Advanced Red Teaming (ART) Framework

for the financial sector

DeNederlandscheBank

EUROSYSTEEM

Contents

1 Introduction: start with ART

1.1 What is ART?

The Advanced Red Teaming (ART) framework is a comprehensive framework that empowers a wide range of financial institutions and their ICT third party service providers to conduct advanced ethical hacking activities, called red teaming, driven by high-level threat intelligence. As an evolution of the Threat Intelligence Based Ethical Red teaming (TIBER) framework, ART gives participating institutions the freedom to select and customise various modules, ensuring that each test aligns with their specific learning objectives and cybersecurity posture. These modules address different facets of cyber resilience testing such as network and application security, endpoint security, data exfiltration, ransomware deployment, but also institutions' crisis management response to such attacks.

By enabling financial institutions to choose the modules that best fit their cybersecurity posture, maturity level and available resources, ART allows them to optimise their cybersecurity efforts and investments. This approach results in a tailored ART test that aligns with their unique learning goals. ART's modular nature offers participating institutions flexibility and value, reducing documentation requirements while upholding the rigorous cybersecurity testing standards associated with TIBER-EU. If all mandatory steps and deliverables have been followed, and if the managing authority (De Nederlandsche Bank - DNB) concurs, a formal attestation will be issued that the test met the requirements of the ART framework.

1.2 Who is ART for?

The fundamentals of red teaming (RT) 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. However, detailed implementation may vary depending on the nature of the sector. This specific ART framework focuses on the financial sector. There are four groups of potential participants in this sector:

1. Institutions 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.
2. Institutions that have a higher level of cyber maturity, but are not subject to the Threat Led Penetration Testing (TLPT) requirement under the Digital Operational Resilience Act (DORA).
3. Institutions that already perform a TIBER or TLPT test, but are looking for a more frequent testing framework.
4. Non-financial institutions that provide services, software and systems that are critical to the functioning and stability of the financial system.

Note: under specific circumstances, the ART framework can be used to meet the DORA requirements (Article 24).

1.3 Disclaimer and legal

An ART test for the financial sector that is followed by a DNB attestation, can only be conducted by financial institutions and their ICT third party service providers under the guidance of the DNB's Test Cyber Team (TCT-DNB). The financial institution and TCT-DNB draw up a contract which specifies the chosen modules, risks, responsibilities and fees for the duration of the test.

The information and opinions expressed in this document are for information purposes only. They are not intended to constitute legal or other professional advice and should not be relied on or treated as a substitute for specific advice relevant to particular circumstances. The sponsors and authors of this document accept no responsibility for any errors, omissions or misleading statements in this text, or for any loss that may arise from reliance on the information and opinions expressed within it. 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 guides. These works are together licensed under the Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0).

1.4 Abbreviations

For alignment with the DORA TLPT framework (TIBER), similar definitions and abbreviations are used in this framework.

Abbreviation	Meaning
ART	advanced red teaming
BOD	board of directors, also referred to as the executive board
BT	blue team
CMT	crisis management team
CIF	critical or important function
CT	control team
CTL	control team lead
GT	gold teaming
GTL	generic threat landscape
GTTP	gold team test plan
GTP	gold team provider
LPT	limited purple teaming
PT	purple teaming
RFP	request for proposal
RT	red teaming
RTP	red team provider
RTTP	red team test plan
RTTR	red team test report
SOC	security operations centre
SSD	scope specification document
TCT	test cyber team
TI	threat intelligence
TIBER	threat intelligence based ethical red teaming
TIR	threat intelligence report
TIP	threat intelligence provider
TM	test manager
TPSP	third party service provider
TSR	test summary report
TTI	targeted threat intelligence
TTIR	targeted threat intelligence report
TTP	tactics, techniques and procedures

2 ART overview

The objective of this section is to provide an overview of (1) the key participants 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 Sections 3, 4 and 5.

2.1 Key participants

An ART test involves many different parties, each with their 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 Section 3.1 and the separate guides supporting the ART framework.

Name	Participation	Role
control team lead (CTL)	Mandatory	Test owner. The CTL has the final say in key decisions, for instance regarding the scope of the test, go/no go decisions and learning goals. The CTL is also responsible for various practical matters, such as meetings and planning.
control team (CT)	Mandatory	Team to support the CTL in delivering the deliverables and gathering the necessary information for the preparation, execution and closure of the test.
threat intelligence provider (TIP)	Mandatory	Responsible for providing threat intelligence and custom scenarios preceding the red teaming phase. The threat intelligence can be provided by the institution's own internal threat intelligence resources or can be purchased from an external provider. The TIP and RTP can be the same organisation.
red team provider (RTP)	Mandatory	The RTP is responsible for the actual red teaming based on test scenarios developed earlier in the test.
gold team provider (GTP)	Optional	Responsible for developing, facilitating and evaluating the gold teaming (GT) session, based on the results of the red team test. The presence of a GTP is optional, and only required if this module is selected. The GTP can be the same organisation as the RTP.
blue team (BT)	Mandatory	The team responsible for the institution's cyber defence. It should not be aware of the test until it has been completed.
board of directors (BOD)	Mandatory	At least one of the board members, or a formal delegate – the so-called C-level sponsor – is part of the CT and must formally approve deliverables and/or take decisions as required during the test.
test cyber team (TCT)	Mandatory	The TCT ensures that the test is performed in a uniform and controlled manner, in accordance with the requirements of the ART framework. The members of the TCT are provided by DNB.

2.2 Key steps and milestones of the ART phases and modules

The ART test comprises three distinct phases: preparation phase, test phase and closure phase as illustrated in Figure 1 below.

The phases are characterised by specific prerequisites that must be met before progressing to the subsequent phase. This section provides an overview of the various steps involved, the key milestones to be achieved and the average time allocation for each phase.

