape used in future tests.

### 6.2 The 360 feedback

During the 360 feedback meeting, the CT, TCT and GT come together to review the ART exercise. The TM arranges and facilitates the workshop. The goal is to further facilitate the learning experience of all those involved in the process and to improve future exercises.

### 6.3 Remediation plan

Based on the test outcomes, the entity should create a remediation plan. The ART documentation can be used to support the business case for implementing improvements to mitigate the vulnerabilities identified during the ART test. The TI report, RT report and GT report can serve as input for the remediation plan. Further input could come from the CT and organisational findings. DNB encourages entities to share their remediation plan with their supervisors. The TCT is not involved in the creation of the remediation plan.

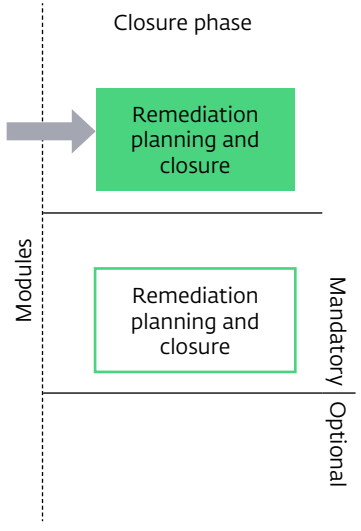


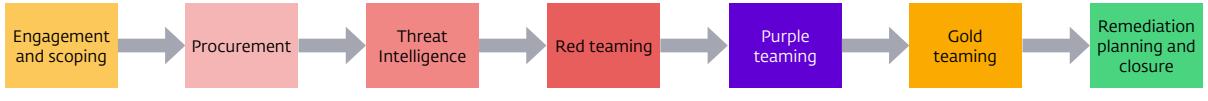

## 6.4 Reporting to C-level

It is of the utmost importance that the entity's board is informed of the threats, test results and the remediation plan (risk mitigation measures). If desired by the CTL, the TCT attends the presentation of the results and findings to the board. The TCT must stress the importance of board involvement, support and accountability in executing the remediation plan.

## 6.5 Finalising the test and attestation document

After the test has been completed, the results have been shared and the purple teaming is finished, the CTL should make sure that any traces of the test are cleaned up. This means that any traces of malware used during the test should be removed, and that the participating teams remove all test data. The RTP should assist the CTL, and all communication groups should be dissolved (unless they are still needed). After this is done, the CTL and the TCT agree that the ART test has ended. If the test has been carried out in accordance with the requirements of the ART framework, the TCT will provide the entity with an attestation document. At this moment, the CTL informs the supervisor(s) that an ART test has taken place. This is a mandatory part of the ART process.

# 7 Annex



# Annex 1

| | |
|---|---|
| ART | Advanced red teaming |
| BoD | Board of directors |
| BT | Blue team |
| CMT | Crisis management team |
| CF | Critical functions |
| CT | Control team |
| CTL | Control team lead |
| DNB | De Nederlandsche Bank |
| DORA | Digital Operational Resilience Act |
| GT | Gold team |
| GTL | Generic Threat Landscape |
| GTP | Gold teaming provider |
| HLS | High-level scenario |
| IT | Information technology |
| NDA | Non-disclosure agreement |
| OSINT | Open-source intelligence |
| PT | Purple team |
| RFP | Request for proposal |
| RT | Red team |
| RTP | Red team provider |
| SOC | Security operations centre |
| TCT | Test Cyber Team |
| TI | Threat intelligence |
| TIBER | Threat intelligence-based ethical red teaming |
| TIP | Threat intelligence provider |
| TLPT | Threat-led penetration testing |
| TM | Test manager |
| TTI | Targeted threat intelligence |
| TTP | Tactics, techniques and procedures used in a cyberattack |

# Annex 2

The main difference between ART and TIBER lies in the potential duration and intensity of the testing process. ART aims for a shorter cycle by limiting the time and resources spent on the threat intelligence and red teaming phases. While both the RT and TI phases can be extended if necessary, the minimum requirements are lower than in TIBER. Second, ART provides entities with the flexibility to shift the focus of the test from red and purple teaming only (network and application security) to physical intrusion exercises and/or gold teaming (incident response). While a minimum level of threat intelligence and red teaming is mandatory, the test scope should be tailored to fit the entity's specific needs and requirements. This results in different options for the threat intelligence and red teaming phases, different levels of purple teaming and the formal introduction of gold teaming.

