Adversary simulation: what you need to know

Adversary simulation ('red teaming') tests your ability to prevent, detect and respond to cyber attacks.

This guidance is for organisations wanting to understand if adversary simulation (sometimes known as 'red teaming') is right for them. It is designed for system owners and security professionals within medium to large-sized organisations.

Providers of adversary simulation services may also find this guidance useful, especially those considering joining the NCSC's assured Cyber Adversary Simulation (CyAS) Scheme, as the methodology within the scheme aligns to this guidance. The scheme assures commercial organisations providing adversary simulation services which meet the NCSC's rigorous technical standards.

In this guidance:

What is adversary simulation?

Adversary simulation is designed to test an organisation's defences against a realistic cyber attack by systematically testing an organisation's ability to prevent, detect and respond to a range of cyber attack scenarios.

Also known as 'red teaming', adversary simulation involves replicating the actions an adversary is likely to use in a real attack, and aims to safely achieve one or more specified outcomes that would have a significant impact on business functions if a malicious attacker were to achieve them. It can be conducted by internal teams carrying out their own due diligence, or by external providers specialising in this type of service.

Adversary simulation tests whether your organisation's technical defences are aligned and operate as expected, and reveals where you may have gaps in your security posture. It also evaluates whether your team can identify threats early, triage them quickly and appropriately, and escalate where necessary.

Adversary simulation differs from penetration testing as it focuses on the efficacy of an organisation's technical controls and detection, whereas penetration testing is more focused on identifying all technical vulnerabilities within an organisation's IT systems.

> **Note:**
> The NCSC would expect all responsible adversary simulation teams to ensure that their resources, tools and capabilities have safety features/guardrails to minimise risks to the customer systems and other users.

What organisations benefit from adversary simulation?

Adversary simulation is best suited to organisations with a mature understanding of the cyber risks they face. To get the most from adversary simulation, these organisations will:

- have identified, assessed, and be regularly reviewing their risks

- have well-established mitigations and defences in place
- have robust network monitoring and detection systems

Adversary simulation should verify that your network monitoring and detection teams (whether provided by an internal team or third-party suppliers) can detect unusual activity and respond appropriately.

While adversary simulation is appropriate for any organisation meeting the above conditions, it is particularly useful for larger organisations, or for those operating within critical national infrastructure or UK government.

> **Note:**
> Organisations which fall outside these parameters (perhaps operating relatively small and simple networks, or still developing an understanding of risk) may find more value from the services provided by companies assured under other NCSC-assurance schemes.

How long does an adversary simulation engagement take?

A typical adversary simulation engagement will normally take between 8 and 12 weeks to complete, depending on its size and scope, although a full spectrum engagement may last in the region of 16 weeks. The engagement should be proportionate to the size, nature, and complexity of the objectives.

Rather than setting the assessment period at the bare minimum, we encourage organisations to allow sufficient time for thorough testing of real-world adversary tactics, techniques, and procedures.

Adversary simulation best practice

The NCSC's approach to adversary simulation differs from most industry frameworks in that it is **capability led**, where the testing approach is not constrained by attempting to simulate specific threat actors or pre-defined threat scenarios.

This means that the team undertaking the adversary simulation is responsible for carrying out reconnaissance activity on the customer's organisation to inform their attack plans, rather than using commercially-procured threat intelligence and threat scenarios to drive the approach.

Through reconnaissance activities, the team develops an understanding of the customer organisation's technology, exposure, and threat landscape. This ensures that the testing approach reflects the specific risks relevant to the customer and realistically models the threats they are likely to face.

The NCSC encourages the application of an adversarial mindset when seeking to achieve testing objectives. Teams should maintain and use a suite of tools and capabilities to achieve a range of objectives – as agreed with the customer – rather than only seeking to mimic specific indicators of compromise or tactics, techniques and procedures (TTPs).

Known and reported TTPs are inherently retrospective and may not fully reflect current adversary tradecraft. Consequently, testing should not be limited to documented techniques when a competent adversary simulation team can emulate more contemporary and representative attack behaviours.

> **Note:**
> The NCSC's Cyber Adversary Simulation (CyAS) Scheme assures commercial organisations providing adversary simulation services which meet the NCSC's rigorous technical standards. As the national technical authority for cyber security, our approach to adversary simulation (which is reflected in the CyAS scheme) is outlined in the sections below.