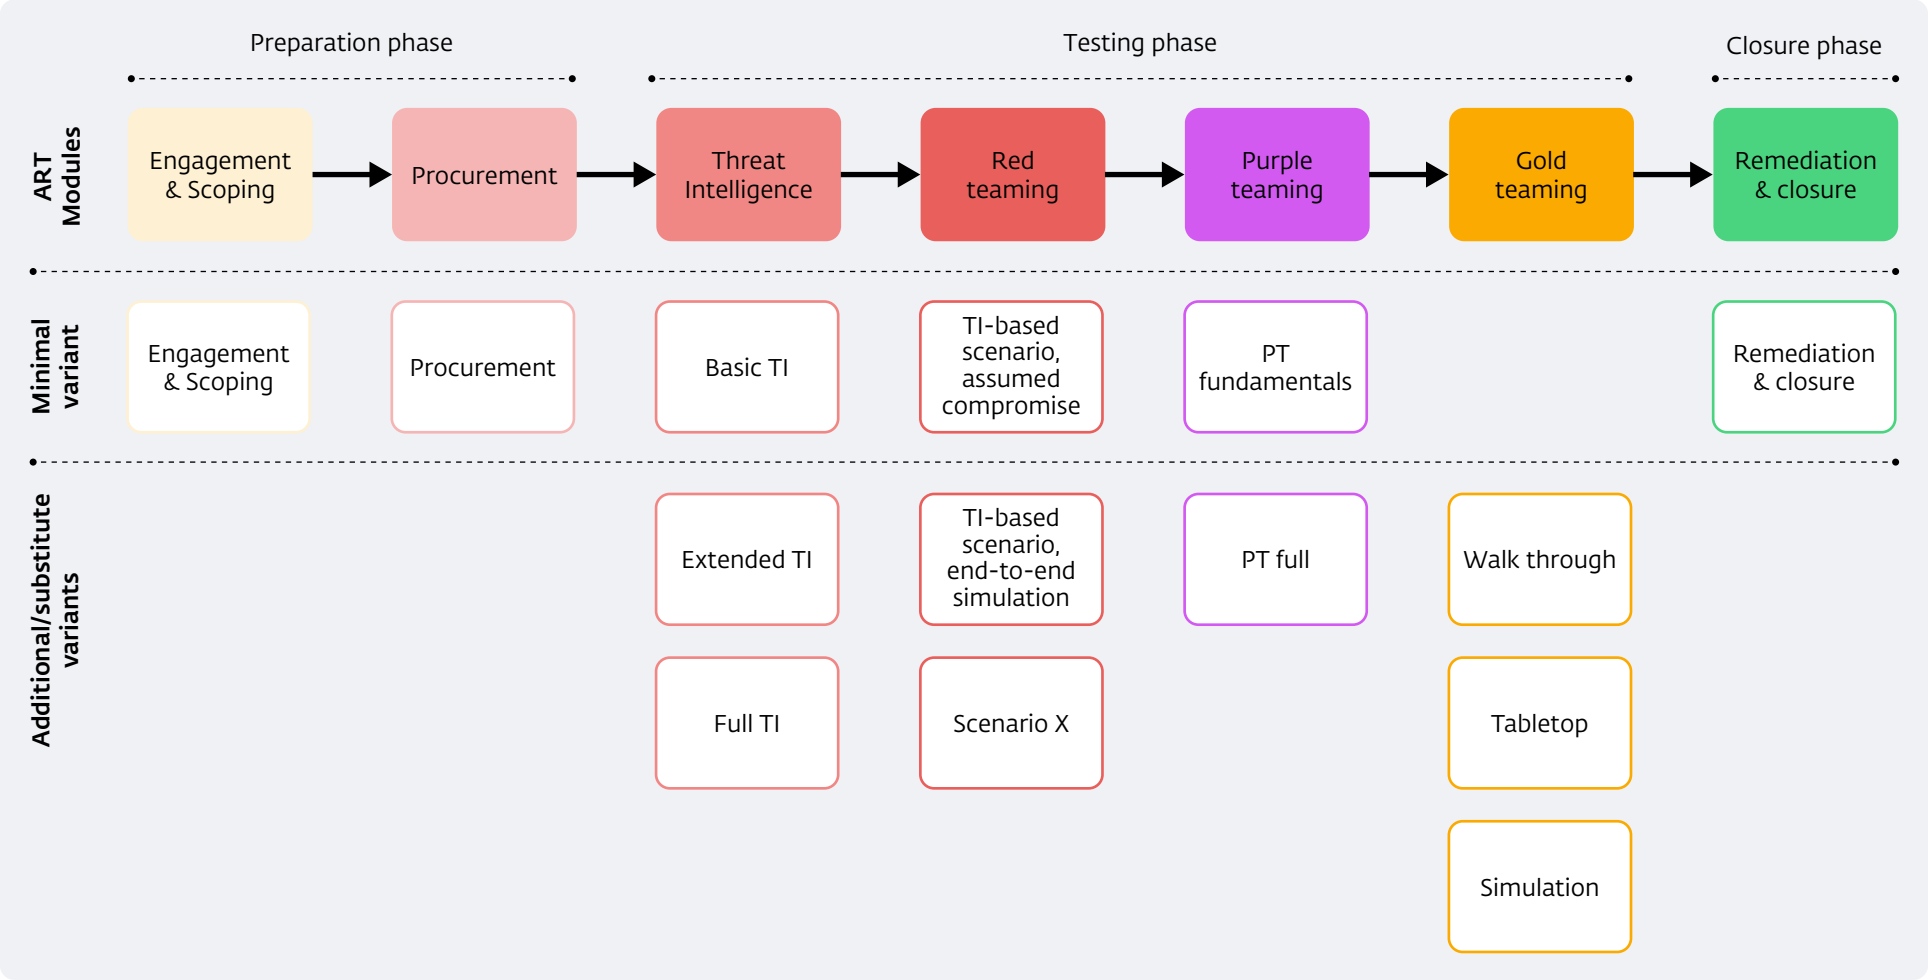

ART phases and modules

Engagement and scoping

What is it?
In the engagement and scoping phase, the TCT and the CTL collaborate to define the parameters of the ART test. This involves identifying the institution's learning objectives, selecting the appropriate modules and their variants, determining the composition of the CT, scheduling the test, assessing the current level of cyber resilience and ascertaining whether the institution has a comprehensive overview of its CIFs, processes and systems.

Milestones and main deliverables
- C-level commitment;
- Agreement on the selected modules and their variants for the test;
- Signed contract between the institution and DNB for guiding the test;
- A filled in and approved SSD.

Average duration: 8-12 weeks

Procurement

What is it?
In the procurement phase, the CTL reaches out to several TI, RT and/or GT providers to request quotations for the ART test as defined in the engagement and scoping phase. The ART service 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 depending on the complexity of the institution's procurement procedures. Even though the modules and variants are only agreed upon during the engagement and scoping, it is recommended to prepare for the procurement phase as early as possible. Depending on the institution's procurement procedures this phase could take longer than the average 6-8 weeks.

Milestones and main deliverables
- Successful procurement procedure based on the test requirements;
- Signed contract between the institution and the RTP (and external TIP and GTP, if applicable);
- Leg-up inventory and preparation.

Average duration: 6-8 weeks

Threat intelligence

What is it?
In this phase the TIP – which can be an internal TI resource, the RTP or an external TIP – formulates one or more threat intelligence-based scenarios that will form the basis for the RTTP and the execution of the test.
The required timeframe, scope and depth of the research, as well as the resulting TIR/TTIR depend on the number of scenarios selected and the specific TI variant chosen.

Milestones and main deliverables
- A business overview meeting (optional) where the institution provides information about its CIFs, the supporting IT systems, business processes, customers and other relevant developments for the test;
- Delivery of TIR/TTIR with one or more TI-based scenarios that serve as the foundation for the RTTP. These are a result of the performed TI research and analysis;
- The approval of the TIR/TTIR and a formal 'go' during the go/no go meeting, formalising the start of the red teaming in the test phase.