Compared to "normal" red teaming and internal penetration testing, ART offers the following advantages, in line with the TIBER framework:

- board attention is assured;
- tested systems are live production systems underpinning critical functions;
- ART testing is always intelligence-driven;
- a high level of control is assured by the TCT;
- the test follows an objective and proven framework;
- other competent authorities recognise the test.

| Subject | TIBER | ART |
|---|---|---|
| Duration (indicative) | | |
| Total test duration | 9-12 months | 6-9 months |
| Threat intelligence | 6-8 weeks | 2-8 weeks (modular) |
| Red teaming | 10-12 weeks | 6-12 weeks (modular) |
| Extent of the modules | | |
| Threat intelligence | Full | Modular |
| Number of scenarios | Two plus a Scenario X | A minimum of one scenario |
| RT test | In-through-out fully | Assumed compromise possible |
| Testing on live systems | Yes | Yes |
| Purple teaming | Full | Standard + additional module |
| Gold teaming | To be determined | Optional (modular) |
| Involved parties | | |
| Test Cyber Team | Mandatory | Mandatory |
| Control team | Mandatory | Mandatory |
| Board engagement | Mandatory | Mandatory |
| Red team provider | Mandatory | Mandatory |
| Threat intelligence provider | Mandatory | Optional |

# Annex 3 Adapting ART for use in other sectors

ART for the financial sector has been created by DNB for the Dutch financial sector and its critical third parties. However, cyber resilience does not stop at the borders of a sector or country. Therefore, DNB allows the application of its ART framework and the underlying guides in other sectors and countries. Adjustments to the framework to align it as closely as possible with the circumstances and learning needs of the country or sector are encouraged. However, there are certain conditions attached to the adoption and adaptation of ART

- This document, the "ART framework", contains material to which DNB, the European Central Bank (ECB) and the Bank of England (BoE) own copyrights, as licensed by BoE under the Creative Commons Attribution 4.0 International License (i.e. BoE's CBEST Intelligence-Led Testing document, the "Licensed Material"). This license granted by BoE inter alia contains a disclaimer of warranties. DNB has made changes to the Licensed Material, to which changes DNB owns the copyrights.
- DNB also owns the copyrights to other additions made by DNB as contained in the ART guide. These works are together licensed under the Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0).
- As mentioned, the ART framework and its underlying guides need to be adapted to the specific circumstances and requirements of the country or sector where it is going to be applied. The responsibility for these adaptations lies with the country or sector itself.
- Regardless of the sector or country where ART is to be applied, the framework must always include the following elements:
  - The test must be conducted in secrecy. Only a very limited group (the Control Team) should be aware of the test.
  - The test must be performed on live systems of the financial entity.
  - The test must always include at least 1 scenario involving ethical hacking.
  - The scenario must always be based on threat intelligence.
- ART tests should be supervised by an experienced Test Control Team (TCT) that can operate independently from the tested institutions and the threat intelligence provider/red team provider/gold team provider. Independent supervision significantly contributes to the quality of the test and the comparability of results.
- Recognition of ART tests across different sectors and countries is possible but needs to be formalized between the respective countries and sectors.

# Annex 4 The different modules in more detail

ART offers multiple mandatory and optional modules, which are explained in this annex.

