CHECKDONE IT

Grip on IT Security




# Deception Technology

## From attacker misdirection to measurable cyber resilience

### Whitepaper

| | |
|---|---|
| Doc.nr. WP20260921-01 | Version: 1.0 (Final) |
| Author: Hans van den Biggelaar | © Checkdone IT B.V. |

Nieuwe Eerdsebaan 8a, 5482 VS Schijndel • +31 73 549 5958 • www.checkdone-it.eu • info@checkdone-it.eu

CHECKDONE IT

# Deception Technology

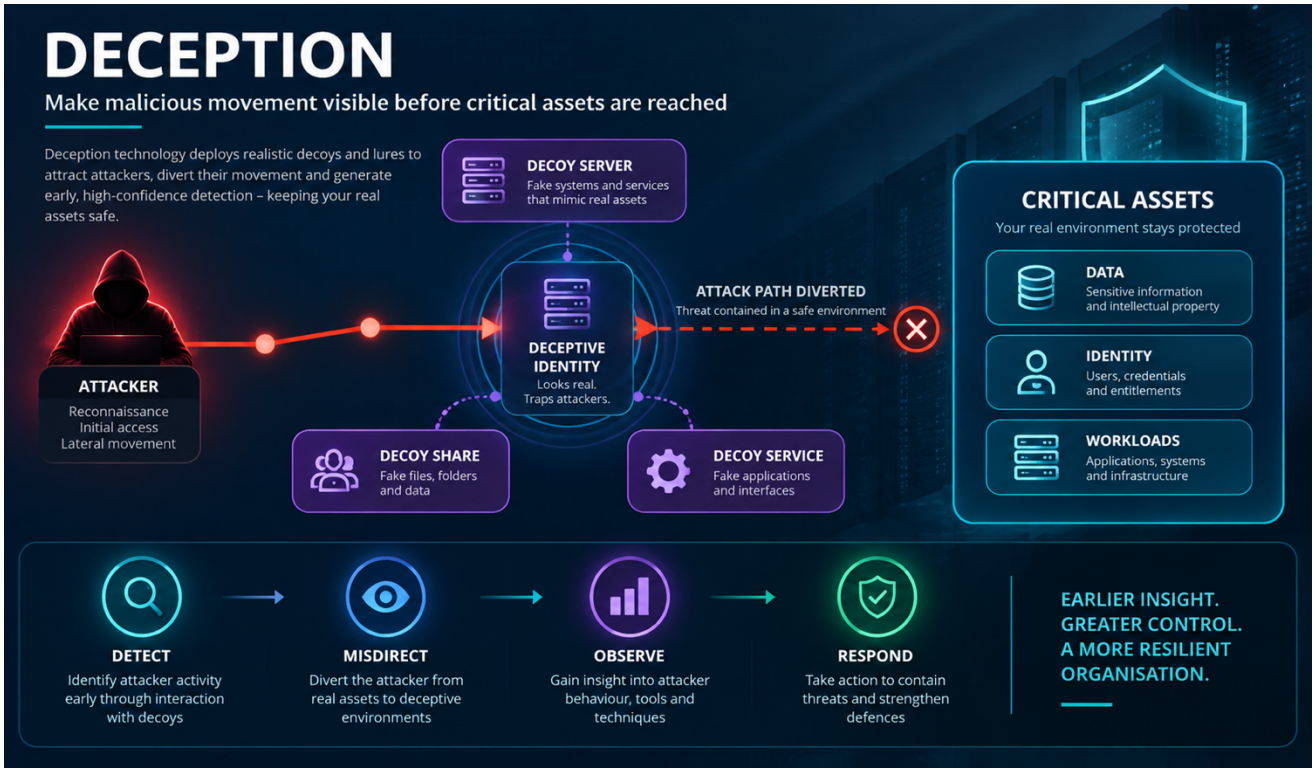
## From attacker misdirection to measurable cyber resilience

A practical, vendor-neutral framework for governance, security and IT, with the relationship to Zero Trust, Continuous Threat Exposure Management (CTEM), compliance frameworks and EU regulation.

### Why this white paper?

Deception technology changes a simple assumption in defensive security: not every system, identity, credential or data object visible to an attacker has to be real. By introducing controlled and believable decoys, lures and false paths, an organisation can make unauthorised exploration easier to detect, harder to trust and more informative for defenders. The purpose is not to "trap" ordinary users, but to create security signals around actions that should have no legitimate business reason to occur.

This white paper explains what modern deception platforms do, where they add value and where they do not. It places deception within a broader security model rather than treating it as a stand-alone product category. Particular attention is given to Zero Trust, CTEM, SOC and incident response, and to the governance, policy and operational prerequisites required for safe and effective use.



# Contents

**Introduction: detection after prevention** ..... **4**

**Deception: not a trap, but a defensive capability** ..... **4**

**How modern deception platforms work** ..... **5**

Believability without production dependency ..... 5

Distributed lures and identity artefacts ..... 5

Interaction telemetry and investigation ..... 6

Response and validation ..... 6

**Core capabilities and use cases** ..... **7**

**The role of deception in the wider security landscape** ..... **8**

From observing the environment to shaping it ..... 8

Prevention, detection, response and validation ..... 8

A connective capability rather than another silo ..... 9

What deception does not replace ..... 10

**Relationship with Zero Trust** ..... **10**

Assume breach becomes observable ..... 11

Least privilege and blast-radius reduction ..... 11

Verify explicitly with richer context ..... 11

Continuous trust, not permanent trust ..... 12

Deception as a Zero Trust validation mechanism ..... 12

**Relationship with CTEM** ..... **13**

From vulnerability to exploitable path ..... 13

How deception contributes across the CTEM cycle ..... 14

Scoping: start with business impact ..... 14

Discovery: expose hidden relationships ..... 14

Prioritisation: add behavioural evidence ..... 15

Validation: move from assumption to evidence ..... 15

Mobilisation: turn evidence into reduction ..... 16

From exposure management to measurable resilience ..... 16

**Bringing the three together** ..... **17**

**SOC, incident response and threat intelligence** ..... **17**

From deception signal to operational value ..... 17

Integration in the SOC workflow ..... 18

Use cases in the SOC and incident response ..... 19

From operational signals to threat intelligence ..... 19

**Governance and policy: the foundation ..... 20**

Policy integration ..... 20

Rules of engagement ..... 21

**Compliance and EU laws and regulations ..... 21**

NIS2 and national implementation in the Netherlands and Belgium ..... 21

GDPR / AVG and national data protection legislation ..... 22

DORA and other EU requirements ..... 22

**Mapping deception to common frameworks ..... 23**

**Implementation prerequisites ..... 24**

**Implementation roadmap ..... 25**

**Measuring value and maturity ..... 27**

A practical maturity model ..... 28

**Practical considerations and limitations ..... 28**

**Conclusion: from hidden movement to actionable evidence ..... 29**

Role of Checkdone IT ..... 29

**Glossary ..... 30**

**Sources, reference frameworks and disclaimer ..... 31**

**Control over cyber resilience without any vendor strings attached ..... 32**

# Introduction: detection after prevention

Most information-security programmes invest heavily in preventing compromise: strong authentication, endpoint protection, vulnerability management, network controls, secure configuration and user awareness. These measures remain essential, but none can guarantee that an attacker will never obtain a foothold. Stolen credentials, vulnerable systems, compromised suppliers and legitimate tools can all provide ways into an environment.

That is why cyber resilience also depends on what happens after the first control is bypassed. Can the organisation recognise reconnaissance, credential abuse and lateral movement before critical data or services are affected? Can defenders distinguish a meaningful attacker action from the enormous volume of normal activity? And can they respond quickly enough to contain the incident?

Deception technology addresses this part of the problem. It deliberately creates resources that look useful to an attacker but are not required by normal business processes. Examples include decoy servers and services, false credentials, synthetic accounts, fake file shares, documents, database artefacts, cloud resources and breadcrumbs that appear to reveal a path towards something valuable. Interaction with these assets is monitored closely. Because legitimate users and applications should not need them, an interaction can provide a high-quality signal of suspicious activity.

This idea is related to the traditional honeypot, which NIST describes as a system or resource designed to be attractive to intruders. Modern deception extends the concept into a distributed security capability: many types of believable artefacts can be placed across identity, endpoint, network, application, cloud and data layers, with central management, telemetry, integrations and response workflows.

**The central idea**
Deception is most valuable when it is connected to risk, attack paths, monitoring and response. The technology can create unusually high-confidence signals, but those signals only become cyber resilience when ownership, processes and action are in place.

# Deception: not a trap, but a defensive capability

The word deception can sound offensive or adversarial. In an enterprise security context, however, the objective is defensive: shape what an intruder can discover, create uncertainty about what is real, detect unauthorised movement and collect evidence without exposing production assets. NIST's cyber-resilience guidance describes misdirection through honeypots, honeynets and decoy files, and disinformation through false credentials and tokens. The underlying purpose is to direct adversary activity towards controlled resources and learn from it.

A mature implementation therefore has two simultaneous goals. The first is detection: make attacker behaviour visible at an early stage and with a low level of ambiguity. The second is influence: make the attacker spend time validating false information, following dead ends or revealing techniques before reaching real assets. The defender gains time and context; the attacker loses certainty.