Average duration: Between 2 and 8 weeks

Red teaming

What is it?
In this phase, the RTP translates the TI-based scenario(s) into a practical RTTP, structured according to the MITRE ATT&CK framework. Once the C-level sponsor and the CTL have given their final approval to the RTTP, the RTP proceeds to execute the actual red teaming activities of the test. The duration of this phase varies based on the specific variants selected for the red teaming part of the test.

Milestones and main deliverables
- A draft version of the RTTP;
- An actionable and formalised RTTP;
- A go/no go meeting where all parties involved formally approve the RTTP;
- Reaching the RT's predetermined flags.

Average duration: Between 6 and 12 weeks

Purple teaming

What is it?
During purple teaming (PT), the RT and the BT set up a collaborative workshop session where they discuss the scenarios executed step by step. The execution of the PT can be based on the RTTR and the BT observations. The two teams work together to share insights in weaknesses detected and attack paths used during the test. The goal of this collaboration is to enable learning and enhance the institution's cybersecurity posture. The duration of the PT session depends on the selected PT variant. It is important to note that the timing of the execution of the PT is determined by the characteristics of the specific ART test. If the ART test also involves GT, the sequence of the PT and GT might be customised to the learning goals of the institution.

Milestones and main deliverables
- The RTTR in which the RT describes in detail which actions it has taken during the active RT phase;
- Hosting the actual PT session with the BT and the RT.

Average duration: Around 2 weeks

Gold teaming

What is it?
GT is an optional module which tests the crisis management capabilities of the institution and which yields learnings to increase this capability. In this module the GTP continues from where the RT and/or PT concludes. In this phase, the consequences of the simulated cyberattack scenario are elevated to CMT level to assess how the institution's crisis management structure responds to a cyber crisis. The inclusion of the GT module allows the institution 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 GT exercise is determined by the specific variant selected. It is important to note that the timing of the execution of the GT is determined by the characteristics of the specific ART test. If the ART test also involves GT, the sequence of the PT and the GT might be customised to the learning goals of the institution.

Milestones and main deliverables
- A GT kick-off meeting during which the scope, learning goals and set-up of the GT phase are defined;
- A GTTP based on the chosen variant and selected red teaming scenarios;
- The formal approval of the GTTP during the GT go/no go meeting, formalising the start of the GT phase;
- Successful execution of a dry run (tabletop exercise and simulation variants only);
- Execution of the GT exercise;
- Delivery of an action list, observation report or evaluation report, i.e. the outcome of an evaluation including points for improvement. The content of this report depends on the GT variant selected.

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 institution drafts the test summary report (TSR), including a plan to address and resolve the vulnerabilities, weaknesses and observations identified during the ART test. 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 requirements – the TCT provides the attestation document.

Milestones and main deliverables
- Delivery of the TSR, including the recommendations for improvements;
- Cleaning up of all traces of the test;
- A 360 feedback session;
- A 360 feedback report;
- Delivery of the attestation document;
- A board meeting, where the outcomes of the ART test are reported to the BOD. If desired, the TCT can also attend this meeting;
- Remediation plan, where the remediation actions are detailed. The TCT is not involved in the creation of the remediation plan, but encourages the institution to share their remediation plan with their supervisors after the ART test.

Average duration: Around 2 weeks

2.3 Key mandatory and optional modules and their variants

ART offers both mandatory and optional modules consisting of a minimum variant and additional or substitute variants. This section provides a high-level overview of these modules. For a more extensive explanation of each module and their variants, please consult the specific ART guides.

Why are there mandatory and optional modules?
Financial institutions 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 approach to define the scope of the ART test. Some ART test modules are mandatory, to ensure that the test meets the minimum standards. Besides these mandatory modules, institutions may to choose from several optional modules and their variants to make the ART test fit their budget, learning goals and security posture.

Who chooses the modules?
In the earliest stages of an ART test, before procurement, the CTL and the TCT meet to discuss a number of topics (see Section 3). An important topic is the scope of the ART test. Although the decision as to which modules to include ultimately lies with the institution itself, the process leading up to this decision is a collaboration between the CTL and the TCT. The following factors play a role in this process:

- The institution's characteristics, such as size, cybersecurity posture and internal TI and crisis management capabilities;
- How this test relates to other forms of cyber security tests, either already completed or in the pipeline;
- The institution's budget;
- The institution's previous experience with threat-led penetration testing;
- The institution's learning objectives;
- The institution's ambitions.