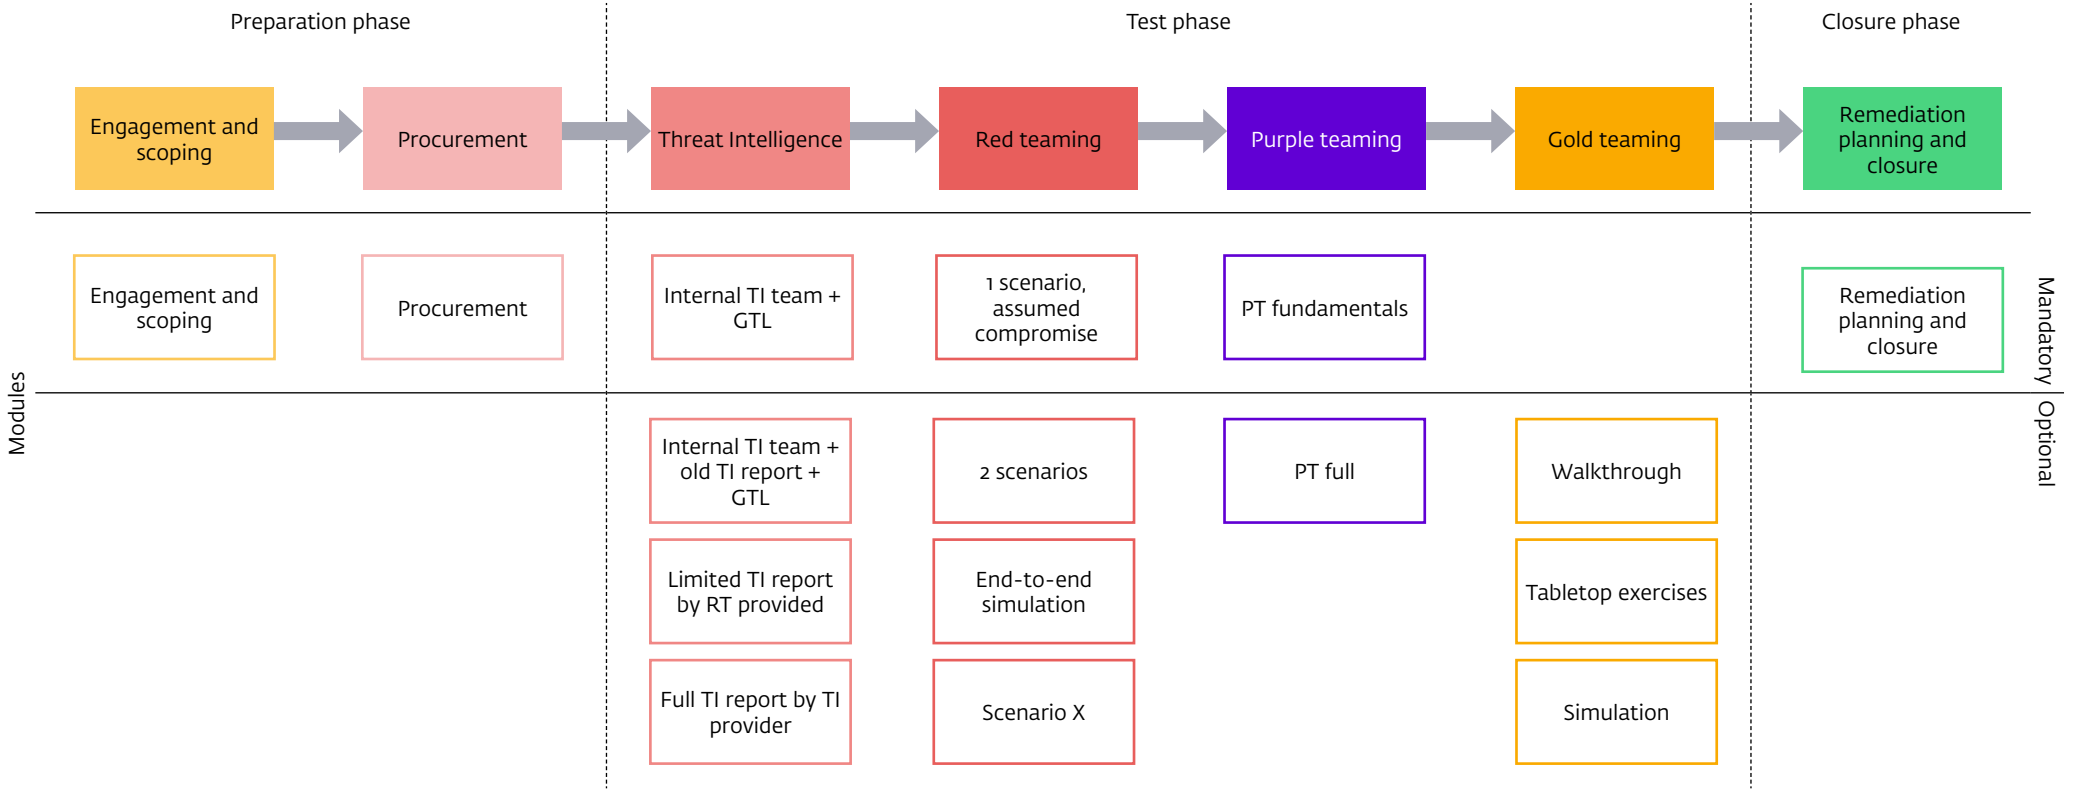

## Threat intelligence modules

**Mandatory:**

1. **Use of the GTL for scenario creation by internal TI experts**
   The Generic Threat Landscape (GTL) is a document that describes the threat landscape of the Dutch financial sector. It is created and updated yearly by the TCT and distributed to the CT during the early stages of the engagement. It shows which threat actors are relevant for Dutch financial entities, reflects on why these actors might attack certain critical functions of an entity and proposes a number of generic scenarios. Based on the intelligence from the GTL, from the GTL, an internal TI expert from the CT can create a basic TI rport that fits can create a scenario that fits the entity or pick a scenario from the GTL to be executed by the RT. If possible and desired, the CTL can involve an internal TI expert in this step to improve the TI scenario(s). The TCT verifies, based on the TI guide, that the TI scenario produced by the CT fits the requirements for an ART test.

   Pros
   + Short lead time; saves time for other phases
   + Minimal investment in terms of people, time and financial resources

   Cons
   - Potentially inaccurate/too generic if the entity does not have a strong intelligence position
   - Requires the entity to invest in creating its own scenario
   - Not all entities have the resources and expertise to create realistic TI-led scenarios
   - Less objective and makes little use of an independent review
   - Results not found due to a limited TI phase can lead to delays in the RT phase

**Optional:**

2. **Use of existing TI report**
