

A Buyer's Checklist

for Evaluating NDR Solutions

The most effective system for
comparing NDR solutions

Introduction

The Network Detection and Response (NDR) market has been saturated with noisy AI hype and vendor claims, making it a growing challenge for CISOs and IT decision-makers to find truly effective solutions that are best suited to their needs.

This guide presents a structured evaluation framework and practical checklist to support informed, objective comparisons of NDR vendors based on measurable capabilities. While each team and organization will have unique use cases, these criteria will act as a starting point for evaluating and comparing vendors in the market.

02	Part 01: Methodology
03	Analysis methodology
04	Part 02: Assessment Categories
04	Organizational roadmap and ROI assessment
05	Scope and deployment
07	Network behavioral analysis
09	Threat detection
11	Triage and investigation
12	Containment and response
13	Additional functionality
14	Bonus Section: How to evaluate AI
16	Part 03: The Evaluation Process
16	Step 01: Commercial evaluation
17	Step 02: Technical evaluation
18	Conclusion

Methodology

How to use this document

This checklist is designed to help you assess and compare modern NDR solutions.

Part 2 of this document will contain the assessment categories split into sections: organizational roadmap & ROI, scope and deployment, network behavioral analysis, detection, alerting, and investigations, threat detection, and additional functionality.

Each category will have a description and a list of questions to ask an NDR vendor. We recommend highlighting the capabilities you are looking for and using the questions we provide as part of your evaluation criteria.

Here's how you can make the most out of this document:



Initiate vendor discussions

Use the checklist as a **conversation starter** with potential vendors. Ask open-ended questions based on the criteria to better understand how each solution aligns with your organization's priorities.



Compare solutions

Use the checklist to **compare different solutions** in a structured manner, to identify where each solution meets, exceeds, or falls short of your requirements.



Clarify and validate concerns

If you have concerns or uncertainties about specific features, refer to the relevant criteria in the checklist to **guide deeper discussions** and request further details or demonstrations. Ensure that vendors are providing accurate and substantiated information.



Prioritize capabilities

Consider assigning a priority level - High, Medium, or Low - to each criterion based on your organization's risk profile and security objectives. This will help **focus** your evaluation **on the capabilities that matter most**.

To ensure a fair and objective assessment of network security solutions, it's **useful to adopt an A/B comparison approach** – evaluating vendors side-by-side against the same set of criteria.

This method helps eliminate bias, highlights real capability differences, and creates a level playing field for evaluation.

Here are some factors to consider to ensure a fair comparison:

Deployment

Establish the same scope: Ensure all solutions have the same visibility across relevant subnets and desired network entities.

Establish a concrete timeline: Test all solutions in the same environment, for the same number of days. Similar timelines allow for the comparison of learning periods and evolution of learning, which may differ depending on the solution.

Tuning and configurations

Compare on optimal performance: Consult your selected vendors on what an optimal deployment looks like and adjust accordingly, especially if you have a complex network or unique outcomes you'd like to achieve.

Consider the long term: Adjust the solution with the long term in mind. A short fix or configuration process will not reveal the future state of the configuration requirements.

Data visibility: Ensure all solutions can retain and have search capabilities on all the data throughout the evaluation period.

Evaluation

Coordinate your decision criteria: All vendors must be held to the same previously established evaluation criteria.

Understand vendor threat classification: Different vendors will have different threat scoring, tags, explanations, and criteria. Make sure to fully understand the nuances to ensure an equivalent comparison.

Consider active vs. passive functionality: Some capabilities may not be active during a trial or proof of concept but could have a substantial impact for a live deployment. Response action are one example of this. Consider how any passive features would affect how the solution handles attacks in the real world when they are made active.

Maintain communication: Communicate with the corresponding vendors if results do not seem accurate. This will show the vendor's capability for support and remediation.

Conclusion

Evaluate the solution based on results: Compare results based on your agreed success criteria and metrics. These could include reduction of alert noise, number of critical incidents raised, and total analyst investigation time saved across the test period.

Evaluate qualitative areas: Go beyond the numbers to consider the entire user experience, relationship with the vendor's team, what support is available, integrations, and the vendor's wider platform capabilities.