*Tabel 1: Deception concepts*

| Concept | What it means in practice |
|---|---|
| Decoy | A controlled system, service, application, share, database, cloud resource or other asset that resembles a real target. |
| Lure / breadcrumb | Information deliberately made discoverable so that an attacker is encouraged towards a decoy, for example a shortcut, configuration reference or synthetic secret. |
| Honeytoken / deceptive credential | A token, credential, API key, account or other identity artefact with no legitimate operational use. An attempted use is security relevant. |
| High-interaction deception | A decoy that supports enough realistic interaction to observe attacker behaviour and techniques in more detail. |
| Deception telemetry | Events and context generated when deceptive assets are discovered, accessed, used or attacked. |
| Attack-path validation | Using controlled simulations and deception signals to test whether expected attack routes and defensive controls behave as intended. |

# How modern deception platforms work

A deception platform normally combines orchestration, realistic decoys, distributed lures, telemetry and integration. The exact architecture differs, but the operating model is consistent: place deceptive resources where an attacker is likely to look, make them plausible in the local environment, observe any interaction and connect the resulting signal to investigation and response.

## Believability without production dependency

A decoy must be convincing enough to attract malicious exploration but isolated enough that compromise of the decoy cannot become a new route into production. Names, services, operating-system characteristics, directory information and apparent relationships should fit the surrounding environment. At the same time, decoys should not become dependencies for legitimate applications or business processes.

## Distributed lures and identity artefacts

The strongest coverage often comes from placing lightweight artefacts near real users and assets. These can include false credentials, synthetic identities, RDP or SSH references, browser artefacts, mapped resources, fake documents, database references, API tokens and other breadcrumbs. The objective is to

create detection points along likely attacker workflows: discovery, credential access, privilege escalation and lateral movement.

## Interaction telemetry and investigation

When a deceptive resource is touched, the platform can record the source, destination, protocol, command sequence, credential used, process context and other available evidence. Higher-interaction decoys can capture richer behaviour. This telemetry should be forwarded to existing SOC tooling so that it can be correlated with identity, endpoint, network and vulnerability context.

## Response and validation

Integrations can trigger playbooks such as isolating a host, disabling or restricting an account, revoking a session, blocking a connection or opening a high-priority case. Automation should be proportional to confidence and business impact. The same environment can also support controlled attack-vector and control validation, helping teams verify that detections, policies and SOC procedures operate as intended.

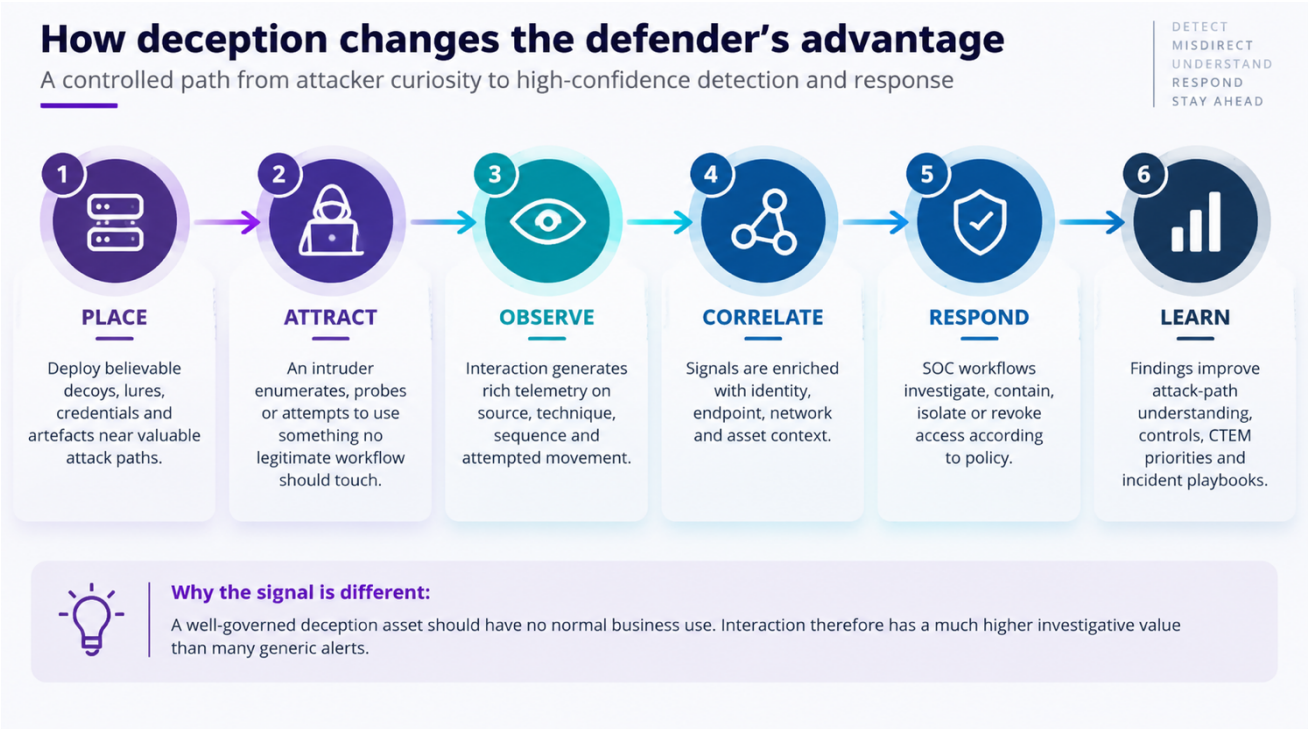



*Figure 1: Deception creates a controlled path from interaction to learning and improvement.*

# Core capabilities and use cases

*Tabel 2: Core capabilities and use cases*

| Capability | Typical value |
|---|---|
| Network and service decoys | Detect scanning, service enumeration, exploitation attempts and lateral movement towards systems that should not receive legitimate traffic. |
| Identity deception | Expose credential theft, password spraying, privilege discovery and attempts to use synthetic or planted credentials. |
| Endpoint lures | Place local artefacts that reveal hands-on-keyboard activity, discovery and credential hunting on compromised endpoints. |
| Data deception | Use fake documents, shares, records or database objects to detect unauthorised browsing and data-access behaviour. |
| Cloud deception | Create decoy resources, credentials, keys or service artefacts in cloud and hybrid environments to detect misuse and movement. |
| OT/IoT deception | Where technically and operationally appropriate, emulate devices or services to gain visibility without placing intrusive agents on sensitive production assets. |
| Threat investigation | Capture attacker techniques and sequences to support triage, scoping and lessons learned. |
| Security-control validation | Use controlled interactions or simulations to confirm that monitoring, segmentation and response controls behave as expected. |

**High confidence does not mean infallible**
Legitimate activity can still touch a deception asset through scanning, backup, vulnerability assessment, asset discovery or misconfiguration. These expected interactions should be systematically documented, filtered or excluded through design, rather than being assumed away.

# The role of deception in the wider security landscape

Modern information security is built on layers. Identity controls determine who should have access, endpoint and network controls protect systems and communications, vulnerability and exposure management identify weaknesses, segmentation restricts movement, and monitoring technologies attempt to recognise suspicious behaviour. However, it should be noted that none of these capabilities eliminate the possibility that an attacker will obtain a valid credential, compromise a trusted endpoint or exploit an overlooked weakness.

Deception occupies a distinctive position within this layered model. Its primary function is not to prevent the initial compromise. Instead, it modifies the environment that an attacker encounters after gaining a foothold. By introducing believable but controlled systems, identities, credentials, services and data artefacts, defenders can create points at which unauthorised exploration can be revealed.

This aligns with the broader concept of cyber resilience. Prevention remains a key priority, but resilience also depends on the organisation's ability to continue defending when a control has failed or been bypassed. Deception is a key factor in this process, as it creates controlled false paths, attracts malicious exploration and converts attacker interaction into evidence that defenders can use. Deception can therefore be viewed not merely as another detection technology, but as a mechanism for continuing to defend an environment when preventive controls have already been bypassed.

## From observing the environment to shaping it

Most monitoring technologies observe activity taking place in the real environment. EDR observes endpoints, NDR analyses network behaviour, identity security monitors authentication and identity activity, while SIEM and XDR correlate information from multiple sources. These capabilities must continually distinguish malicious behaviour from large volumes of legitimate activity.

Deception changes this dynamic. Rather than observing only genuine assets and attempting to determine whether activity around them is malicious, it deliberately introduces believable resources for which legitimate interaction should be absent or strictly controlled. A connection to a production server may have thousands of valid explanations, making malicious activity difficult to distinguish from normal behaviour. An attempt to use a synthetic privileged credential that has never been assigned to a legitimate user, however, has very few legitimate explanations. This makes interaction with deceptive assets a particularly strong security signal, helping defenders identify potentially malicious activity earlier and with greater confidence.

### The security signal is different

The value is not simply another alert. Deception creates security-relevant moments with deliberately little or no legitimate activity around them. When correctly designed, those moments provide the SOC with a high-confidence starting point for correlation with endpoint, identity, network, cloud and vulnerability information.

## Prevention, detection, response and validation

Deception crosses the traditional boundaries between security disciplines. From a detection perspective, it can reveal reconnaissance, credential harvesting, privilege discovery, lateral movement and unauthorised access. From a response perspective, interaction with a deceptive asset can provide the

confidence and context required to initiate investigation, revoke a session, isolate an endpoint or restrict an identity.

Its role can extend further into validation. If an organisation believes that segmentation prevents movement between two security zones, that privileged credentials are adequately protected or that a particular attack path is monitored, deception can help determine whether those assumptions hold in practice. Controlled exercises and observed interaction provide evidence about whether security architecture, monitoring and response processes behave as intended.

**From controls to assurance**

A mature programme moves from asking "Have we implemented the required controls?" towards asking "Do our controls actually prevent, detect and contain the attacker behaviours that matter to us?"