   In addition to the GTL, the TI expert within the financial entity can use an existing TI report written specifically for the entity by an external company. This TI report should be no older than 24 months. The TCT and CT should discuss the extent to which the TI report is applicable to the ART test. Based on this TI report (and the GTL), the CT can create a scenario that fits the entity. If possible and desired, the CTL can involve an internal TI expert to improve the TI scenario(s).

   Pros
   + Short lead time; saves time for other phases
   + Intelligence better tailored to the entity than when using just the GTL

   Cons
   - The older the report, the more outdated the intelligence is likely to be
   - Still requires a translation from the CT into actual TI scenarios

3. **Limited target intelligence report and scenario(s) produced by a RTP**
   A number of RT providers are capable of producing TI reports in addition to their RT services. The CT could hire the RTP to write a limited TI report, which the RT attack plan would be based on. This limited TI report would include threat-led intelligence scenarios, but exclude targeted threat intelligence as mentioned in option 4. The TCT must compare the limited TI report to the requirements in the TI guide.

   Pros
   + TI produced by an expert outside party
   + Intelligence more up to date than in an older report (option 2)
   + Intelligence is tailored to the needs of an ART test

   Cons
   - A limited TI report offers less information than a full TI report (option 4)
   - One party producing both TI and RT poses a risk with regard to the separation of duties

4. **Target intelligence and scenario(s) produced by a TIP**
   This option assumes that a professional threat intelligence provider (TIP) is hired to produce a full TI report, similar to TIBER. Reports written by an expert TIP will produce the most added value to the ART engagement, but are more costly and time-consuming. The CT, in collaboration with the TCT, should determine which option is the best fit for the test and the entity.