Review evaluation criteria and results: Review internally but also ensure to communicate the results to your selected vendors.

Assessment Categories

The following checklist outlines key capabilities that organizations should consider when evaluating modern NDR solutions. Use this as a structured point of comparison between vendors, not just to confirm the presence of a feature, but to assess the depth, effectiveness, and approach each solution takes in delivering it. Some vendors may claim that they have a particular feature or capability, however when you dig deeper, they may not equate to relevant security outcomes.

While not every feature needs to be checked off by a single provider, this checklist can help you establish your expectations for your network security stack.

Organizational roadmap and ROI assessment

Criteria	Questions for vendors
Business case parameters	Does the vendor include ROI modeling, total cost of ownership, licensing model clarity, and operational resource burden?
Transparency around how AI is used for security operations	- Does the vendor articulate how AI is used in detection, alerting, and response. Includes details on model types (e.g., supervised, unsupervised, rule-based), training data sources, how false positives are handled, and mechanisms for analyst oversight? - Is the vendor compliant with one or more of the following regulations regarding AI: - ISO/IEC 42001:2023 - AI Management System Standard - ISO/IEC 23894:2023 - AI Risk Management - Does the vendor document AI limitations, explainability features, and how updates are communicated or governed?
Vendor transparency & roadmap alignment	Does the vendor have full product documentation, regular product release cadence, long-term support model, and customer influence over roadmap?
Product scalability & performance metrics	Does the vendor support growth across environments with metrics like throughput, sensor density, latency, and ingestion capacity?
Pricing scalability	Does the vendor support a scalable licensing model with volume-based and multi-year discounts?
Deployment flexibility	Does the vendor support multi-environment deployment (i.e: on-premises, virtual networks, containerized environments, hybrid cloud, remote devices, email, OT devices, air-gapped environments)?
Licensing flexibility	Does the vendor provide flexible licensing options to support changes in infrastructure, for example the migration of on-premises IPs to cloud workloads?
Core capabilities	Which capabilities are included as part of the vendor's core product? Which features are priced separately or require additional licensing?

Scope and deployment

Aspect	Description	Key points to consider
Out-of-band monitoring	Passively monitor network traffic to avoid any impact to the production environment	▪ Port mirroring via SPAN session or TAP, typically deployed into a core switch ▪ Analysis of raw network traffic (packets) ▪ Passive network monitoring instead of in-line deployment ▪ Automatically identifies all network assets without relying on users or active scanning ▪ Solution should be agentless by default ▪ Quality of data ingested, e.g. is it purely source and destination IP? Typically, the richer the data the better the analysis
Integrations	Connecting third-party technologies to enhance detection, response and investigation workflows.	▪ Not reliant on, but can integrate with, tools such as: ▪ Firewalls, VPNs, SASE and ZTNA solutions ▪ EDR and XDR solutions ▪ SIEM and SOAR platforms ▪ Microsoft 365 ▪ Cloud services such as Microsoft Azure and AWS ▪ IAM platforms including Microsoft Entra ID, Okta and Duo ▪ SaaS applications such as Dropbox, Salesforce, Slack and Zoom ▪ Vulnerability management tools such as Rapid7 and Tenable ▪ Workflow and ticketing solutions including ServiceNow and Jira ▪ Enrichment from threat intelligence feeds and known CVEs
Data retention	Retention of raw network packets and metadata for analysis and investigation	▪ Ability to retain: ▪ PCAPs ▪ Flows and key metadata ▪ Log data for deeper investigation
Air-gapped and highly secure / segregated networks	Does not lose core functionality in air-gapped environments	▪ Option to operate entirely independently with no outbound connections ▪ Can hold all data on-site within a supplied appliance ▪ Not reliant on any external connectivity for: ▪ Detection and response ▪ Incident investigations ▪ Continuous AI model updates ▪ Third party threat intelligence ▪ PCAP storage and search ▪ Creation of custom detection and response models ▪ Custom configurations and edits to system and detection and response models