## A connective capability rather than another silo

The greatest value is achieved when deception connects existing capabilities rather than becoming another isolated security console. Identity context can make a deceptive credential more believable and help interpret its use. EDR can identify the process that accessed a lure. NDR can show subsequent movement. Exposure-management information can identify where deception should be placed. Threat intelligence can help determine which attacker techniques should be represented.

The relationship is bidirectional. Deception telemetry can enrich existing security capabilities with evidence of actual adversary behaviour, while those capabilities provide the context that turns a deceptive interaction into a meaningful investigation. This creates a practical security loop: understand risk, reduce exposure, restrict access, detect movement, investigate behaviour, contain the threat, validate controls and improve the environment.

Deception is not responsible for every step. Its role is to make attacker activity visible at strategically chosen points and turn that activity into evidence that other security processes can use.

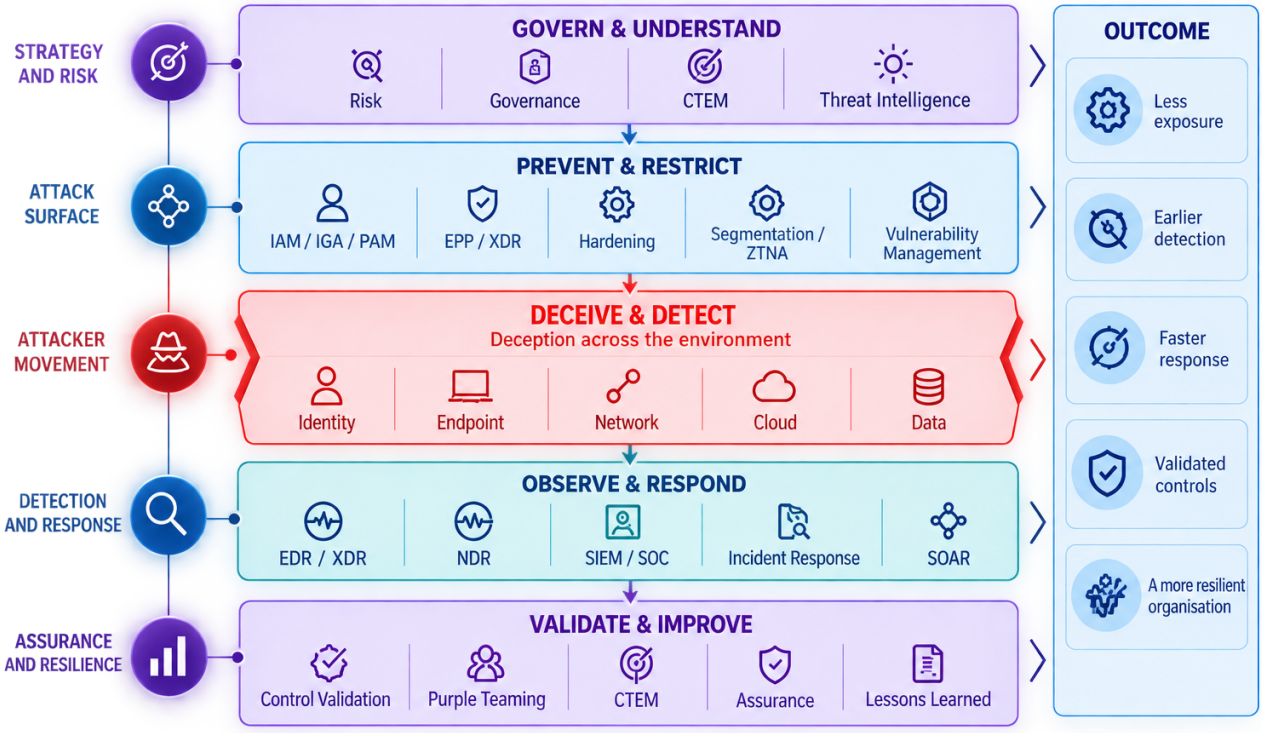



*Figure 2:Where deception fits in a layered security model.*

### What deception does not replace

**This distinction is important.**

Deception does not replace secure configuration, vulnerability management, IAM, IGA, PAM, EDR, NDR, segmentation, SIEM, threat intelligence or incident response. Nor does it compensate for weak governance. Instead, it provides an additional defensive layer for the situation in which one or more of those controls have failed, been bypassed or are operating on incomplete information.

That is why deception is best positioned as a connective cyber-resilience capability: one that links architecture, exposure, detection, response and assurance around the behaviour of an attacker.

## Relationship with Zero Trust

Zero Trust and deception address different parts of the security problem. Zero Trust governs how access to resources should be granted and maintained; deception helps reveal what happens when an adversary nevertheless obtains access or begins operating inside the environment. Zero Trust removes implicit trust and replaces it with explicit, minimal, context-dependent and revocable decisions. Deception strengthens that model by creating additional evidence about whether an apparently legitimate identity, device or session is behaving in a way that deserves continued trust.

The relationship becomes clearest when deception is mapped against the three practical Zero Trust principles: verify explicitly, use least privilege and assume breach.

## Assume breach becomes observable

"Assume breach" does not mean treating every user or system as compromised. It is an architectural principle: preventive measures can fail, credentials can be stolen and trusted systems can become compromised. Security architecture must therefore limit the consequences and detect malicious activity even after an attacker has established an initial foothold.

Deception gives this principle an observable dimension. Consider an attacker who has compromised a legitimate endpoint and acquired valid credentials. Authentication may succeed. The endpoint may still appear compliant. Network traffic may initially resemble legitimate administration. Traditional controls therefore have to infer malicious intent from context and behavioural anomalies.

Deception involves actions that a legitimate user should have little or no reason to perform, such as accessing a false administrative share, attempting to authenticate with a planted credential, enumerating a synthetic privileged account or following a breadcrumb towards a decoy management system. Therefore, the attacker may disclose the compromise through their chosen method of exploration, even when the initial authentication appears legitimate.

## Least privilege and blast-radius reduction

The Principle of Least Privilege limits what an identity, workload or device can access and preferably limits that access to the period for which it is required. Segmentation and Zero Trust access controls then restrict the paths through which resources can be reached. Deception does not enforce least privilege, but it can help determine whether least privilege and segmentation are delivering the expected outcome.

Deceptive assets can be positioned around privileged administration paths, identity infrastructure, management interfaces, sensitive data locations and trust boundaries. Attempts to discover or reach these resources can indicate that a compromised identity or endpoint is moving beyond its expected operational scope. This is particularly useful for identifying the difference between authorised access and authorised behaviour. A compromised administrator account may technically possess extensive permissions. That does not mean that every action performed with those permissions is legitimate.

Deception adds behavioural evidence to the access model and can reveal when valid privileges are being used to explore paths that should not form part of the user's normal workflow.

## Verify explicitly with richer context

Explicit verification is stronger when decisions are based on more than username, password, MFA status or device compliance. Identity risk, device posture, location, workload characteristics, requested resource and previous activity can all contribute to an access decision. Deception adds another category of context: interaction with something that should not have been used at all.

A device that has attempted to use a honeytoken, an identity that has accessed a decoy resource or a session associated with systematic discovery of deceptive assets represents a materially different risk condition. Where architecture and integrations permit, that evidence can contribute to actions such as step-up authentication, session revocation, access restriction or endpoint isolation. This does not mean every interaction should trigger automated containment. The response should remain proportional to confidence, business impact and the organisation's predefined policy.

## Continuous trust, not permanent trust

This exposes a deeper relationship between the two concepts. Zero Trust is strongest when trust is understood as temporary and revocable rather than as a status obtained once at login. Deception provides additional evidence that can challenge an existing trust decision after authentication has already succeeded. In that sense, deception helps make continuous trust evaluation tangible: access can begin as legitimate and later become suspect because behaviour around deceptive resources indicates compromise.

CISA's Zero Trust Maturity Model similarly treats Visibility and Analytics, Automation and Orchestration, and Governance as cross-cutting capabilities supporting Zero Trust. Visibility informs policy and response, while automation can use contextual information from multiple sources to support security decisions.

Deception fits naturally into this model: it adds high-value visibility, which can inform automated or human response, within a governed policy framework.