   Pros
   + Scenarios are tailored to the entity by dedicated professionals
   + The targeted threat intelligence could help the entity before the test has even started

   Cons
   - More costly compared to the other options
   - More time-consuming than the other options

## Red teaming modules

In the red teaming phase (RT phase), the red team provider (RTP) simulates an intelligence-led attack on a specified target (systems and services that underpin one or more critical functions). The RT phase is the main event in an ART test, and it is not optional. Entities can, however, choose the level of intensity they think fits best by adjusting the number of scenarios as well as the duration of the test. The RT phase of an ART test should last at least six weeks for one scenario, or 10 for two or more scenarios. Together with the TM and the provider(s), the CT decide the appropriate duration for the entity.

**Mandatory:**

1. **One scenario and assumed compromise as starting point**
   The most basic, and therefore mandatory, option for the RT phase is a test that includes one scenario and starts with an assumed compromise, thus skipping the in phase of the test. This way, the RTP can focus on identifying and exploiting vulnerabilities that might be missed in a more conventional penetration test.

   Pros
   + The main focus of the test is on the through and out phases
   + Prevents overlap with other (or earlier) in phases
   + Saves time

   Cons
   - Valuable insights might be missed by skipping the in phase

**Optional:**

2. **Two scenarios**
   Given the wide range of potential threats and threat actors that could target an entity, it is beneficial to simulate two (or even more) attack scenarios in one test. It is important to ensure that these scenarios differ significantly and involve a wide range of techniques, tactics and procedures (TTPs).

   Pros
   + Attacks with more TTPs can be simulated, providing a greater learning experience
   + It is more (cost) efficient to test two scenarios in one test than to carry out two separate tests

   Cons
   - The test will likely take longer and cost more
   - The entity will have more issues to remedy

3. **End-to-end simulation**
   Instead of using assumed compromise as the starting point and skipping (part of) the in phase of the test, the entity can intensify the attack by choosing to simulate the entire attack chain from start to finish (also referred to as end-to-end or in-through-out simulation). This can help identify potential weaknesses in the entity's processes and security defences that might be overlooked when assuming compromise. This allows the entity to more accurately identify its greatest security risks.

   Pros
   + The test is more comprehensive and realistic
   + It can provide the entity with a better understanding of its risks and weaknesses

   Cons
   - The test will likely take longer and cost more

4. **Scenario X**
   A Scenario X can be included in addition to the planned scenario(s). The goal of a Scenario X is to emulate attacks that may be expected in the near future. This scenario can be focused on innovative techniques and tactics the RTP (and TIP) expect to emerge, and it may take into account societal developments or developments in the threat landscape that will impact the entity in the future. A Scenario X uses findings from the earlier scenario(s) and hence is developed during the RT phase. The ultimate goal of a Scenario X is to target a critical function, often using a highly creative approach. The use of a Scenario X is decided on and approved by the CT and TCT halfway through the test.

   Pros
   + Can be based on the threat-led scenario using more factual knowledge gained in the test phase
   + Allows the RTP to test based on findings it has gathered in the first weeks of the test

   Cons
   - The test will likely take longer and cost more

## Purple teaming modules

The purple teaming phase (PT phase) is usually conducted after the RT phase of a test has been completed. If the circumstances of the RT phase call for it, the PT phase can start earlier. During the PT phase, the RTP and the BT meet to go over the steps taken by the RT during the test and discuss the measures taken (or not taken) by the BT. The goal of this exercise is to gain a better understanding of each team's techniques and approaches and to identify any weaknesses in the entity's defences.

**Mandatory:**

1. **Purple teaming fundamentals**
   The most basic, and therefore mandatory, option for the PT phase is a one-day purple teaming exercise. During PT fundamentals, the RT and BT share intelligence, review the simulated attacks and analyse the findings, before proposing ways to improve the entity's defences. A PT fundamentals option is suitable for ART tests where the RT phase was relatively compact. If multiple scenarios are tested, PT fundamentals will not suffice. Towards the end of the RT phase, the TM and the CT decide whether PT fundamentals is adequate or not.