Aspect	Description	Key points to consider
Cloud environments, hybrid cloud, IaaS, PaaS, SaaS	Coverage and visibility of fully cloud-based and hybrid cloud environments	▪ Natively extends network coverage to cloud environments and can view both through a single platform ▪ Ability to deploy in fully cloud-native and hybrid environments, and during transition from on-prem to cloud ▪ Ingestion of flow logs (e.g., AWS VPC Flow Logs, Azure NSG Flow Logs) ▪ Agentless by default, or leverage VPC mirroring or virtual taps ▪ Dynamically maps Azure architecture, including ephemeral resources like containers and serverless applications ▪ Deployment options for virtual machines and in container orchestration environments ▪ Integrates with services like AWS Security Hub, IAM, CloudTrail, and Azure equivalent to understand identity, control plane and network context
OT/CPS and IoT devices	Coverage and visibility of Operational Technology/Cyber Physical Systems and Internet of Things devices (IoT, IIoT, mIoT, BMS, SCADA, HMIs, PLCs)	▪ Natively covers converged and segregated IT/OT systems, IoT, IIoT, and remote edge devices in a single platform ▪ Can be deployed in fully air-gapped environments ▪ Can perform deep packet inspection across OT-specific protocols ▪ Active discovery for OT asset inventory and vulnerability insights ▪ Coverage of all levels of the Purdue Model, from Level 0 physical processes to Level 5 enterprise networks ▪ Solution is vendor agnostic ▪ Performs continuous risk assessment and vulnerability correlation by analyzing real-time asset behavior, network exposure, and exploitability
Remote worker endpoints	Full network visibility for remote worker devices	▪ Full network-level visibility of remote worker endpoints and travelling devices, in addition to EDR alerts ▪ Coverage of small satellite offices or shared workspaces with no/limited network infrastructure

Network behavioral analysis

Aspect	Description	Key points to consider
Artificial intelligence and machine learning	Uses multiple AI techniques for threat detection, investigation and response	- Multi-layered approach to AI that is layered sequentially and hierarchically, including: - Unsupervised machine learning - Supervised machine learning - Bayesian learning - Clustering algorithms - Ensemble methods - Graph Neural Networks (GNNs) - Natural Language Processing (NLP) - Domain-Specific Language Models (DSLMs) - Security-specific custom Large Language Models (LLMs)
Self-Learning AI	AI that learns the unique environment it is deployed into, adapting and learning by itself	- Builds an understanding of what is 'normal' for each unique environment without relying on pre-defined rules and signatures - Detection of anomalous activity based on what is considered normal behavior, rather than historical attack data - Drastically reduces ongoing manual detection engineering compared to traditional tools - Can automatically adapt to changes in the environment without requiring manual updates or configuration - Automatically sorts devices into groups and clusters by their behavioral similarity - Solution is completely configurable and allows for fine-tuning where desired
Deployment of AI	Location of AI for ongoing learning and processing of data	AI is deployed locally, and learns based on each unique environment it is deployed into
Compliance	Detection of compliance related activity	- Ability to detect compliance policy violations such as unauthorized usage of GenAI tools - Ability to create edit custom compliance models

Aspect	Description	Key points to consider
Encrypted traffic analysis	Can identify anomalous and suspicious activity in encrypted network traffic	▪ Proven ability to detect both known and novel threats in encrypted traffic without requiring decryption
Decryption	Decryption of encrypted network traffic	▪ Ability to decrypt network traffic if desired, e.g. for compliance purposes ▪ Native and third- party decryption options ▪ Not reliant on decryption to detect anomalous and malicious behavior
Inbound and lateral network traffic analysis	Analyzes both North-South and East-West network traffic	▪ Does the vendor analyze incoming network traffic for various threats to prevent security breaches ▪ Does the vendor analyze internal network traffic to detect and respond to threats within the organization such as lateral movement and insider threats
Retention	Storage of network traffic, events and other associated metadata Retains raw network traffic and metadata for retrospective analysis and to ensure compliance with legal requirements	▪ Retains raw network traffic and metadata for retrospective analysis and to ensure compliance with legal requirements
Asset visibility	Identification and management of network entities	▪ Automatically discovers and maps network assets in a single location ▪ Ability to view detailed asset information and network activity logs for each device ▪ Automatic and manual tagging of devices for easy identification, search and management