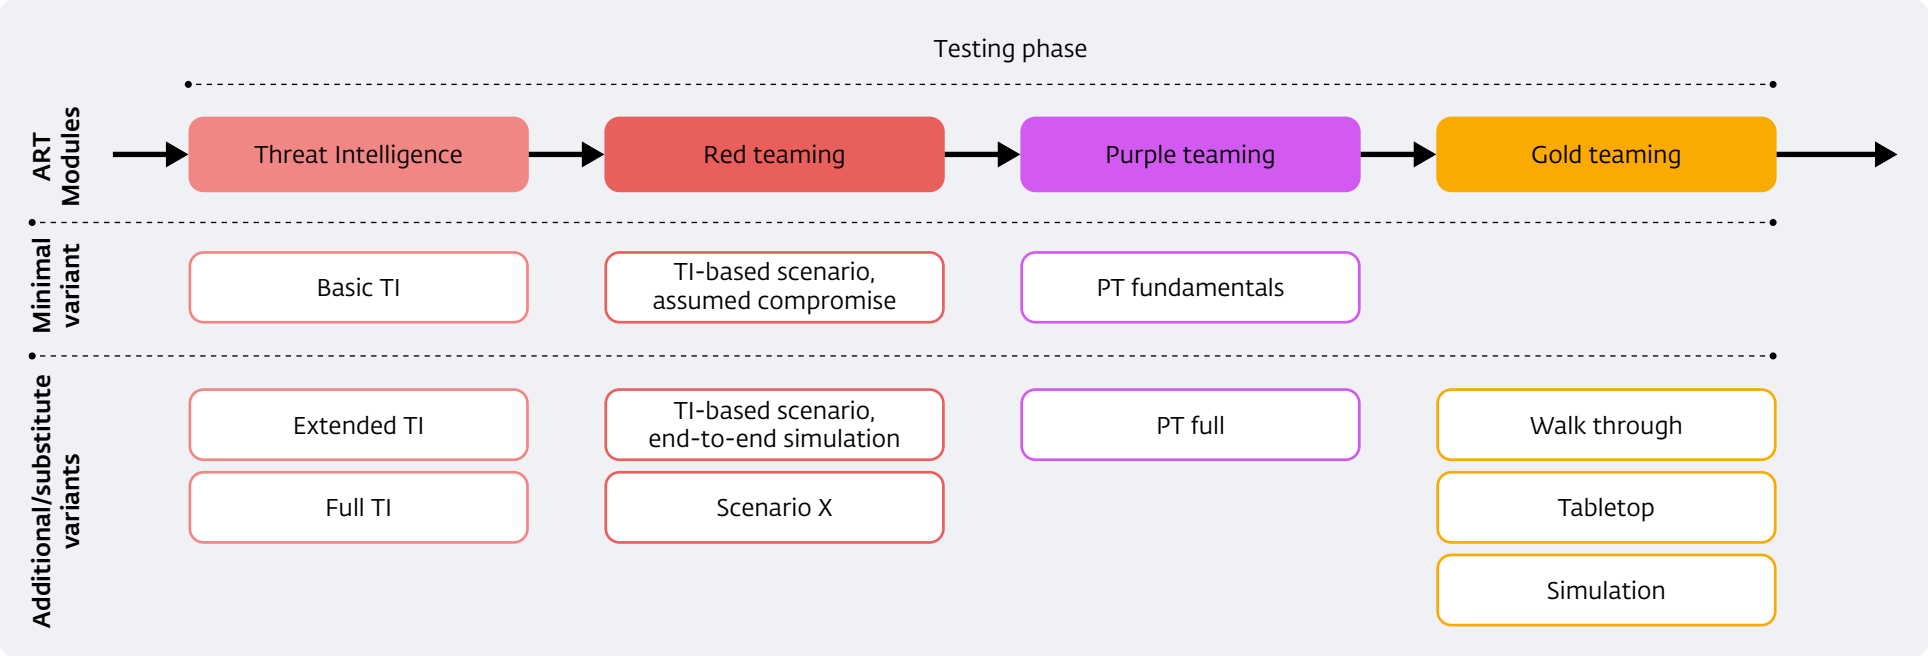

Threat intelligence

Basic TI

Basic threat intelligence is the minimum TI variant in ART. The main sources of input for the basic TI are the generic threat landscape (GTL), the SSD and – if available – TI reports from previous tests or analyses that specifically describe the institution that is being tested. The GTL is an intelligence report that is updated by TCT-DNB on an annual basis. It contains the outcomes of a threat assessment of the Dutch financial sector and provides the most relevant high-level scenarios and threat actors. While it is not custom-tailored to any specific organisation, it can serve as a starting point for the threat intelligence phase of any chosen threat intelligence module within ART.

Information from a previous TI report, the SSD or the GTL cannot be copied directly into the basic threat intelligence report (TIR). Although the TI team in this variant does not need to access additional TI sources, they do need to perform an analysis on the intelligence available to see (1) how the different sources fit together, (2) which information is relevant for the test and why and (3) how the information fits into a logical report. The resulting TIR includes a mandatory executive summary, business overview, limited threat landscape, actor selection and possible scenarios.

Internal TI resources can execute this variant if they meet the requirements as set out in the ART TI guide. Otherwise, external providers like TIPs or RTPs may be engaged, depending on their capabilities. Details can be found in the ART TI guide and ART service procurement guide. This variant can be fast and cost-efficient, but it lacks a tailored, in-depth view of the organisation, which may lead to overlooking key threats.

Extended TI

In the extended TI variant, the GTL, the SSD and a previous TIR (if available) are still the primary sources. The main difference with basic TI is that these sources now need to be enriched with current data from independent sources and expert insights that the TI team (the individuals conducting the TI for this test) gathers. As in the basic TI, the extended TIR contains the following mandatory elements: executive summary, business overview, threat landscape, actor selection and possible scenarios. All elements are expected to more elaborate, have more of an analytic nature and be more detailed than the basic TI variant.

Both internal resources or an external RTP or TIP may perform this work, depending on their capabilities. Institutions should consult the ART TI guide and ART service procurement guide when sourcing external support. The extended TI variant offers a broader and more current (external) intelligence perspective on the financial institution without going all-in with a full TTIR that is potentially not fully utilised by the institution.

Full TI

In the full TI variant, a specialised external provider (TIP or RTP with TI capacity) must be hired. The TI team also uses the GTL, SSD and previous TIRs (if available) but in this variant, these are not the main sources of intelligence. The TI team is expected to have an extensive intelligence position with own and external TI sources. The goal is to create a highly accurate and current threat landscape, detailing likely threat actors and their tactics, techniques and procedures (TTPs). Additionally in this variant, the provider conducts a targeted threat intelligence (TTI) analysis of the institution's attack surface, including exploitable public information such as look-alike domains or leaked data.

As in the basic TI and extended TI, the full TI variant has the same mandatory elements in the TIR, i.e. an executive summary, business overview, threat landscape, actor selection and possible scenarios. However, in this variant, the full TIR includes more actors (longlist), more scenarios, a deeper and more detailed analysis and an additional TTIR section.

Though this is the most comprehensive and expensive variant, it offers the most valuable insights for the institution and the RT. Procurement of the TI team must align with the requirements described in the ART TI guide and ART service procurement guide.

Most likely real-life application of the TI variants

	Internal resource	RTP	TIP
Basic TI	most realistic option, pending capabilities	most realistic option, pending capabilities	allowed, but not likely in practice
Extended TI	allowed, pending capabilities	most realistic option, pending capabilities	allowed, but not likely in practice
Full TI	not allowed	allowed, pending capabilities	most realistic option, pending capabilities

Red teaming

TI-based scenario, assumed compromise

The active RT phase of an ART test consists of at least one TI-based scenario where the in-phase can be skipped. This is called the 'assumed compromise' scenario and is the minimum variant of the RT module. If desired, the skipped in-phase can still be simulated later, such as at the end of the RT phase or during the PT phase. For example, it is possible to efficiently simulate the in-phase after the through- and out-phases with the knowledge gained during these latter two phases.

TI-based scenario, end-to-end simulation

To make an ART test as realistic as possible, an institution can choose to incorporate all the steps of the TI scenario into the actual red teaming. This means that the in-, through- and out-phases are all fully simulated during the RT phase in the same order in which they are realistically executed in a real-life situation. This method of RT generally requires more time and resources, but it provides the most comprehensive view of the institution's cyber resilience across all aspects.

Scenario X

A scenario X variant can only be included in addition to the other variant(s). Scenario X it not a full substitute for an assumed compromise or end-to-end scenario. The goal of a scenario X is to emulate attacks that may be expected in the near future or that are solely based on a specific learning goal of the institution. This scenario can focus on innovative techniques and emerging tactics. Scenario X can use observations from scenarios that have been executed earlier in the test, meaning it can be defined during the RT phase. If based on the learning goals of the institution, the scenario can be defined upfront together with the definition of the TI-based scenarios. The ultimate goal of scenario X is to target a CIF, often by using a highly creative approach.