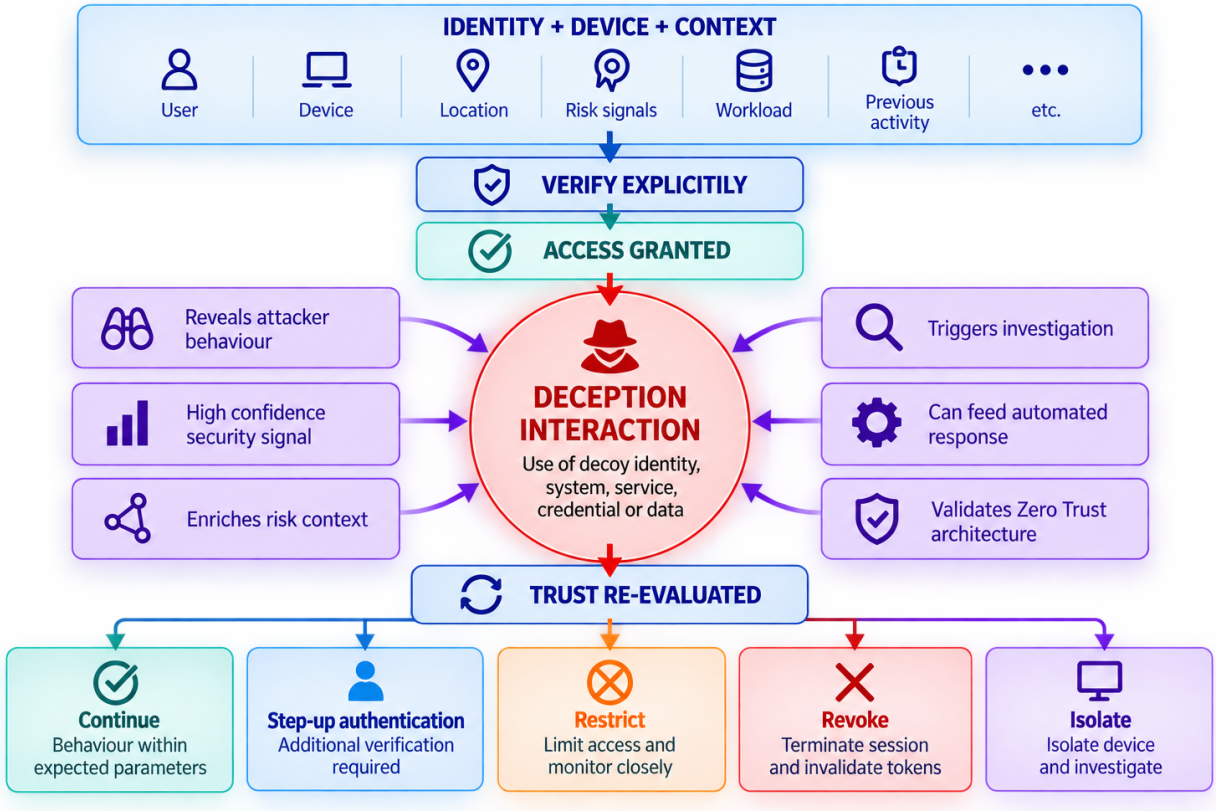



*Figure 3: Deception makes dynamic trust observable (Zero Trust).*

## Deception as a Zero Trust validation mechanism

There is a further relationship that deserves emphasis: deception can help determine whether the Zero Trust architecture works as designed. If a deceptive asset represents an attractive lateral path towards a critical resource, controlled security testing can answer practical questions about discoverability, reachability, segmentation, telemetry, SOC escalation and the speed with which access can be revoked.

- Can an endpoint in this zone discover the deceptive resource?
- Can the identity authenticate or otherwise interact with it?

- Does segmentation block the attempted path where it should?
- Does the attempt generate the expected identity, endpoint and network telemetry?
- Does the SOC recognise, enrich and escalate the event correctly?
- Can access be restricted or revoked quickly enough without unacceptable business impact?

The result is more than detection. It is evidence about the effectiveness of the Zero Trust architecture itself.

# Relationship with CTEM

Continuous Threat Exposure Management (CTEM) represents a shift from managing individual vulnerabilities towards continuously understanding which exposures can actually contribute to meaningful attack paths and business risk. It connects scoping, discovery, prioritisation, validation and mobilisation into a continuous programme rather than treating vulnerabilities as an undifferentiated list.

Deception contributes by adding attacker-facing evidence to that cycle. It does not replace CTEM and it is not another CTEM stage. Its value lies in helping organisations test assumptions, observe real or simulated attacker behaviour and feed the resulting evidence back into exposure priorities and remediation decisions.

Continuous Threat Exposure Management represents a shift from managing individual vulnerabilities towards continuously understanding which exposures can actually contribute to meaningful attack paths and business risk.

Gartner's current exposure-management material continues to emphasise moving beyond vulnerability-centric approaches towards broader exposure context, prioritisation, adversarial validation and remediation. Its 2026 CTEM research also emphasises connecting exposure information with SecOps so that exposure intelligence leads to coordinated and measurable risk reduction.

## From vulnerability to exploitable path

A vulnerability does not exist in isolation. Its practical importance depends on factors such as asset criticality, reachability, identity relationships, privileges, configuration, security controls and the paths through which an adversary could exploit it. CTEM therefore asks a broader question than traditional vulnerability management: **which combinations of exposure could allow a realistic attacker to reach something that matters?**

**The additional question deception can answer**

Which of those technically possible paths are attackers actually discovering, testing or attempting to use?

Technical possibility and observed adversary interest are not the same thing. A theoretically exploitable path may never be reachable from likely attacker footholds. Conversely, a seemingly minor weakness may form part of an attractive chain towards privileged identities or critical data. Deception telemetry can provide behavioural evidence that helps distinguish between these situations.

This distinction is important because technical possibility and observed adversary interest are not the same thing.

A theoretically exploitable path may never be reachable from likely attacker footholds. Conversely, a seemingly minor weakness may form part of an attractive chain towards privileged identities or critical data. Deception telemetry can provide behavioural evidence that helps distinguish between these situations.

## How deception contributes across the CTEM cycle

| CTEM activity | Contribution from deception |
|---|---|
| Scoping | Use critical services, crown jewels, identities and trust boundaries to determine where deceptive coverage provides meaningful evidence. |
| Discovery | Reveal reconnaissance, lateral movement and overlooked credentials or relationships that inventory and scanner data may miss. |
| Prioritisation | Add behavioural evidence to severity, business criticality, reachability, exploitability and attack-path context. |
| Validation | Use decoys and controlled attack-path tests to verify reachability, segmentation, detection quality and SOC readiness. |
| Mobilisation | Translate findings into remediation, playbook changes, privilege reduction, segmentation and control tuning; then verify the effect. |

**Avoid a category mistake**
A deception platform is not a complete CTEM programme. CTEM requires governance, exposure discovery, prioritisation, validation and mobilisation across people, process and technology. Deception is one useful source of evidence and validation within that cycle.

### Scoping: start with business impact

CTEM begins by defining what matters. Critical business services, sensitive information, privileged identities, important workloads and significant trust relationships determine which exposure scenarios deserve attention. This should also determine deception placement.

Rather than deploying decoys uniformly across the environment, organisations can identify likely attacker routes towards crown jewels and place deceptive identities, credentials, systems, services or data artefacts where interaction would be particularly meaningful. **The result is risk-driven deception coverage rather than decoy density.**

### Discovery: expose hidden relationships

Exposure discovery identifies assets, identities, vulnerabilities, misconfigurations, external attack surface and other conditions that could contribute to compromise.

Deception adds a behavioural discovery mechanism. Interaction may expose relationships or attack paths that inventories and scanners did not identify, including unexpected lateral movement, credential reuse, unmanaged administrative routes or overlooked trust relationships.

**It can therefore reveal not only what exists, but also how an adversary attempts to navigate what exists.**

## Prioritisation: add behavioural evidence

Traditional vulnerability severity is only one input into risk prioritisation. CTEM adds business importance, exploitability, reachability and attack-path context. Deception can add observed behaviour.

For example, two technically similar exposure paths may have very different operational significance if one repeatedly attracts reconnaissance or credential activity while the other remains inaccessible from realistic footholds. That does not automatically make every deception interaction the highest priority. Rather, it introduces another evidence source into the prioritisation decision. Together, these provide a richer basis for determining what should be addressed first.




*Figure 4: Evidence source into the prioritisation decision.*

## Validation: move from assumption to evidence

Validation is where deception has one of its strongest relationships with CTEM. Security teams routinely make assumptions:

- This segment cannot reach the management network.
- These credentials cannot be obtained from ordinary endpoints.
- This attack path would be detected.
- This control prevents lateral movement.

Controlled attack simulation, purple-team activity and deception can turn those assumptions into testable hypotheses.

A deceptive resource placed at the end of a suspected attack path can help determine whether the route is discoverable, reachable and monitored. If a controlled test reaches the resource without triggering the expected controls, the organisation has identified a gap not merely in theory, but in the operation of its defensive architecture. **Deception therefore helps CTEM move from "we believe this exposure matters" towards "we have evidence showing whether the path and controls behave as expected".**

### Mobilisation: turn evidence into reduction

Finding and validating an exposure creates no risk reduction until something changes. The mobilisation stage translates findings into patching, configuration changes, privilege reduction, segmentation, detection improvements, architectural changes, revised SOC playbooks or accepted risk. Deception provides a useful feedback mechanism after those actions have been implemented. Lures can be repositioned, attack paths retested and telemetry monitored to determine whether the remediation changed attacker reachability or defensive visibility.

This creates a closed loop. That closed loop is more important than any individual deception alert.