Threat detection

Aspect	Description	Key points to consider
Detection models	AI models to detect suspicious and malicious activity	▪ Wide variety of stock detection models to cover both security and compliance use cases ▪ Ability to view, edit and customize all models directly in the UI, without requiring additional development, vendor support or external connectivity ▪ Ability to create entirely custom models to support specific use cases such as compliance or edge cases outside of the standard model deck
Novel threat detection	Uses AI-driven anomaly detection to spot emerging or zero-day attacks with no prior indicators.	▪ Proven protection against threats such as: ▪ Novel ransomware and new variants ▪ Zero-day vulnerabilities ▪ Insider threats ▪ Misuse of legitimate tools ▪ Not reliant on: ▪ Threat intelligence feeds ▪ Comparing data across other customer environments ▪ Rules or signatures ▪ CVE ingestion ▪ Vendor should provide at least 3 documented and detailed examples of the technology detecting zero-day threats before they are publicly disclosed, without relying on rules, signatures, training data, or any prior assumptions
Known threat detection	Matches anomalous activity indicative of known threats	▪ Detection of known threats such as ransomware and malware ▪ Can ingest and use threat intelligence feeds if desired
Existing compromises	Detection of pre-existing threats and compromises in an environment	▪ Identification of unusual behavior in relation to similar devices and peer groups ▪ Identification of suspicious activity indicative of existing compromise
Account takeover & lateral movement	Detects compromised accounts moving across systems to escalate access and control.	▪ Does the product analyze internal network traffic for anomalies? ▪ Can it correlate network activity with compromised credentials or session hijacking? ▪ Analysis of inbound, outbound, and lateral network flow ▪ Behavioral modeling to identify anomalous communication patterns ▪ Simulations of high-fidelity events ▪ Cross-domain correlation across network, endpoints, identity, SaaS, etc.
C2 communications	Identifies malicious command-and-control traffic, even when hidden in DNS, HTTPS, or other common protocols.	Can the vendor detect: ▪ Beaconing patterns (regular interval callbacks) ▪ Use of uncommon protocols (e.g., DNS tunneling, ICMP-based C2) ▪ Connections to known malicious IPs/domains ▪ Encrypted C2 over non-standard ports or unexpected services

Aspect	Description	Key points to consider
Malware and payload delivery	Detects the transfer of malicious files or code across the network before execution.	- Can the vendor detect signs of: - Malicious file transfers (via SMB, FTP, HTTP) - Suspicious executable transfers or dropped payloads - Delivery via uncommon or deprecated protocols (e.g., TFTP, Telnet)
Insider threat behavior / legitimate credential abuse	Flags unusual activity by authorized users that could indicate malicious intent or negligence. Detects misuse of valid credentials, whether stolen or willingly shared	Can the vendor detect signs of: - Privilege abuse - Unusual data access or transfer - Abnormal authentication or session behavior - Data transfers/exfiltration to unknown or untrusted external hosts - Large outbound transfers to rare or new external destinations - Data sent over covert channels (DNS, HTTPS to unknown domains) - Unusual protocol usage (e.g., FTP/SFTP from endpoints that don't normally use it) - Timing anomalies to evade DLP (e.g., trickling)
CVE exploitation	Taking advantage of known vulnerabilities in systems, services, or applications.	- Detection of anomalous activity across network entities, such as: - Unusual login activity - Anomalous data transfers - Connections to rare endpoints - Unusual internal network scanning - SMB reconnaissance - Lateral movement - Privilege escalation
Suspicious scanning & reconnaissance	Identifies probing and mapping of network assets that often precede an attack.	Detection of internal network scanning and reconnaissance techniques
Data exfiltration detection	Alerts on unauthorized movement of sensitive data to external or cloud destinations, even if encrypted.	Detection of connections to rare endpoints Detection of connections and uploads to file sharing sites such as WeTransfer, Dropbox, etc. Detection of large internal downloads or anomalous connections to internal locations