Configuring the RT module

All of the above variants can be combined to define the full size of the RT module. A financial institution can choose to only execute one scenario in the RT phase, but it is also possible to execute multiple scenarios. However, when selecting only one scenario to be performed, the assumed compromise variant should be seen as the minimum variant. Performing only scenario X is not allowed. If there is enough diversity in actors, TTPs and objectives in the different scenarios, valuable additional learnings can be achieved from executing a second or third scenario.

Purple teaming

PT fundamentals

The minimum variant for the PT phase is a one-day PT exercise: the PT fundamentals. During this exercise the RT and BT share intelligence, replay and/or review the simulated attacks and analyse the observations. The RT also proposes ways to improve the institution's defences. A PT fundamentals variant is suitable for ART tests with a relatively compact RT phase.

PT full

The full variant is a PT exercise of more than one day. During this variant of the PT module, the RT and BT share more intelligence and review the simulated attacks and analyse the observations in greater depth, before proposing ways to improve the institution's defences. This variant is better suiting ART tests where 2 or more scenarios have been executed during the RT phase.

Gold teaming

Walk-through session

The walk-through session is the most low-key and accessible GT variant in ART. It is therefore the minimum variant when choosing this optional module. It can be used by institutions with no or very limited experience in crisis management. A walk-through session can also be a good choice for institutions that have recently had significant changes in their crisis management structure and personnel. A walk-through session is a discussion-based meeting aimed at validating plans, processes and procedures.

Tabletop exercise

The tabletop exercise is an accessible, discussion-based GT variant. It is a good choice for institutions 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 practice crisis management capabilities (based on learning goals) in a low-stress environment.

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 practice crisis management capabilities (based on learning goals) under stress, by confronting team members with a realistically simulated scenario as performed in the RT activities, unfolding in real time.

2.4 Key documentation

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

- ART framework
- ART service procurement guide
- ART control team guide
- ART threat intelligence guide
- ART red teaming guide
- ART purple teaming guide
- ART gold teaming guide
- ART quality assurance checklist
- ART scope specification document
- ART 360 feedback format
- ART test summary report format

These documents can be found at
Advanced Red Teaming (ART) | De Nederlandsche Bank

3 Organising an ART test

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

The most important stakeholders in a test are the:
- Control team (CT), including the control team lead (CTL);
- Threat intelligence provider (TIP);
- Red team provider (RTP);
- Gold team provider (GTP) (if the GT module is selected);
- Institution's blue team (BT);
- Institution's board of directors (BOD) and it's representative, being the C-level sponsor;
- Test Cyber Team (TCT) from DNB.

3.1 Key participants in an ART test

Control team, including the control team lead
The CT is the team that manages the institution's involvement in the test. The CT members are the only staff within the institution who are fully aware of the test. The CT consists of a CTL and a mandatory substitute, subject matter experts and, if necessary, representative(s) of ICT third party service providers. If possible and desired, and depending on the chosen TI variant, the CTL can involve an internal TI expert in the TI phase of the test to improve the TI scenario(s). When the GT module is chosen, the CT should also contain a GT lead. At least one C-level sponsor, being a board member or a formal delegate, is also part of the CT, but does not necessarily receive daily updates. This person is informed on important developments by the CTL.

For more information on the CT, please consult:
Advanced Red Teaming (ART) | De Nederlandsche Bank

Threat intelligence provider
Depending on the TI variant selected, the TIP may be internal threat intelligence experts, an RTP with the capability to produce threat intelligence, or a specialised external threat intelligence provider. In any of these cases, the TIP is responsible for providing (targeted) threat intelligence during the test phase and, if necessary, provides additional intelligence during the RT phase. The main product delivered by the TIP is the TIR as appropriate to the TI variant selected.

For more information on the TIP, please consult:
Advanced Red Teaming (ART) | De Nederlandsche Bank

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

For more information on the RTP, please consult:
Advanced Red Teaming (ART) | De Nederlandsche Bank

Gold teaming provider
The GTP is responsible for developing, facilitating and evaluating the selected GT exercise, based on the scenario used during the RT phase of the test. The GTP can be the same as the RTP, but only if the RTP has the required qualifications to organise and execute this component of the test. The main products delivered by the GTP are the gold team test plan (GTTP), a crisis scenario, exercise materials and the GT report (for the tabletop exercise and simulation variants).

For more information on the GTP, please consult:
Advanced Red Teaming (ART) | De Nederlandsche Bank

Blue team
The blue team (BT) is the institution's defensive team. This is usually a security operations centre (SOC), but it can also be another organisational unit. The BT should not be made aware of the test until the active RT is finished. However, a situation may arise where the BT finds out about the test (or parts of the test) before it has been completed. After the red teaming is over, the BT can be made fully aware of the test. Together with the RT, it will evaluate the observations and expand the learning experience during the PT session. Besides technical personnel, such as SOC staff and IT administrators, the PT is relevant to personnel that was not part of the CT, but played a role during the test. This can vary from staff that received phishing emails to personnel whose accounts might have been compromised during the test.

Institution's board of directors
Throughout the test, the board of directors (BOD) plays an important role in a number of ways. At least one of the board members – or a formal delegate of the BOD – is part of the CT 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 C-level sponsor involved and up to date during the test. The other board members are not aware of the test and thus are only involved during the closure and learning phase. This can either be during the PT or GT sessions, or when the test is finished. After each test, the CT and the BOD must allocate time to allow the CT to present the observations and proposed improvements.

Test Cyber Team
The role of the Test Cyber Team (TCT) is to make sure that institutions are tested in a uniform and controlled manner, in accordance with the requirements of the ART framework. The TCT appoints a test manager (TM) and a backup TM for each test who works closely together with the institution'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. The TM has the right to escalate major deviations from the test scope or scenario to the TCT's manager, who the TM directly reports to. Next to test management capacity, the TCT also has TI and GT capabilities. These resources play an active role during the TI phase or GT exercise of the test phase by supporting the TM during the validation of the TI deliverables. The TI resources also provide the annual GTL supporting the TI for a test.

3.2 The control team lead's project and risk management responsibilities

The CTL is responsible for managing the ART test as a project along with the accompanying risks. This means that the CTL is, among other tasks, responsible for planning the mandatory meetings, agreeing on communication methods, managing escalations during the test, keeping track of risks, drafting a high-level overall planning for the entire test and sharing it with those involved. Project management also involves ensuring that internal stakeholders, such as the C-level sponsor, are involved in the test in a timely manner and that the RTP and TIP deliver according to the contractual agreements and planning.