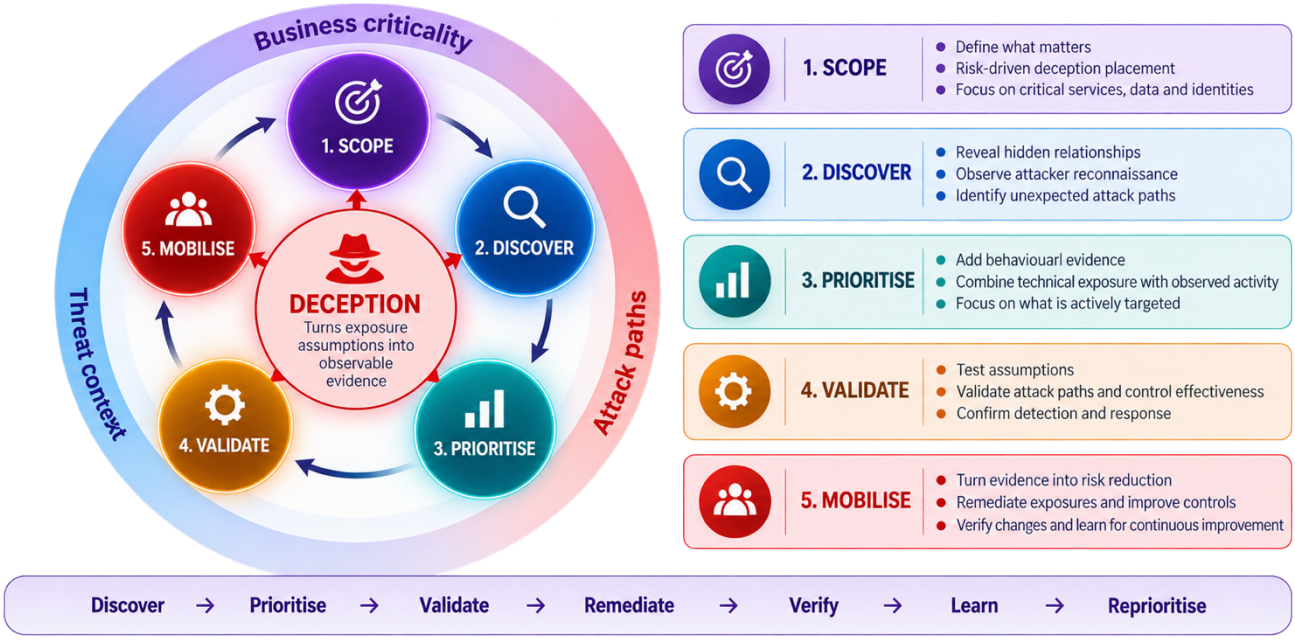



*Figure 5: CTEM deception feedback loop.*

## From exposure management to measurable resilience

The strategic relationship between CTEM and deception goes beyond finding vulnerabilities. CTEM determines where exposure matters. Deception helps determine whether attacker behaviour makes that exposure visible in practice. Security controls reduce or contain the exposure. The SOC responds when malicious behaviour occurs. Validation determines whether those controls actually worked. Governance then uses the resulting evidence to adjust risk decisions and priorities.

This is also why deception should not be measured primarily by the number of decoys deployed. More meaningful indicators include the percentage of critical attack paths instrumented, validated control failures, time from interaction to triage, time to containment and the number of material security improvements resulting from deception findings.

### Avoid a category mistake

A deception platform is not a complete CTEM programme. CTEM requires governance, exposure discovery, prioritisation, validation and mobilisation across people, process and technology. Deception is an evidence and validation capability within that cycle.

# Bringing the three together

Zero Trust determines how trust is established, limited and continuously reconsidered. CTEM determines which exposures and attack paths matter most and drives their reduction. Deception introduces controlled uncertainty for the attacker and high-value evidence for the defender. Together, they create a closed loop in which access is constrained, exposure is continuously evaluated, attacker movement becomes observable and findings are translated into measurable security improvement.

**Core message**
Zero Trust constrains trust. CTEM manages exposure. Deception makes adversary interaction with that environment observable and testable.

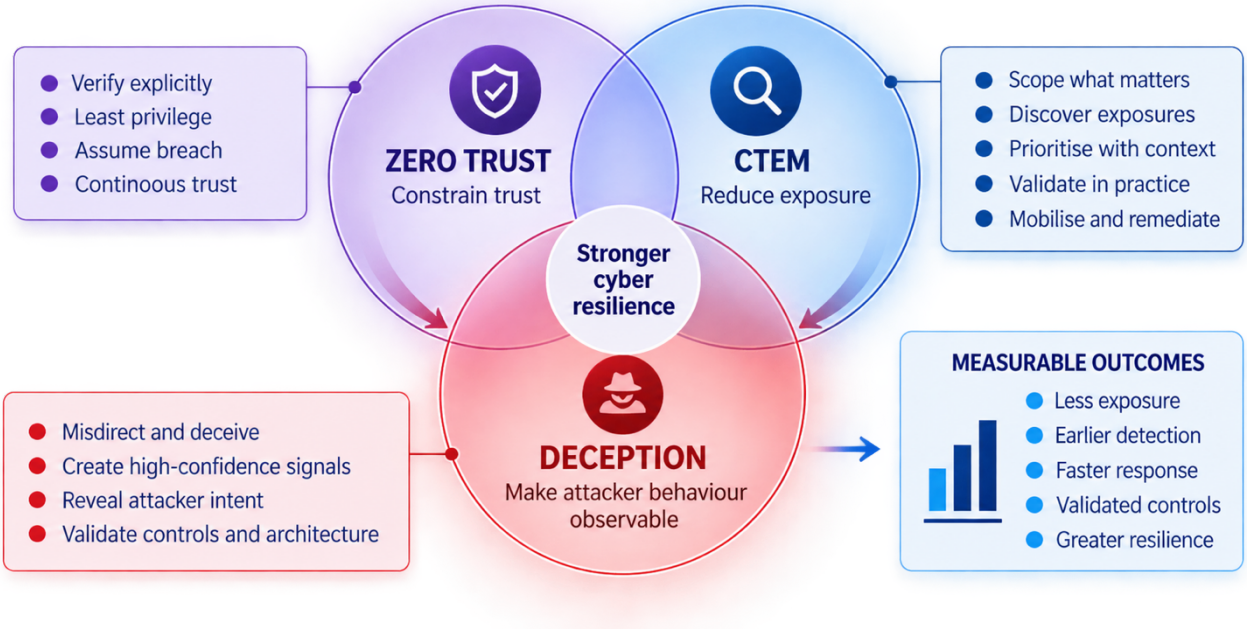



*Figure 6: Zero Trust, deception and CTEM form complementary parts of a closed security loop.*

# SOC, incident response and threat intelligence

## From deception signal to operational value

The operational value of deception technology is realised in the Security Operations Centre (SOC) and incident-response process. A deception interaction is not the end product, but the beginning of an investigation. When integrated with identity, endpoint, network, exposure and threat-intelligence data, deception signals can provide high-confidence indicators of suspicious activity, focus investigation and help defenders decide whether containment is required.

This is an important distinction. A platform that produces high-quality alerts but remains outside established triage and response processes simply becomes another console. Before deployment, the organisation should therefore define severity, ownership, enrichment, escalation, containment authority and evidence handling. The objective is not to create a separate deception workflow, but to make deception part of the existing detection-and-response operating model.

A typical investigation starts by validating whether the interaction can be explained by authorised activity such as vulnerability scanning, discovery tooling, backup processes, red-team exercises or administrative testing. If not, the SOC correlates the source with identity, endpoint and network telemetry, determines whether the deceptive resource was merely discovered or actively used, scopes related activity and assesses attacker intent. Higher-interaction telemetry may reveal tools, commands and techniques, but collection should remain proportionate to the defensive purpose and governed accordingly.

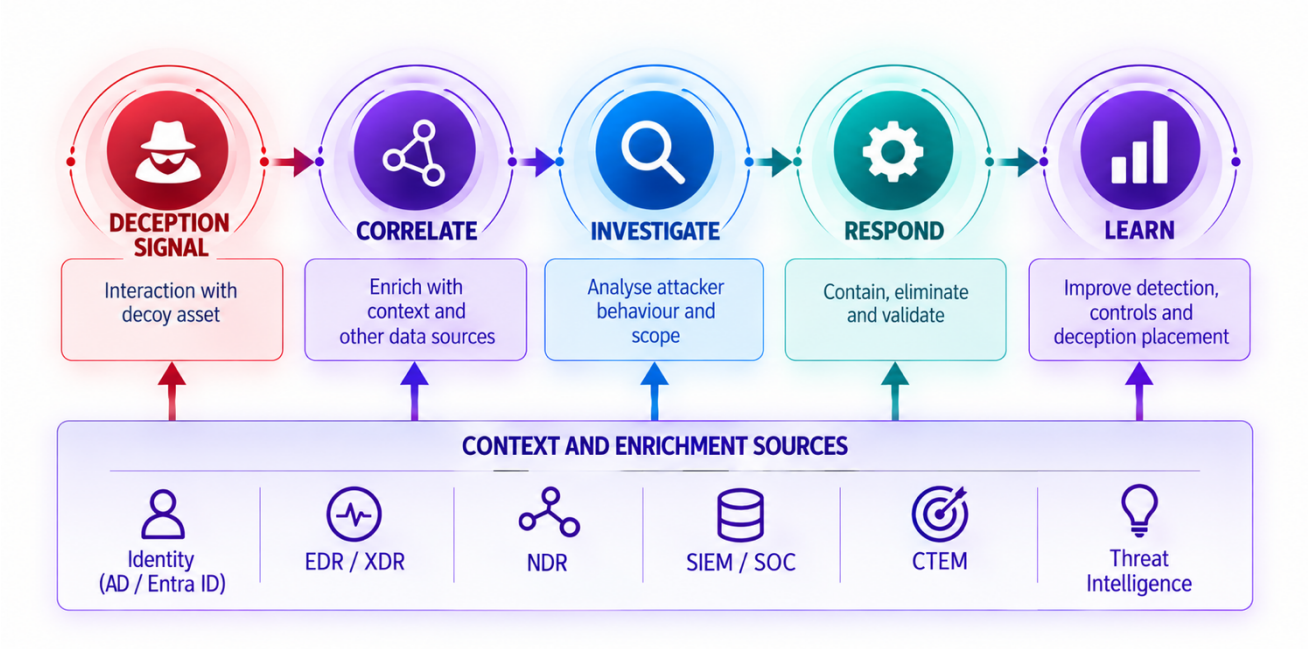



*Figure 7: From deception signal to response and learning.*

## Integration in the SOC workflow

Deception signals should enter the same case-management as SIEM, XDR or SOAR processes used for other security events. Their value comes from signal quality and context, not from isolation. Correlation with identity data can show whether a credential belongs to a real user or a synthetic identity Endpoint telemetry can reveal the originating process, NDR can expose lateral movement, CTEM data can show whether the source sits on a known attack path and threat intelligence can help interpret observed tools and techniques.

| # | Level | Characteristic |
|---|---|---|
| 1 | Detect | Receive the deception interaction, validate signal integrity and perform initial enrichment. | A credible event with known source, target and interaction type. |
| 2 | Correlate | Correlate with identity, endpoint, network and exposure context. Assign severity and scope. | A prioritised case with enough context for investigation. |

| # | Level | Characteristic |
|---|---|---|
| 3 | Investigate | Analyse attacker behaviour, identify related activity and attack paths, determine likely intent and impact. | An evidence-based understanding of what happened and what may be affected. |
| 4 | Respond | Contain compromised assets or identities where warranted, block attacker access and preserve relevant evidence. | Controlled containment with minimum unnecessary business impact. |
| 5 | Learn | Update detections, deception placement, CTEM priorities, hardening and response playbooks. | A closed learning loop that improves the wider security environment. |

## Use cases in the SOC and incident response

The exact use cases depend on architecture, threat model and critical assets, but several patterns recur across environments. Deceptive credentials can reveal credential theft and attempts to move laterally. Decoy systems and services can expose internal reconnaissance, privilege escalation and unauthorised movement. Decoy data sources can indicate attempts to discover or exfiltrate sensitive information. Cloud deception can make unauthorised exploration of cloud resources more visible, while deceptive assets associated with privileged or unusual access paths can help identify malicious insider activity or compromised accounts.