Triage and investigation

Aspect	Description	Key points to consider
Automated triage	Elimination of manual triage and investigation work that commonly leads to alert fatigue	- Automated triage and investigation of network events to surface only the most anomalous and interesting activity - Technology learns by itself to continually update its understanding of normal - Does not require human intervention, prompting or scripting
Alert prioritization and management	Prioritizes alerts based on criticality	- Automatically prioritizes alerts to surface the most critical activity for human analysts - Categorizes alerts from critical to informational and compliance - Can assign alerts to specific users - Can view and manage alerts via a mobile app, as well as via a browser - Can connect to third-party solutions e.g. ServiceNow, Jira
Alert correlation	Correlation of multiple related alerts into single incidents	- Correlation of alerts from within the vendor's own platform, such as NDR, OT and email - Correlation of alerts from third-party technologies including: - EDR/XDR - SASE and ZTNA - SaaS applications - Grouping of multiple, related alerts to surface genuine incidents, without human interaction
AI-led investigations	End-to-end investigations performed autonomously by an agentic AI system that is not GenAI -based	- Automatically combines relevant alerts into incidents, either existing or creates a new one, and escalates critical cases for human review. - Agentic AI performs investigations to the quality of a Level 2 SOC analyst - No reliance on prompt-based chatbots, GenAI or human intervention - Investigations can be triggered manually or automatically by external sources - Full reports are generated for each incident, showing a timeline and clear explanation of the investigation process completed by the AI - Available to all customers, not just specific licensing tiers or as part of a managed service
Threat hunting and advanced analysis	Additional capabilities to support threat hunting, deeper investigations, forensic analysis and incident response.	- Ability to download and analyze PCAPs manually if desired - Advanced searching and analysis of network traffic using simple queries, with the ability to create and save custom queries - Ability for advanced users to use complex query syntax and build complex structured searches - Advanced search functionality is included as part of the core product capabilities and does not require an additional license, incur an additional cost or require external connectivity

Containment and response

Aspect	Description	Key points to consider
Autonomous response	Autonomously takes the best course of action to respond to threats	▪ Functions 24/7, autonomously and at machine speed ▪ Can choose and apply the most effective response action itself without human interaction ▪ Can autonomously escalate the severity of response actions as a threat evolves ▪ Can be configured to require human confirmation before taking a response action if desired ▪ Can be configured by subnet, times, days, etc. ▪ Can fully customize response models and parameters in detail ▪ Does not rely on static rules, signatures ▪ Has proven examples of containing novel, known and insider threats at the earliest stages
Native response	Ability for the vendor to deliver response actions natively without relying on any third-party technologies	▪ Can natively respond to network threats without an EDR or firewall integration ▪ Can deliver targeted actions to block specific traffic without causing business disruption ▪ Can enforce only what is considered normal activity for a network entity, while blocking everything else ▪ Can also be applied based on what is considered normal for a peer group of devices ▪ Can operate in flat networks with little to no segmentation
Third-party response	Ability to integrate with third-party technologies to enhance native response capabilities	▪ Can integrate with third-party solutions to extend response action capabilities including: ▪ EDR/XDR solutions ▪ Firewalls ▪ Microsoft 365 and Entra ID ▪ SIEM and SOAR
Threat containment	Stopping threats from progressing at the earliest stages	▪ Can act at the earliest signs of threatening activity, rather than waiting for an EDR alert or manual investigation to take place ▪ Contains threatening or suspicious behavior without disrupting normal business activity ▪ Containment duration can be extended where required to give analysts time to react ▪ Can be used to autonomously quarantine any new network entities that are not recognized