3.3 Overview of mandatory documents

One of the objectives of ART is to limit mandatory documentation where possible. A lower documentation load prevents an unnecessary burden on the parties involved. Nevertheless, a certain degree of documentation is essential, for example because it forms the foundation for improvements within the tested institution. The following documentation needs to be produced or maintained during an ART test:

Name	Author	Goal
ART contract	TCT and CT	The ART contract between TCT and the institution outlines the legal foundation, the scope of the ART test, the statement of work and other procedural agreements that underlie the test. This contract is the basis for the invoicing of the TCT support activities.
Risk register	CTL	The risk register is used to identify and manage risks that may arise from performing the ART test as a whole. This log could contain both the organisational risks for the institution itself, as well as the technical risks when performing the RT activities. It can also be split, depending on contractual agreements, where the CTL documents and manages the organisational risks of the ART test and the RTP documents and manages the technical risks during the active RT.
Scope specification document (SSD)	CT	A document that identifies vital systems, business processes and assets that will be included in the test. The SSD ensures that the CT, TCT, RTP and TIP have a clear understanding of key areas, systems and processes, and that they stay within the predefined scope.
Threat intelligence report (TIR)	TIP	The TIR is a document that outlines the financial institution's threat landscape. It also identifies risks and crown jewels and offers a business overview and attack scenarios. When the full TI variant is selected, this TIR should also contain the information regarding the TTI analysis. Based on the TIR, the RTP makes its RTTP.
RT test plan (RTTP)	RTP	A document in which the TI-based scenario(s) and the optional scenario X are translated into a technical plan, complete with TTPs of the selected threat actors using MITRE. Additionally, the document should include descriptions of potential leg-ups, risk management for the technical risks during RT and the expected timeline for execution.
RT test report (RTTR)	RTP	An RTTR is a comprehensive document that sets out the observations and insights from the RT phase.
GT test plan (if GT module is selected)	GTP	A GT test plan describes the approach for the preparation and execution of the GT. It includes the scope, learning goals, high level scenario and risk management of the GT exercise.
GT MSEL (if GT module is selected)	GTP	The MSEL is a chronological timeline of scenario information and scripted events that will be injected into the TTX or SIM. This MSEL will be validated during the dry run and used for the execution of the TTX or SIM.
GT report (if GT module is selected)	GTP	A GT report is a comprehensive document that sets out the observations and insights from the GT phase. The GT report can be an overview of actions or improvement points (in the walk-through), an observation report (for a table top exercise) or an evaluation report (in case of a simulation).
Test summary report (TSR)	CT	The TSR is a concise document that provides an overview of the key observations, 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.
360 feedback report	TCT	A summary of the key observations from the evaluation of the ART test.
Attestation document	TCT	A document attesting that the ART test has been performed and the documentation has been delivered in accordance with the ART framework. The attestation document is not an indication of the quality of the institution's defences or the RTP's capabilities. This document can only be delivered once all steps of the framework have been finalised and all documents have been delivered according to the quality standards.

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 Sections 4, 5 and 6.

Name	Attendees	Objective(s)
Initiation meeting(s)	TCT and CT	Meeting(s) between the CT and the TCT before the formal start of the ART test in which the scope, the modules and other fundamental prerequisites are discussed and contracted between the TCT and the institution. The discussion on these topics can be held during one or more meetings where the documentation is iteratively discussed and provided with feedback.
Scoping meeting	TCT, CT (including the C-level sponsor) and, if procured, the TIP and RTP	During the (final) scoping meeting, the SSD is agreed by the TCT and the institution's C-level sponsor. The scoping meeting can be combined with the launch meeting.
Launch meeting	TCT, CT, TIP and RTP	Formal launch of the ART test. During the launch meeting, the following topics are discussed: ■ The ART process and documentation; ■ Other TCT members involved; ■ Stakeholders, roles and responsibilities; ■ The overall project planning and risk management. The launch meeting can be combined with the scoping meeting. The end of this meeting marks the formal start of the ART test phase.
Business overview workshop (optional)	TCT, CT, TIP and RTP	Workshop given by the institution's business expert to support the RTP/TIP in its understanding of the institution. This workshop can optionally be held when starting the activities in the test phase, e.g. as a separate meeting at the start of the TI phase.
Weekly update meetings during the test phase	TCT, CT, TIP (during TI phase) and RTP (during RT phase)	During the entire test phase, there are – at minimum – weekly update meetings where the TIP/RTP provides an update on the progress made in the preceding week. The activities for the upcoming week are also discussed.
Go/no go TIR meeting	TCT, CT (including the C-level sponsor) and TIP or RTP	After the TIP delivers the TIR, a meeting is held where: ■ The TIR is formally approved; ■ The proposed (TI-based) scenarios are discussed and definitive test scenario(s) are selected.
Go/no go RTTP meeting	TCT, CT (including the C-level sponsor) and RTP	After the RTP has created the RTTP, a meeting is held to formally approve the RTTP and start the RT phase of the test.
PT session	CT, RTP Optional: TCT	During the PT session the RTP and the BT of the financial institution replay the attack. The BT and RTP subsequently focus on the areas with the most learning opportunities.
GT kick-off meeting (if GT module is selected)	TCT, CT, RTP and GTP	During the GT kick-off, the scope, learning goals and set-up of the GT phase are defined. This meeting also marks the formal start of the GT phase.
GT dry run (if GT TTX or SIM is selected)	TCT, CT and GTP Optional: RTP and TIP	If the tabletop exercise or the simulation variant are selected, a GT dry run is undertaken to approve that everything is in place for successful execution.
Board meeting	CT and BOD Optional: TCT, TIP, RTP and GTP	After the operational part of the test has been concluded, a board meeting can be held to communicate the results and the impact of the test to the BOD. If desired, the TCT will attend this meeting.
360 feedback session	TCT, CT, TIP, RTP, (GTP, if this module was selected)	During the 360 feedback session, all parties that were actively involved evaluate the test process and possible improvements in the ART framework. The operational results of the test are not being evaluated.

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 institution 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 potential damage can be restored. Furthermore, the institution 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 consider the following risks (non-exhaustive):

- 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 external providers, the institution makes sure that there is mutual agreement on the following aspects (non-exhaustive): 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 all 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
An ART test should mimic the (previously 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 an ethically acceptable point – to make the test as realistic as possible. This ethically acceptable point will be different for every institution and up to the CT to define. 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, without explicit approval of the holder of those attributes.

Code names
To prevent the leakage of sensitive information, code names must be used. 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 throughout all documentation related to the ART test, at least in document titles and throughout the documents. 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 communication, meeting invites and documentation between the parties involved in the test.

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 institution to set up contractual agreements with the RTP regarding, the inviolability of their employees' privacy, for instance. 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 institution, it cannot stop the test. However, it does have the authority to remove/deny the ART label, which means the test will not be recognised as an official ART test and no attestation document will be delivered. 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 the following situations (non-exhaustive):

- 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 institution, 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 institution 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 Preparing an ART test

The preparation phase

"Preparing an ART test" refers to the preparation phase of the ART test, which includes the engagement, scoping and procurement of external parties. This section covers these three aspects of the preparation phase.