**High confidence does not mean automatic guilt**

Interaction with a well-designed deceptive resource can be a strong security signal, but context remains essential. Known scanners, authorised testing and configuration mistakes can still create legitimate explanations. Response should therefore be risk-based and proportionate rather than automatically destructive.

## From operational signals to threat intelligence

Deception can also contribute to internal threat intelligence. Observed behaviour may reveal attacker tooling, techniques, command sequences, credential-use patterns and preferred attack paths. That information can enrich detection engineering, inform purple-team exercises and guide where future deception should be placed. Where appropriate and lawful, selected indicators can also be shared with trusted partners or sector communities.

The relationship is bidirectional. Threat intelligence helps determine which services, identities, artefacts and attack paths should be made believable to an adversary; deception then provides local behavioural evidence about what is actually being explored. This feedback loop is particularly valuable because it connects external knowledge of threats with evidence from the organisation's own environment.

# Governance and policy: the foundation

Deception should be governed as part of the organisation's information security management system, not as an isolated technical experiment. The technology intentionally creates false assets and may collect detailed activity. That requires a defined purpose, clear ownership, acceptable-use boundaries, lifecycle management and integration with existing security policies and processes.

Governance is needed throughout the lifecycle. During design, the organisation determines why deception is being used and which risks it should address. Approval confirms that architecture, privacy, operational and legal considerations have been assessed. Deployment places and integrates deceptive assets under controlled conditions. Monitoring ensures interactions and platform health are visible. Response defines what happens when a signal is generated. Review assesses effectiveness and changing risk, while retirement removes stale or unnecessary assets and credentials.

| Lifecycle stage | Governance focus |
|---|---|
| Design | Define objectives, scope, use cases, ownership, success criteria and prohibited uses. |
| Approval | Review risk, privacy, architecture, operational impact and segregation of duties. |
| Deployment | Implement securely; document decoys, lures, integrations, logging and expected legitimate interaction. |
| Monitoring | Track platform health, interactions, false or expected triggers and changes in the environment. |
| Response | Apply approved triage, escalation, containment, evidence and communication procedures. |
| Review | Assess value, control effectiveness, attack-path coverage and required changes. |
| Retirement | Remove obsolete decoys, credentials, routes and integrations without leaving unmanaged artefacts. |

## Policy integration

A mature deployment should connect to existing policy rather than create a parallel governance structure. Logging and monitoring policy should define which deception events are retained, where they are forwarded, retention periods, time synchronisation and evidential integrity. The Incident response policy should define alert severity, triage, escalation, containment authority and forensic preservation. A privacy policy should address purpose limitation, minimisation, lawful basis, access and retention. Change and vulnerability management should cover the deception infrastructure itself, while supplier management should address service location, subprocessors, support access, incident notification and exit requirements where a third party is involved.

### Rules of engagement

An organisation should explicitly state that deception is a defensive monitoring capability for its own authorised environments. It should not be used to hack back, retaliate against external systems or collect data beyond what is necessary for the defensive purpose. High-interaction environments require particularly strong network segmentation and egress controls so that a compromised decoy cannot become a pivot point or a platform for activity against third parties.

**Policy first, technology second**

Every deceptive asset should have an owner, a purpose, an expected legitimate-interaction profile, a logging destination, a response path, a review point and a retirement process. Governance should remain valid for the full lifecycle, not only at deployment.

## Compliance and EU laws and regulations

Deception technology is not a compliance framework, and deployment does not by itself demonstrate compliance. Its value lies in supporting specific control objectives such as detection, incident handling, effectiveness assessment, access monitoring, attack-path validation and evidence. The organisation must still determine which legal and sector-specific obligations apply and whether the chosen implementation is appropriate, proportionate and properly governed.

A useful way to position deception is to separate requirement, control, operation, evidence and assurance. Laws, regulations and frameworks define requirements and expected outcomes. Organisational and technical controls implement those expectations. Deception can then provide operational evidence that selected controls are functioning or reveal where assumptions fail. That evidence can support assurance, but it does not transfer accountability from governance to the technology.

### NIS2 and national implementation in the Netherlands and Belgium

NIS2 requires organisations within scope to implement appropriate and proportionate technical, operational and organisational measures. These cover areas such as risk analysis, incident handling, business continuity, supply chain security, secure acquisition and maintenance, effectiveness assessment, cyber hygiene, cryptography, access control and, where appropriate, multi-factor or continuous authentication. Member States transpose these requirements into national legislation, which means that the precise legal and supervisory context differs by country.

In the Netherlands, NIS2 is implemented through the **Cyberbeveiligingswet (Cbw)**, which entered into force on 15 August 2026. The Cbw requires organisations within scope to manage cybersecurity risks through appropriate and proportionate measures and introduces obligations relating to registration, incident reporting and governance. In Belgium, NIS2 is implemented through the **NIS2 Law of 26 April 2024**, which entered into force on 18 October 2024 and is supplemented by the Royal Decree of 9 June 2024. The Centre for Cybersecurity Belgium (CCB) acts as the national cybersecurity authority and has an important role in implementation, coordination and supervision.

Deception can support several of the security objectives underlying NIS2 and both national implementations, particularly incident detection and handling, assessment of control effectiveness, network and information system security, access monitoring and evidence for continuous improvement.

It should, however, be treated as a contributing security capability rather than as evidence of compliance in itself or as a substitute for the broader organisational, governance and risk-management duties imposed by NIS2 and national legislation.

## GDPR / AVG and national data protection legislation

The General Data Protection Regulation (GDPR), known as the Algemene Verordening Gegevensbescherming (AVG) in Dutch and the Règlement général sur la protection des données (RGPD) in French, applies directly throughout the European Union and therefore provides a common data-protection foundation across the Benelux. It requires organisations to implement appropriate technical and organisational measures to protect personal data, taking into account the risks associated with its processing. These requirements are complemented by national legislation, including the Uitvoeringswet Algemene verordening gegevensbescherming (UAVG) in the Netherlands and the Data Protection Act of 30 July 2018 in Belgium.

Deception can contribute to this security objective by helping organisations detect unauthorised access, credential misuse, lateral movement and suspicious interaction with systems or data. Deceptive data objects and credentials can also provide early warning of attempts to access or exfiltrate information. This can strengthen monitoring and incident-detection capabilities around environments in which personal data is processed.

Care is nevertheless required. Deception telemetry may itself contain personal data, such as usernames, IP addresses, device information or behavioural information. Its collection, retention and use should therefore be governed by appropriate purposes, access controls, retention periods and data-minimisation principles. Deception can support the security of personal data and provide useful evidence of control effectiveness, but its deployment does not in itself demonstrate compliance with the GDPR/AVG or applicable national data-protection legislation.

## DORA and other EU requirements

For financial entities in scope, DORA places strong emphasis on ICT risk management, detection, incident response and testing. Deception can strengthen anomaly detection and provide practical evidence during testing, particularly when connected to established response workflows. The Cyber Resilience Act has a different focus: it primarily concerns products with digital elements and their manufacturers. It is therefore not an implementation framework for an organisation using deception internally, although its product-security and vulnerability-handling requirements may be relevant when selecting technology and assessing suppliers.

### Compliance contribution

Deception is strongest as evidence of operation and effectiveness: it can show that suspicious movement is detected that response paths work and that architectural assumptions hold in practice. Whether this evidence is sufficient for a particular legal or regulatory obligation remains a governance and assurance decision.