   For more information on purple teaming, please see the purple teaming guide.

   Pros
   + For smaller tests, one PT day might be enough
   + Saves time and money compared to PT extended

   Cons
   - Participants may learn less and there may not be as much cross-pollination

**Optional:**

2. **Purple teaming extended**
   The extended option for the PT phase is a purple teaming exercise of more than one day. During PT, the RT and BT share more intelligence and review the simulated attacks and analyse the findings in greater depth, before proposing ways to improve the entity's defences.

   Pros
   + Participants will learn more and there will be more cross-pollination

   Cons
   - The PT phase will take longer and cost more
   - For smaller tests, PT extended might not be needed

## Gold teaming modules

**Optional:**

**Walkthrough session**

This is a low-key and accessible GT variant. It can be used by entities with limited experience in crisis management. A walkthrough session can also be a good fit for entities that have recently seen significant changes in their crisis management structure and personnel. During a walkthrough session, all steps of the crisis management process, from detection to closure, are discussed and completed in detail. A walkthrough session addresses the roles and expectations of CMT members, as well as actions and measures to take in specific crisis scenarios. It is a discussion-based session with the purpose of validating the crisis management processes and increasing the CMT's knowledge of how to act if a specific crisis unfolds. The walkthrough is facilitated by a process supervisor or facilitator.

Pros
+ The entity gains insight into its crisis management processes
+ Module can be used by all entities

Cons
- The learning experience will be less thorough than with the other GT options

**Tabletop exercise**

This GT variant is an accessible discussion-based exercise and a good fit for entities with a CMT that already has some experience in crisis management but that do not want to submit their CMT to a full simulation. The goal of a tabletop exercise is to practice crisis management in a low-stress environment. The CMT members gain knowledge and skills on an individual level, but the exercise also trains their ability to collectively respond to challenges and work effectively as a team. During tabletop exercises, participants are trained in specific crisis management capabilities, such as gaining situational awareness, communicating actions and statements, information management and decision-making (depending on the exercise goals). At the start of the tabletop exercise, all participants receive the same scenario information. After this, the exercise starts, and more role-specific information can be shared with individual participants. A tabletop exercise is always facilitated by a facilitator, trainer and/or observer for training and evaluation purposes.

Pros
+ The entity trains its crisis management processes
+ The entity practises in a team setting
+ The module tests the adaptability of the CMT

Cons
- Only for entities with some experience in crisis management
- Not as realistic as a full simulation

**Full simulation**

The simulation is the most elaborate and challenging GT variant. It is intended for experienced crisis teams that want to test their capabilities in an "as real as possible" situation. In an interactive crisis simulation, the entity's crisis management team experiences what it is like to be confronted with a real crisis. The goal of a simulation is to practise and train crisis management capabilities (based on learning goals) under stress, by confronting team members with a realistic crisis scenario unfolding in real time. Under time pressure, the CMT members must decide what to do to mitigate the impact of the crisis. Scenario information and events (called injects, such as phone calls, emails, messages on social media channels and the news) can be inserted in multiple ways. There are two subvariants:

a. **Simulation without counterplay** In this subvariant, static or pre-defined scenario injects are sent to the exercise participants through various (fictitious) channels from the exercise control cell;
b. **Simulation using counterplay** This subvariant uses dynamic scenario injects that are sent to the participants from the exercise control cell. Counterplay events are based on actions performed by the CMT, in order to make the exercise more realistic.

For both subvariants, the simulation set-up includes an exercise bubble containing the exercise participants, facilitator, trainer and observer. The scenario injects and/or responses are sent to the

exercise bubble from the exercise control cell, which is located in a separate room or location. Please note that a successful simulation requires profound and meticulous preparation, execution and evaluation.

Pros

+ Optimal learning experience for participants and highest level of cross-pollination
+ As close as an entity can get to "the real deal"

Cons

- A full simulation takes longer and is more costly
- The risk of escalation is slightly higher

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

Follow us on:

DeNederlandscheBank
EUROSYSTEEM