4.1 Engagement

The main goal of the engagement steps is to define the institution's learning objectives and to get commitment from all parties involved. During this phase, the TCT and the CTL also determine which modules are to be included in the test. Another objective of this early phase is that the institution makes sure that all relevant internal stakeholders in the ART test are involved and aligned. The institution also ensures that the TCT is engaged. With the guidance of the TCT, the institution 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 the initiation meeting(s), which are the first meetings in the preparation phase of the ART test. During the engagement steps, 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 phases, the CTL and the TCT meet to discuss the scope of the ART test and to decide which modules to include. Although the decision as to which modules to include in the ART test ultimately lies with the institution itself, the road leading up to this decision is a collaboration between the CTL and the TCT. The final decision depends on a number of factors, including:

- The institution's characteristics, such as size, cybersecurity posture and internal TI and crisis management capabilities;
- How this test relates other forms of cyber security tests, either already completed or in the pipeline;
- The institution's budget;
- The institution's previous experience with threat-led penetration testing;
- The institution's learning objectives;
- The institution's ambitions.

Meetings

- Initiation meeting(s).

Milestones and main deliverables

- C-level commitment;
- Agreement on the selected modules and their variants for the specific ART test;
- Signed contract between the institution and DNB for guiding the ART test;
- Establishment of a CT;
- Agreement on code names and communication channels.

4.2 Scoping: critical functions and systems

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

Institutions 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 process. These flags will later serve as goals in the test scenarios, which are based on relevant threat intelligence.

During the scoping process, the institution must complete the ART SSD. In addition to defining the scope, the ART SSD lists the key systems and services that underpin each CIF. 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. If desired, the TCT can assist the CT by participating in workshops to create the SSD.

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, preferably face-to-face.

Milestones and main deliverables

- Delivery of the SSD;
- SSD approved by C-level sponsor and TCT.

4.3 Procurement

Based on the agreed scope, the institution's timetable and modules/variants selected, the CTL starts the procurement of an RTP and/or TIP and/or GTP. The ART service procurement guide can assist the CT with this task. If a variant of the TI module is chosen where the internal TI team will take on the role of TIP for the purpose of this test, the CT needs to make arrangements that these experts can fulfil that role without knowledge from the rest of the organisation.

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.

If desired, the request for proposal (RFP) can be shared with the TCT. The TCT can assist in verifying that the RFP contains all the necessary elements listed in the ART service procurement guide. Given the lengthy nature of procurement processes, this process should start as soon as possible, ideally even before the preparation phase. The TCT can arrange contact between the institution and other TIBER/ART institutions to request references for the provider. The TCT can also assist in conducting resume checks of individuals mentioned in the quotes supplied by the providers.

As soon as the (external) providers have been procured, the CT starts drafting an overall planning of the ART test. Also, a schedule of meetings is set up to be held between the institution, RTP and TCT and TIP (and GTP if applicable). Apart from the mandatory meetings, the TCT and CT should have regular meetings to discuss the progress made and the risks that are identified by the CT and the way they are managed. After procurement has been completed, the launch meeting can be held to align all stakeholders in the test. During that meeting, a number of practical agreements are made regarding the frequency of meetings during the TI and red teaming phases, communication channels, documentation and responsibilities. And although

the scenario is not yet clear, the institution starts working on a leg-up inventory and – where possible – prepares leg-ups.

Meetings

- Launch meeting.

Milestones and main deliverables

- Signed contract with RTP and/or TIP and/or GTP;
- Planning of the ART test;
- Setting up a risk register and risk management activities;
- Agreement on the communication channels and setting up the schedule of meetings during the ART test;
- Leg-up inventory and preparation.

5 Running an ART test

The test phase

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

5.1 Threat intelligence

The TI phase involves gathering information about the institution and identifying potential threats that are specific for this institution. This information is used to simulate real-world attack scenarios in the later stages of the test. During the engagement steps of the ART test, one of the variants of the TI module is selected by the institution. Despite the differences between the TI variants, the procedures and mandatory steps remain mostly the same.

To facilitate the TI phase and outcome, a business overview workshop can be organised by the CT for the TIP. It is of the utmost importance that the TIP understands not only the technical side, but also the vital business processes of the institution. The business overview workshop can therefore help the TIP to fully understand the information in the SSD and to make a better assessment of the threats that apply to the institution so that it can present one or more realistic threat scenarios.

During this part of the test, the TIP regularly informs the CT and the TCT of its observations and progress. At an appropriate time, a draft version of the TIR is provided for discussion and feedback. To ensure that the report continues to meet the requirements as set out in the framework/TI guide, the TCT can give feedback on how to align the report with those requirements. 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 TIR meeting, attended by the TCT, RTP, TIP and CT, as well as the C-level sponsor. During this meeting, the TIP presents its observations and the corresponding TIR. The CT/C-level sponsor and TCT approve the selected scenario(s) to initiate the preparation of an RTTP by the RTP, based on the TIR.

Approvals
The TIR should be approved by the CT (including the C-level sponsor) and the TCT in the Go/no go TIR meeting and confirmed in writing.

Meetings
- Business overview workshop (optional);
- Go/no go TIR meeting;
- Regular progress meetings.

Milestones and main deliverables
- Finalised and approved version of the TIR;
- Go for the creation of the RTTP, based on the selected scenario(s).

5.2 Red teaming

The RT phase consists of a number of steps. First, the selected scenario(s) from the TIR are converted into an RTTP. Once this plan is approved by the C-level sponsor, TCT and CT, the RT executes it by attacking the institution's live systems. The goal is to identify weaknesses, vulnerabilities and potential gaps in the institution's defences, providing insights it can use to improve its cybersecurity posture and incident response readiness.

The RTTP

In the RTTP, 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 scenario(s) drafted in the TI phase to ensure realism;
- 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 financial stability in the Netherlands;
- Also include some elements that test the institution's response, showing whether the attack would immediately be detected or could have a fair chance of succeeding.

The attack scenario(s) 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 attacks seen in the past and are therefore representative for the threat actor mimicked.

If a scenario X is added, that scenario could be 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 institution in the future. While the ultimate objective of a scenario X is to compromise a CIF, it attempts to do so using a highly creative approach.

Rules of engagement

The RTTP should include rules of engagement, in which the RTP lays down the rules it will follow. The rules of engagement should contain the following (non-exhaustive):

- 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 safeguarded in compliance with relevant legislation.

The red teaming

After the RTTP is approved, the RTP starts with the actual RT activities. During this phase, it performs an intelligence-led RT 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 RTTR 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(s) should be played out according to the chosen variant of the scenario(s). This could be as an assumed compromise or end-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 institution's CT and the TM. The TM is updated on progress at least once a week by the RTP and the CT. Face-to-face 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. Institutions have also had very positive experiences when a member of the CT was onsite with the RTP for some time during the test. 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.