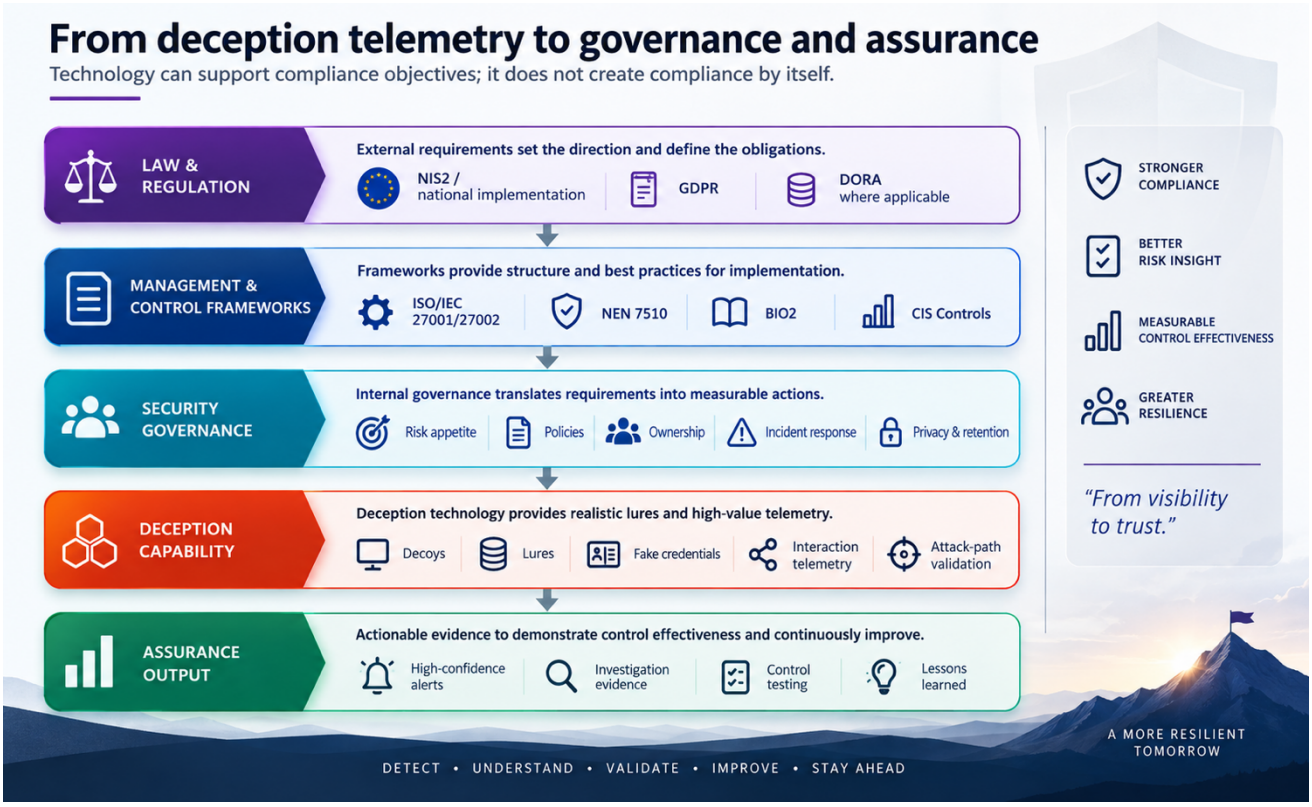



*Figure 8: Deception telemetry can support assurance, but compliance remains a governance responsibility.*

# Mapping deception to common frameworks

Framework mapping is most useful when it shows both contribution and limitation. The table below therefore does not claim that deception satisfies a framework requirement. It shows where deception can provide supporting capability or evidence and where broader controls remain necessary.

| Framework / area | Possible contribution from deception | Example evidence | Important limitation |
|---|---|---|---|
| NIS2 / Cbw | Detection, incident handling, control-effectiveness testing, access and network monitoring. | Interaction logs, SOC cases, exercised runbooks, validated attack paths. | Does not replace risk management, governance, continuity, supply-chain or training duties. |
| ISO 27001 / 27002 | Supports monitoring, event analysis, incident management, access-control assurance and technical testing. | Alert records, correlation evidence, response results, control-validation exercises. | Certification depends on the complete ISMS and applicable controls. |

| Framework / area | Possible contribution from deception | Example evidence | Important limitation |
|---|---|---|---|
| NEN 7510 | Can strengthen monitoring and detection around identities, systems and sensitive health information. | Evidence of detection, investigation and tested response around critical information assets. | Healthcare-specific governance, privacy and continuity obligations remain separate. |
| BIO2 | Supports monitoring, incident response and assurance of technical controls in public-sector environments. | SOC evidence, attack-path tests and documented remediation. | Must be embedded in the organisation's broader baseline and risk-management approach. |
| CIS Controls | Can complement account, network, logging, monitoring and incident-response safeguards. | Detection of credential misuse, lateral movement and control-testing outcomes. | Does not replace foundational inventory, configuration, vulnerability or recovery safeguards. |
| DORA | Can support anomaly detection, incident-response readiness and testing of ICT controls. | Detection events, response timing, exercised controls and improvement records. | Only one component of the wider ICT risk-management and resilience framework. |

**Evidence is not the same as assurance**
An alert, log or test result only becomes useful compliance evidence when its control objective, scope, ownership, retention, integrity and follow-up are understood and documented.

# Implementation prerequisites

Deception can be deployed relatively quickly from a technical perspective, but successful implementation depends on organisational readiness. The minimum prerequisites are less about having a perfect security stack and more about knowing what is important, who will act on signals and how the new capability fits existing architecture.

| Prerequisite | Why it matters |
|---|---|
| Defined business and security objectives | Prevents deployment volume from becoming the goal; use cases should link to risk and critical processes. |
| Asset and identity visibility | Believable placement and meaningful investigation require knowledge of systems, owners, identities and dependencies. |

| Prerequisite | Why it matters |
|---|---|
| Critical-asset identification | Determines where early-warning coverage matters most and where deception must never create operational dependency. |
| Threat scenarios and attack paths | Guides which lures and decoys are plausible and where they should be placed. |
| Segmentation and egress control | Prevents a compromised decoy from becoming a pivot or a platform for external abuse. |
| Central logging and SOC workflow | Ensures deception events are correlated, triaged and acted upon rather than stranded in a separate console. |
| Incident response runbooks | Defines what happens after a high confidence interaction and who may contain systems or identities. |
| Change and configuration management | Prevents stale lures, unrealistic decoys and unmanaged credentials from reducing trust in the capability. |
| Privacy and retention decisions | Sets proportionate collection, access and retention for detailed activity telemetry. |
| Skills and ownership | Requires people who can design believable deception, interpret attacker behaviour and maintain the environment. |

## Implementation roadmap

A phased approach reduces risk and makes value measurable. Begin with a small number of high-value use cases rather than broad coverage. A good first pilot often focuses on lateral movement towards critical server segments, misuse of privileged credentials, ransomware-style discovery or unauthorised access to sensitive data locations.

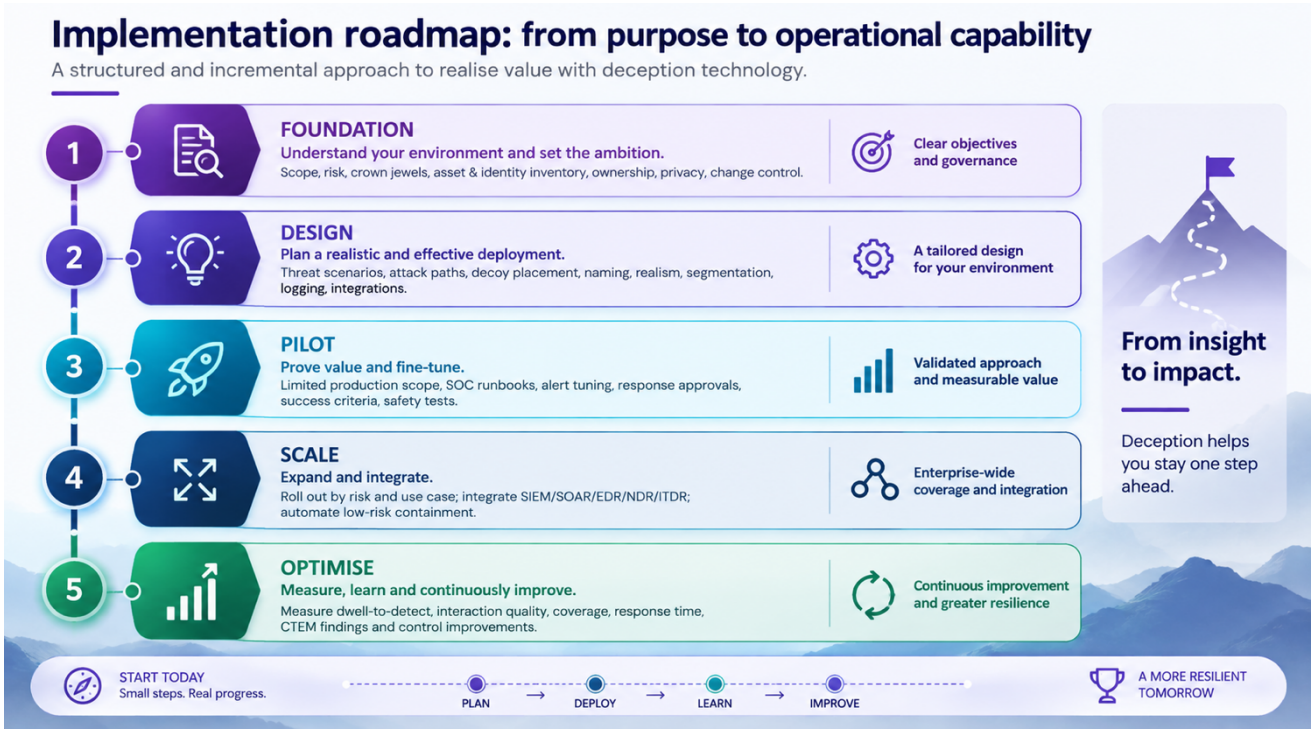



*Figure 9: A risk-driven implementation path for deception technology.*

| # | Phase | Objective | Key activities | Exit criterion |
|---|---|---|---|---|
| 1 | Foundation | Establish purpose and governance. | Agree scope, ownership and success criteria; identify crown jewels and attack paths; confirm privacy, logging, change and response requirements. | Approved use cases, owners and governance boundaries. |
| 2 | Design | Translate risk into a controlled architecture. | Define decoy placement, lures, synthetic identities, naming, realism, segmentation, egress controls and integrations. | Documented design and mapped SOC response for each interaction type. |
| 3 | Pilot | Validate value in a bounded production scope. | Test benign tooling, exercise runbooks, assess evidence quality and verify that response actions do not create unacceptable impact. | Demonstrated signal quality and repeatable operational handling. |

| # | Phase | Objective | Key activities | Exit criterion |
|---|---|---|---|---|
| 4 | Scale | Expand coverage by risk. | Add critical attack paths, cloud, identity, sensitive data and third-party access; automate repetitive enrichment where appropriate. | Coverage aligned with prioritised risks and sustainable operations. |
| 5 | Optimise | Create a continuous improvement loop. | Use incidents, purple-team exercises, threat intelligence and CTEM results to reposition lures, improve realism and retire stale assets. | Evidence that deception measurably improves controls, response and resilience. |

## Measuring value and maturity

Deception maturity is not the number of decoys deployed. A large estate of stale or unmonitored decoys can create less value than a small, well-governed deployment on critical attack paths. Measurement should therefore combine coverage, signal quality, operational response and improvement outcomes.