'Full spectrum' vs 'assumed breach' approaches

The NCSC's approach supports two different approaches to evaluating an organisation's ability to detect and respond to a cyber attack, depending upon the attacker's starting location:

- a **full spectrum** approach starts from outside of the network, and evaluates an organisation during an end-to-end attack which attempts to breach the perimeter and pursue the agreed objectives
- an **assumed breach** approach starts from a point within the customer network, and simulates a threat once an attacker has managed to gain an initial foothold

The most suitable approach depends on what the customer wishes to achieve. Both approaches will help an organisation understand more about their resilience to cyber attacks, but the approaches differ in terms of what the customer can learn about their security. For example, an assumed breach approach will not tell the organisation what information they may deliberately or inadvertently be making publicly available (which could be of interest to an attacker).

For organisations with a mature security posture, assumed breach can provide greater value by bypassing the initial access phase and concentrating on the consequences of a successful compromise, specifically whether an attacker can expand their access beyond the initial point of entry and reach high-value targets.

Regardless of the approach taken, the NCSC's preferred methodology is split into 3 phases: **prerequisites**, **testing**, and **reporting**.

Phase 1: Prerequisites

The prerequisites phase typically involves **scoping**, **initial passive reconnaissance** and **preparation**.

Scoping

At the scoping meeting, the team providing the adversary simulation and the customer discuss and formally record the objectives of the engagement, commercial arrangements and timescales. Objectives should focus on identifying the key functions which, if compromised, would critically impact the ability of the organisation to operate effectively.

The adversary simulation team may help shape objectives during scoping, but the onus is on the **customer** to make clear what the organisation's essential functions are, and what the overall objectives of the adversary simulation should be.

The scoping meeting also allows the customer to share information critical to the safe performance of the adversary simulation. The customer should provide only the minimum information necessary to begin the reconnaissance stage of the exercise. For example, the primary domain name of the organisation, or the system name. This is unlike penetration testing, where detailed information about the system or network is shared with the testing team to enable them to prepare for the test.

Initial passive reconnaissance

For a full spectrum engagement, the adversary simulation team will build up an understanding of the customer organisation by undertaking reconnaissance. This understanding will inform the preparation and initial access phases - for example, tailoring the pretext to be used as part of a social engineering campaign based on information gained about the customer. The NCSC expects this work to be completed by the adversary simulation team, and should **not** solely be the result of a procured threat intelligence product.

Passive open source intelligence gathering is done with stealth, that is without interacting with the target organisation. This can include, but is not limited to:

- using online databases
- search engines
- looking up IP/DNS registrations
- certificate information
- relevant social media

Assumed breach engagements may also require limited passive reconnaissance to be undertaken in preparation for the testing phase. For example, reviewing the customer's website to ensure that any domains the adversary simulation team plans to register match the customer organisation.

Preparation

With information from the scoping exercise and the passive reconnaissance, the adversary simulation team are able to develop tools and techniques to use later in the active testing phase. This may include registering of domains to align to the attack plan or social engineering pretext, though having some generic domains 'ready to go' may be of value.

During this phase, the customer must create some legitimate low-privilege level user accounts (as agreed in the scoping section) that could be used as a contingency if the 'active' testing from outside of the perimeter is unsuccessful. The accounts need to be as representative of a real account as is possible (group memberships, user metadata) and should not lead to the engagement being compromised before it has even started.

Provisioning users that do not fit with the corporate naming convention may lead to increased suspicion and trigger an internal review by the customer's defensive teams (often called 'blue teams'). Similarly, devices provisioned specifically for adversary simulation often garner an increased level of focus due to them suddenly appearing on the asset register.

Phase 2: Testing

The testing phase typically involves **continual active reconnaissance**, **initial access**, **internal phase** and **clean up**.

Continual active reconnaissance

The adversary simulation team are likely to repeat this process throughout the engagement. As mentioned previously, this work is undertaken by the team providing the service, and should not be the result of a procured threat intelligence product.

Now that the adversary simulation has commenced, the reconnaissance is likely to be more active than passive. Active reconnaissance involves direct interaction and runs the risk of being detected. It can include, but is not limited to:

- limited port scanning
- visiting the customer organisation's websites
- surveying available online services

Initial access

For full spectrum engagements, the adversary simulation team uses reconnaissance findings to identify potential access points to the customer's internal IT systems. Common techniques include phishing or watering hole attacks, which may have been identified in the scoping phase as potential attack vectors. The team may identify additional opportunities during active testing and discuss them with the customer.

For assumed breach engagements, the team will start from a point within the customer network. There are numerous ways to achieve this, for example customer staff may be included as part of the engagement to facilitate the adversary simulation team's initial access to the network.

For both types of engagements, it is important that the team establishes and discusses the escalation process during the scoping phase, so that, if the customer's defensive security team detects the attack, there are some pre-agreed options covering the course of action the customer will take.