Limited purple teaming
During active RT testing and before PT, there might be circumstances that require disclosing (a 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).

LPT is described in more detail in the ART RT guide.

Out phase plan
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. 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.

Completing the RT phase
The output of the red teaming phase is an RTTR produced by the RTP and delivered to the institution. The draft report must be shared within four weeks of the test's completion. It must give an overview of the entire ART process, including the CIFs, the scope, the threat intelligence base of the test, the planned scenarios, the executed scenarios, the test observations and the RTP's advice to the institution, and should be written using the RTTR format. At this point, key members of the institution's BT are informed of the test. If desired, they can write their own report ahead of the PT session. If the BT cannot write a full report based on their own observations or omissions in the monitoring, the RTTR can be shared with the BT.

Approval
The RTTP must be approved by the CT (including the C-level sponsor) and TCT in the Go/no go RTTP meeting, before the start of the actual test phase.

Meetings
- Go/no go RTTP meeting;
- Discussion of the 'out'-activities;
- Weekly updates.

Milestones and main deliverables
- Finalised version of the RTTP;
- Go for the execution of the RTTP;
- Achieving flags and learning goals during the RT phase;
- Filled in RTTR.

5.3 Purple teaming

After the RTP delivers its report, the institution organises a PT session. Often, the PT session is perceived as the most educational part of the test, leading participants to spend more time on this part of the process. The goal of this session is to enhance the learning experience. The PT session can last either one day (PT fundamentals variant) or two days (PT full variant), depending on the scope and duration of the test. Towards the end of the RT phase, the CT and TCT can discuss if the PT variant selected upfront is still the best fit for the learning goals. If needed, the contractual agreements with the RTP might need to be changed, if a change in the PT variant is preferred.

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

Approval
There is no formal approval step in the PT variants.

Meetings
- PT session.

Milestones and main deliverables
- Execution of the PT.

5.4 Gold teaming

GT (sometimes publicly known as a "tabletop" or "crisis management simulation") is a collaborative session with the crisis management team (CMT) of the financial institution. GT allows an institution to validate, train and exercise crisis management in a controlled environment. It is also a perfect opportunity to get (additional) C-level engagement during the test. During a GT exercise, the CMT gathers to discuss and respond to the crisis caused by the completed RT scenario.

Within the ART framework, the optional GT module allows the institution'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 section will give an overview of the different GT variants, the GT planning and the rules of engagements.

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

Walk-through session
This is the most low-key and accessible GT variant. It can be used for institutions with no or very limited experience in crisis management.
A walk-through session could also be a good fit for institutions that have seen significant changes in their crisis management structure and personnel. During the walk-through session, all steps in the crisis management process, from detection to closure, are discussed and completed in detail. A walk-through 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 whose purpose is to validate the crisis management processes and increase the CMT's knowledge of how to act if a specific crisis unfolds. The walk-through is facilitated by a facilitator.

Tabletop exercise
This GT variant is an accessible discussion-based exercise and a good fit for institutions 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 needs thorough preparation (including a dry run) and is always facilitated by a facilitator, trainer and/or observer for training and evaluation purposes.

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 institution'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:

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.

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 TM and depending on the learning goals of the institution. 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 GTTP.

During a GT, only fictional decision-making should take place in a controlled environment. GT is intended for the CMT members of the tested institution. Depending on the scenario and the institution's learning objectives and other response teams can be involved in the GT.

Every GT begins with a GT kick-off meeting and the creation of a GTTP. Ideally, this process should start well ahead of the GT exercise, preferably during the TI phase or the early stages of the ART test. 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 observations of the ART test. For the development of the scenario, it is important to involve the relevant departments of the tested institution.

The GTTP 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 institutions 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 greater if the "pain" from the RT phase is still present. For example, if the RT provider simulated a ransomware scenario during the red teaming phase, the GT 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 capture the flags in the scenario, the GT scenario will assume that it did. The impact that the RT had on the institution 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 RT. While in other instances it is advisable to delay the start of the GT by a couple of weeks to achieve optimal learning. The ideal starting point of a GT will be determined in consultation with the TM and CT.

Rules of engagement

The GTTP should include rules of engagement, in which the GTP lays down the rules it will follow. The rules of engagement should contain the following (non-exhaustive):

- 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

Before the actual GT exercise begins, both the TCT and the CT (including the C-level sponsor) must approve the GTTP. For the TTX and SIM also the MSEL needs to be approved. For all variants the TCT needs to approve the GT report.

Meetings

- GT kick-off meeting;
- GT dry run (depending on the variant).

Milestones and main deliverables

- Finalising and approving the GTTP;
- Approval for the execution of the GT exercise;
- Execution of the GT exercise;
- Delivery and approval of the GT report including points of improvement.

6 Learning from an ART test

Closure phase

The financial institution 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 institutions. DNB aims to facilitate a mutual exchange between institutions by encouraging information sharing through community building.

6.1 Test summary

The test summary report (TSR) summarises the ART process and should draw on the delivered documentation, such as the RTTR and optionally the BT observations, the TIR and the recommendations for remediating actions. The institution could use the ART test summary report format for this purpose. The gathered intelligence and lessons learned from the test can serve as input for the GTL used in future tests.

6.2 360 feedback

During the 360 feedback meeting, the CT, TM and providers 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 parties involved in the process and to improve future exercises. It can also provide feedback for future improvements in the ART framework.

6.3 Remediation plan

Based on the test outcomes, the institution 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 TIR, RTTR and GT evaluation can serve as input for the remediation plan. Further input could come from the CT and organisational observations. DNB encourages institutions to share their remediation plan with their supervisors. The TCT is not involved in the creation of the remediation plan.

6.4 Reporting to the BOD

It is of the utmost importance that the institution's BOD is informed of the threats, test results and the recommendations for remediation actions (risk mitigation measures) in a board meeting. If desired by the CTL, the TCT attends the presentation of the results and observations to the BOD. The TCT must stress the importance of BOD involvement, support and accountability in executing the following remediation actions.

6.5 Finalising the test and attestation document

After the test has been completed, the results have been shared and the PT 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 institution with an attestation document.

Annex 1 Adapting ART for use in other sectors

ART has been created by DNB for the Dutch financial sector and its critical ICT third party service providers. 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 are encouraged to align it as closely as possible with the circumstances and learning needs of the country or sector. However, there are certain conditions attached to the adoption and adaptation of ART.

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 secret. Only a very limited group (the CT) should be aware of the test;

- The test must be performed on live systems of the target institution;
- The test must always include at least 1 scenario involving red teaming;
- The scenario must always be based on threat intelligence;
- ART tests should be guided by an experienced TCT that operates independently from the tested institutions and the TIP/RTP/GTP. Independent guidance 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 formalised between the respective countries and sectors.

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