A useful reporting model connects technical measures to operational outcomes and ultimately to risk and resilience. A detected interaction is a technical event. Faster triage and containment are operational outcomes. Reduced opportunity for lateral movement, improved control assurance and better protection of critical services are the resilience outcomes that matter to management.

| Measure | Example question | Management relevance |
|---|---|---|
| Attack-path coverage | Are the most important lateral movement and privilege escalation routes instrumented with meaningful deceptive signals? | Shows whether coverage follows business risk rather than asset count. |
| Signal quality | What proportion of interactions are expected tooling or operational noise versus genuinely suspicious behaviour? | Indicates whether the capability is trusted and operationally sustainable. |
| Time to triage | How quickly does the SOC validate a deception alert and correlate identity, endpoint, network and exposure context? | Measures whether high-confidence signals translate into faster decisions. |

| Measure | Example question | Management relevance |
|---|---|---|
| Time to contain | How long from a confirmed interaction to containment when containment is warranted? | Connects detection to reduction of attacker opportunity. |
| Control validation | Do exercises demonstrate that SIEM, EDR/NDR, identity, segmentation and response controls operate as designed? | Provides evidence of control effectiveness, not only control presence. |
| Learning loop | What material improvements to access, segmentation, hardening, detection or CTEM priorities result from deception findings? | Shows whether evidence produces measurable improvement. |

### A practical maturity model

| # | Level | Characteristic |
|---|---|---|
| 1 | Initial | Ad hoc honeypots or isolated experiments; unclear ownership and little SOC integration. |
| 2 | Basic | Defined use cases and central management; limited production coverage and documented triage. |
| 3 | Integrated | Deception events correlate with SIEM/XDR/identity/network context; runbooks and change processes are established. |
| 4 | Adaptive | Placement and response are adjusted using threat intelligence, attack paths and CTEM; selected response is automated. |
| 5 | Optimised | Closed-loop governance: deception evidence continuously informs risk, control assurance, architecture and measurable improvement. |

## Practical considerations and limitations

Deception is most effective when its limitations are treated as design requirements rather than afterthoughts. Decoys need hardening, patching, isolation, restricted egress and protected management interfaces. Believability requires maintenance. Stale naming, impossible services or unrealistic credentials reduce value. Legitimate applications and users should never depend on deceptive assets or credentials.

Operational tooling can also interact with decoys. Vulnerability scanners, discovery platforms, backup systems and authorised security testing should therefore be understood and governed. Captured sessions and identity data may be security-sensitive and may contain personal data. Finally, a high-

confidence alert still requires investigation; destructive containment should not be automated without appropriate confidence, governance and testing.

**Measure outcomes, not objects**
Decoy count is an implementation metric. The more meaningful questions are whether important attack paths are observable, whether investigations become faster and more accurate, whether controls are validated and whether findings lead to demonstrable risk reduction.

# Conclusion: from hidden movement to actionable evidence

Deception technology gives defenders a different way to observe an attack. Instead of trying only to distinguish malicious behaviour from millions of legitimate events, it introduces believable resources that have little or no legitimate reason to be touched. That changes the signal-to-noise problem and can reveal reconnaissance, credential abuse and lateral movement before critical assets are reached.

Its strategic value is strongest when it is treated as part of a broader security model. Within Zero Trust, deception makes “assume breach” more observable and provides additional context for revocable trust. Within CTEM, it adds behavioural evidence and validation to exposure management. Within the SOC, it can focus investigation and accelerate containment. Within governance and compliance, its telemetry can help demonstrate testing and control effectiveness - provided ownership, purpose and evidence are properly defined.

The technology is therefore not a replacement for prevention, identity governance, segmentation, endpoint security, vulnerability management or incident response. It is a complementary capability that makes attacker movement more visible and more costly while giving defenders better evidence.

**Take Control of your Cyber Resilience**
Start with the question “which attacker movements do we need to make visible before they reach our critical data and processes?” Then design the deception capability around those risks, not around the number of decoys a platform can deploy.

## Role of Checkdone IT

Checkdone IT supports organisations independently and vendor-agnostically in assessing where deception can add value, defining governance and use cases, mapping the capability to Zero Trust and CTEM, designing the target architecture, selecting suitable technology and integrating it into monitoring and incident response. The focus remains on risk, practical operability and demonstrable cyber resilience rather than on a product category in isolation.

# Glossary

| Term | Description |
|---|---|
| Attack path | A sequence of reachable systems, identities, privileges and weaknesses that could allow an attacker to move from an initial foothold towards a target. |
| Breadcrumb / lure | A deliberately discoverable artefact that points an attacker towards a controlled deceptive resource. |
| CTEM - Continuous Threat Exposure Management | A continuous programme for scoping, discovering, prioritising, validating and mobilising the reduction of exploitable exposure. |
| Decoy | A controlled fake or emulated asset designed to resemble a real system, service, application, identity or data resource. |
| Deception technology | A defensive security capability that uses believable false resources and information to detect, misdirect and observe unauthorised activity. |
| Defence in Depth | The use of multiple complementary security layers so that failure of one control does not automatically lead to serious impact. |
| Honeytoken | A token, credential or data artefact with no legitimate operational purpose whose use can indicate suspicious activity. |
| Honeypot | A system or resource designed to be attractive to intruders and monitored for interaction. |
| High-interaction decoy | A deceptive system that permits realistic attacker interaction to collect richer behavioural evidence. |
| Lateral movement | Techniques used by an attacker to move from one compromised identity or system to others within an environment. |
| PoLP - Principle of Least Privilege | The principle that users, systems and workloads receive only the minimum access necessary, preferably only for the required period. |
| Purple teaming | Collaborative testing in which offensive and defensive security teams work together to validate and improve controls. |
| Telemetry | Technical and operational data that describes events, behaviour, status and context. |
| Zero Trust | A security governance and architecture model that replaces implicit trust with explicit, minimal, context-dependent and revocable decisions. |

# Sources, reference frameworks and disclaimer

Primary public references consulted for the technical and regulatory content include:

- NIST Cybersecurity Resource Center - Honeypot glossary definition.
- NIST SP 800-160 Vol. 2 Rev. 1 - Developing Cyber-Resilient Systems: A Systems Security Engineering Approach (misdirection, disinformation, decoys and honeytokens).
- ENISA - Proactive detection / honeypot guidance and research on digital traps for detecting cyber-attacks.
- Directive (EU) 2022/2555 (NIS2), particularly Article 21 on cybersecurity risk-management measures.
- Nationaal Cyber Security Centrum (Netherlands) - Cyberbeveiligingswet (NIS2), duty of care and the ten measures; Cbw in force from 15 August 2026.
- Regulation (EU) 2016/679 (GDPR), particularly Article 32 on security of processing.
- Regulation (EU) 2022/2554 (DORA), particularly requirements for ICT risk management, anomaly detection and testing.
- Regulation (EU) 2024/2847 (Cyber Resilience Act), for product cybersecurity and vulnerability-handling context.
- ISO/IEC 27001 and ISO/IEC 27002; NEN 7510; BIO2; CIS Critical Security Controls - where applicable to the organisation.
- Gartner research on Continuous Threat Exposure Management (CTEM) - for the programme concept and exposure-management lifecycle.

**Disclaimer**

This white paper is a practical guide and provides general information about deception technology and its relationship with security governance, compliance frameworks and relevant laws and regulations. It is not legal advice, a formal compliance assessment or a substitute for current legislation, sector-specific requirements or licensed standards. Organisations remain responsible for determining which obligations, controls and privacy requirements apply to their specific circumstances.

CHECKDONE IT

# Control over cyber resilience without any vendor strings attached

Checkdone IT is the independent partner that brings structure and control to information security and IT security. We help organisations understand where their risks lie, which measures are genuinely necessary, and how policy, technology and behaviour can work together as a whole.

**INDEPENDENT**

We work in the client's best interests, not for vendors.

**FROM INSIGHT TO IMPLEMENTATION**

From analysis and decision-making to implementation and embedding.

**INTEGRAL**

Technology, policy and behaviour in context.

**PRACTICAL & VERIFIABLE**

Focused on measures that work and evidence that is actionable.

**Grip on IT security**

No sales pitch and no piling on of solutions. Instead, insight, clear choices and control over the whole picture.

**Checkdone IT B.V.**

Nieuwe Eerdsebaan 8a

5482 VS Schijndel

T: +31 (0)73 549 59 58

E: info@checkdone-it.eu

https://www.checkdone-it.eu/