Additional functionality

Aspect	Description	Key points to consider
Reporting	Enhances NDR value by turning detection data into clear, prioritized reports that support security operations and cyber compliance	Can generate specific reports for compliance adherence over time, such as PCI-DSS, GDPR, HIPAA, ISO 27001, NIST, SOC2, MITRE ATT&CK, IEC 62443
Continuous threat and exposure management (CTEM)	Complements NDR by proactively identifying and reducing vulnerabilities, plus displaying attack paths that attackers could exploit in the network	- Generates unique, specific risk scoring for your specific environment, rather than general claims about patch latency or CVE risk - Prioritized view for cyber risk discovery to address the users, devices, and vulnerabilities which pose the most severe compromise risk - Evaluates the potential impact to affected assets during incidents and enables more effective response along the most critical attack paths
Incident readiness and recovery	Supports testing, refinement and validation of incident response capabilities	- Does the vendor provide simulations that reflect your specific architecture, configurations, and tech stack? - Can the solution simulate live, realistic incidents across our actual environment — including cloud, SaaS, and OT infrastructure — not just tabletop scenarios? - Does the system support adaptive or dynamic playbooks that evolve with our environment (e.g., post-cloud migration or M&A activity)? - What are the licensing, resource, and skill dependencies for setup and ongoing management?
Attack surface management	Informs NDR by mapping external -facing assets and exposures, helping prioritize detections tied to high-value or high-risk systems.	- Does the solution use AI techniques to automatically discover unknown or shadow assets without needing IP ranges or seed data? - Can it detect assets associated with your brand across third-party domains, cloud instances, and IoT infrastructure? - Does it continuously crawl and monitor your digital footprint to identify new or changing exposures? - Is it capable of detecting zero-day or high-impact vulnerabilities without relying on periodic scanning? - Can it detect unauthorized or unmonitored domains, services, and assets created outside central IT governance? - Does it account for supply chain and brand abuse risks, including phishing domains or spoofed infrastructure? - Does the ASM solution enhance other NDR capabilities (e.g., spoofed domain defense, endpoint context, AI Analyst correlation)? - Can it provide a unified view when paired with internal telemetry and other risk management tools?
Email security	Integration of email security with NDR	- Ability to automatically correlate cross-domain threats from the inbox with subsequent network activity - Are the tools working together to inform more accurate threat detections and faster investigations? - Provides a defense-in-depth approach from the inbox to the network

How to evaluate AI


AI is being positioned as a silver bullet in cybersecurity, but most security teams are left wondering what's actually under the hood. Many tools claim to use AI yet often rely on just supervised machine learning or generative AI models, which are only effective in narrow, well-defined scenarios and lack the ability to detect novel threats.

When evaluating NDR solutions that use AI, it's essential not to rely solely on the output of a single model or one type of AI (such as generative-AI based LLMs). This will only deliver limited outputs and may generate incomplete or biased results.

This challenge is further compounded by the rise of AI-powered threats, which demand more adaptive and intelligent defenses. Unfortunately, without a clear way to evaluate these capabilities, organizations risk falling behind, or worse, being caught off guard.

The truth is, not all AI is created equal. To be effective, AI in cybersecurity needs to go beyond buzzwords and integrate multiple advanced techniques that work together to address the complexity of modern threats. And security professionals need to understand what types of AI are being used for what use cases to in order to evaluate the AI's effectiveness against their organizational needs and build a stronger security posture.

Attack-centric



Business-centric




This diagram shows a comparison between AI-based security tools that leverage business data to detect anomalies in behavior and tools that are trained on historical attack data to catch threats.

To help you we have created three separate guides to AI in cybersecurity that will help you in every step of your adoption journey:

Deep dive into the application of AI in cybersecurity

Download the AI Arsenal Whitepaper to explore how AI models can be applied to cybersecurity, and how Darktrace's Self-Learning AI layers multiple techniques to deliver proactive, resilient threat defense.


[Learn more ↗](#)



Evaluate AI-based security solutions

Download the CISO's Guide to Buying AI Whitepaper. This white paper looks at how buyers should approach purchasing AI-based solutions. It outlines key considerations for each stage of the AI adoption journey, specific questions to ask vendors, and what to look for in the responses.


[Learn more ↗](#)



Evaluate how advanced your organization is in their AI maturity journey

Download the AI Maturity Model Whitepaper. This paper gives a full deep dive into Darktrace's AI Maturity Model – the only available framework capable of guiding CISOs and cybersecurity department leaders to understand, evaluate, strategize, and advance their progress in AI adoption.

[Learn more ↗](#)



The Evaluation Process

■ Step 01:

Commercial evaluation

The first phase is a commercial evaluation or discovery process.

This involves assessing whether a potential solution aligns with your organization's needs – not just in terms of core capabilities, but also in how it fits your workflows, operational model, and overall value expectations.

Here are our recommendations for what to consider during the process:



Engage with vendors

Maintain open lines of communication with your vendors throughout the evaluation process. Work with them to identify and understand any performance discrepancies. This collaborative approach allows you to test the team's responsiveness while addressing any issues promptly and ensures that you are making an informed decision based on accurate data.