Internal phase

Once the adversary simulation team has established a foothold on the customer network, it works towards the objectives agreed during the scoping phase. This is a crucial phase where the customer's internal capability to **detect** and **identify** anomalous behaviour and **protect** its most critical functions is tested. If the team achieves one or more objectives without detection, they will advise the customer and then agree how to proceed. Note that it is for the customer to determine whether an objective has been met (in conjunction with the evidence provided by the adversary simulation team).

If the objectives have been met, the customer and adversary simulation team may choose to agree additional activities in the remaining timescales. It is important that the necessary permissions are in place for such activity, and that this is

documented. Such activities may include identifying alternative approaches to achieve the same result, or gradually increasing the ‘noise level’ until the detection team becomes aware of the attack.

If the adversary simulation team cannot gain a foothold from an external vantage, the testing may proceed using a de-chain action whereby the customer provides a (typically low-privileged/standard) user account or a device with credentials, from which the team can continue to work towards the agreed objectives.

Using a de-chain action should be a serious consideration if the engagement stalls and the provider is not making progress. This is still valuable as it provides insight on the risks faced from trusted third-parties, or from malicious insiders seeking to access the customer's most sensitive data. However, the customer and the adversary simulation team need to carefully judge the best time for switching approaches, balancing the benefit of faster progress with the consequences of losing insight into the resilience of the systems to attack from outside.

A de-chain action may also be used in other scenarios, such as where an attacker has a capability not available to the adversary simulation team at present (for example zero-day vulnerabilities), or in a scenario where the adversary simulation team have identified an attack path which the customer does not wish them to exploit.

The use of a de-chaining action reflects a mature approach to testing. Rather than being constrained by a single attack path, it enables testing to remain focused on the organisation's most significant risks and the outcomes that matter most. By redirecting effort where it will provide the greatest insight, a de-chaining action helps maximise the value of the engagement and strengthens confidence in the organisation's overall resilience.

Clean up

In this phase, the adversary simulation team remove from the customer’s system all artefacts that were created during the engagement. If they are unable to remotely remove artefacts, the team must document them and provide to the customer details of how to safely remove them. Note that the adversary simulation team may advise that rebuilding a host is the only viable solution.

In addition, as secrecy is no longer necessary, the customer’s internal detection team may wish to more thoroughly review all potential indicators of compromise (IoCs) that might have been flagged but not acted upon during the test window. This will allow the customer to immediately start to learn from the engagement while waiting for the report to be produced.

Throughout the engagement, the adversary simulation team must keep a contemporaneous record of all activity. This provides evidence in the event of any disputes resulting from the test, such as proof that *Technique A* was used on *Host B*. A sanitised version of the record can be provided to the defensive security team to aid in detection of the IOCs and to provide opportunities for staff training.

Phase 3: Reporting

The final phase of the adversary simulation includes report writing and review meetings. The team produces a detailed report covering all stages of the engagement. This will include all of the information gathered in the reconnaissance phase so that the customer can fully understand their internet-facing footprint and can take action to address any concerns.

The report explains how the team achieved objectives, identifies detection opportunities, and recommends improvements.

Importantly, if any part of the adversary simulation relied on social engineering, the adversary simulation team should anonymise the targeted users’ details to avoid harm. Instead, the report should focus on the privileges that were obtained, and how this supported the rest of the attack.

Unlike a penetration testing report, an adversary simulation report does **not** document every potential vulnerability, but instead:

- emphasises how different vulnerabilities or misconfigurations were combined to achieve the attacker's objectives
- recommends improvements to the customer’s network and detection capabilities.

> **Note:**
> All reports produced by suppliers operating as NCSC CyAS Assured Service Providers must include a mapping to the Cyber Assessment Framework.

Alternative NCSC assurance schemes and resources

If your organisation has not reached a level of cyber maturity from which it can detect and respond to a cyber attack (or has more specific requirements), the NCSC offers other resources and assurance schemes:

- **Cyber Essentials:** ensures organisations have fundamental cyber security technical controls in place.
- **Cyber Advisors:** cost effective cyber security advice and practical support aimed at small organisations.
- **Assured Cyber Security Consultancy:** consultancy in Risk Management, Security Architecture, Audit and Review, and Post-Quantum Cryptography for organisations with high-risk or complex cyber security requirements.
- **Exercise in a Box:** a free, practical, self-service toolkit providing ready-made scenarios and supporting materials to help organisations test incident response plans.
- **Cyber Incident Exercising:** a facilitated service led by NCSC-assured providers. Exercises are designed, delivered and tailored to your organisation's specific risks, systems, and leadership needs.
- **CHECK penetration testing:** penetration testing for government departments, public sector bodies and critical national infrastructure using the NCSC's methodology.

PUBLISHED

17 September 2026

WRITTEN FOR

Cyber security professionals

Large organisations

Public sector