Factor in people and partnerships

Consider the human elements that can significantly impact your experience with the solution. Evaluate the local presence of the vendor and the depth of their account management team, including any relationship with your existing Managed Security Services Provider (MSSP) if relevant. Strong partnerships and accessible support can be crucial for long-term success and satisfaction.



Compare similar licensing

Ensure that you are comparing solutions with similar licensing packages. Some vendors offer different packages that may not provide equal capabilities. Be aware of which features are included to ensure a fair comparison, as well as any added-value capabilities.



Think about post-purchase support

Ensure the vendor offers robust post-purchase support, including regular updates, training and certifications, events, and responsive customer service. Read reviews from other customers and industry analysts such as Gartner to determine if customers recommend and see value from the vendor over time, not just during the trial and onboarding phases.

Technical evaluation

The second key step is the technical evaluation, where it's essential to involve the day-to-day users of the product.

This discovery process helps determine how well a solution meets your organization's performance, usability, and integration requirements. It also allows you to assess critical aspects such as deployment architecture, detection efficacy, use of AI, policy enforcement, and compatibility with your existing infrastructure.

Here are the key factors to consider when performing a technical evaluation:



Evaluate based on optimal performance data

Ensure you assess your chosen solution once it has been fully deployed. This is especially true with solutions that use unsupervised machine learning and therefore have an initial learning curve. Engage with your vendor to confirm this and understand the true potential of each solution under ideal conditions, including the deployment of any integrations that you require.



Evaluate the learning curve

Assess not only the detection capabilities at a single point in time but also the rate of improvement and maturity over time. This will give you a more accurate picture of the solution's long-term precision.



AI models and fine-tuning capabilities

Ensure the vendor provides complete transparency in their AI modeling and capabilities. You should be able to customize the detection and response models to make any finer adjustments to suit your needs, plus the ability to create custom models for specific use cases that are not covered in the vendor's default model library.



Clear and transparent explanations

Every detection, investigation and response action should come with a transparent explanation. Understand what detection models were triggered, the corresponding logic behind any automated investigation and a detailed list of response actions taken. This transparency is crucial for trust and effective evaluation.



Combine qualitative and quantitative analysis

Go beyond AI marketing claims by combining both qualitative and quantitative analyses of the solution's capabilities in your specific environment. Assess elements such as visibility, detection accuracy and overall usability, plus quantifiable metrics that you can compare against your existing workflows such as triage and investigation time, total incidents generated, and mean time to respond.



Compare threats consistently

Different vendors may classify threats in various ways, which can lead to inconsistencies in your evaluation. Ensure that you are comparing equivalent metrics by seeking clarification from your vendor when necessary. This will help you make a fair and accurate comparison.



Do not evaluate all false positives equally

Understand that a 'false positive' may be defined differently across the solutions you're evaluating. For example, business-centric tools that use anomaly-based detection will raise alerts for activity that is considered anomalous. They are legitimate detections, but anomalous does not always equal 'malicious' and therefore does not meet the traditional definition of a false positive. Solutions that use AI for anomaly-detection will operate just like a skilled human analyst – flagging suspicious activity based on the data available, then using that as context for further investigation and response. Take the time to understand the context and impact of false positives within each solution to make a more informed evaluation, and if you're not sure, ask your vendor how they define false positives. Test how easy it is to acknowledge alerts that are not considered interesting or of genuine concern to your team.

Conclusion

We hope this checklist provides a practical framework to help you evaluate multiple NDR solutions and identify the one that best aligns with your organization's needs.

Now you have the framework in hand, why not see Darktrace / NETWORK in action and start a free 30-day trial to test our technology your own environment.

[Get a demo ↗](#)

■ **About Darktrace**

Darktrace is a global leader in AI cybersecurity that keeps organizations ahead of the changing threat landscape every day. Founded in 2013 in Cambridge, UK, Darktrace provides the essential cybersecurity platform to protect organizations from unknown threats using AI that learns from each business in real-time. Darktrace's platform and services are supported by 2,700+ employees who protect nearly 10,000 customers globally. To learn more, visit www.darktrace.com.