THREAT WARRIOR

Key Considerations for Selecting an NDR Solution

Technical FAQ

ThreatWarrior Network Detection and Response (ThreatWarrior NDR) is a powerful, extensible solution that leverages cutting-edge unsupervised neural networks and deep learning technology to identify anomalous behavior and potential threats without relying on pre-set rules or labeled data. Our neural networks use meta representations to analyze the entire packet stream and provide a comprehensive view of network traffic, enabling the solution to detect sophisticated threats that signature-based and other legacy technologies may miss. ThreatWarrior NDR also has custom feature hierarchies and real-time learning capabilities that allow it to adapt quickly to changing network environments and identify new threats. Furthermore, the solution's training granularity enables it to focus on specific areas of the network, providing detailed insights into potential threats and reducing false positives.

While we are proud of our solution, we understand that selecting an NDR product that aligns with your organization's unique requirements can be challenging. As such, we have compiled a list of essential questions to ask any vendor before making a purchase decision. These questions will help you understand the critical features and capabilities of a robust NDR solution, such as the ability to collect and analyze data from diverse sources, seamless deployment, advanced threat detection and analysis, intuitive visualization, intelligent automation, and the flexibility to integrate with other security tools. We believe that asking these questions will enable you to make an informed decision and select an NDR solution that will provide comprehensive protection against potential network threats.

Questions to Ask

What data sources are used?

A strong NDR solution needs network traffic data and security data to identify potential threats and vulnerabilities. Collectively, these data sources are critical to building a comprehensive NDR solution that can effectively detect and respond to potential cyber threats:

- **Network traffic data:** Network traffic data (including source and destination of network traffic, the type of traffic, and the volume of traffic) is used to detect anomalies and potential threats.
- **Network device data:** Data on the devices connected to the network – including information on the operating system, applications, and configuration settings – helps to identify potential vulnerabilities that could be exploited by attackers.
- **Network behavior data:** Machine learning and behavioral analysis can identify abnormal patterns of behavior on the network, e.g., deviations from normal behavior, such as unusual communication between devices or unauthorized access attempts.
- **Security data:** Threat intelligence feeds and vulnerability data help identify potential threats and vulnerabilities on the network.

> *ThreatWarrior NDR ingests full packet capture, network flow data and a number of threat intelligence feeds into its suite of analysis engines. Customer/network segment-specific behavioral data is learned, ensuring abnormal behavior generates alerts. Behavior monitoring continues to evolve and improve over time as it learns more about the specific nature of your business and how to protect it.*

How is data retention managed?

Some NDR solutions only track metrics for 30 seconds, five minutes or an hour. NAS is supported for expansion to 24 hours. Short window metrics (only) storage is insufficient for an NDR solution to perform its duty.

Other vendors claim their sensors can handle up to 100 Gbps connections. In and of itself, given modern NICs, this is a near commodity. Processing and storing that volume of data, however, is another discussion. A connection that size, under full load, would transmit no fewer than 8 million packets per second. Ignoring RAID, and assuming an IMIX average packet size of 353 bytes/packet, that equates to 244 TB/day.

> ***ThreatWarrior advocates storage of a rolling 30 days of alerts for analysis purposes, where packet payloads are stored on sensors and made available through the web UI on request. Longer data retention windows can be managed with offline, less expensive storage options.***

How important is encrypted traffic inspection?

Some NDR vendors advocate for a man-in-the-middle encrypted traffic decryptor.

> ***ThreatWarrior NDR does not rely upon decryption to judge the intent of a traffic flow. In fact most customers do not want those detailed packet information reviewed or retained outside of their controls. ThreatWarrior never takes custody of customer information.***

Is there actual network monitoring or just acceptance of data feeds?

Some NDR solutions accept and aggregate data feeds but do not do any monitoring of their own on endpoints or the network. These solutions often have only minimal machine learning, which runs in a batch mode every n hours (n varies by vendor) and rely upon predefined models to detect certain kinds of abnormal behavior and known indicators of a breach.

> ***ThreatWarrior NDR monitors the entire network in real time. Whether malicious or benign, everything has to cross a network to achieve anything. ThreatWarrior NDR observes traffic at the packet level as it flows across the network. Packet traffic is ground truth for what's happening on your network. The solution provides visibility into all traffic - north-south and east-west - to deliver a full picture of network activity.***

Are agents required?

Some NDR solutions require agents, while others do not. Agents are small software programs that must be installed on individual devices or endpoints within a network to collect data and perform analysis.

> ***ThreatWarrior NDR doesn't require that you install any agents on the network. It passively connects by network or virtual tap to provide insight into all activity in your organization without having to download software onto everything you want to protect. This is also important as it means the solution is invisible on the network, so bad actors can't tell you're using NDR and evade it or disable it - as they can with EDR.***

Deployment

Can a solution architecture be based solely on virtual machines (VMs)?

Some NDR solutions claim to discover managed, unmanaged, IoT devices (both on- and off-air) without requiring hardware devices or software agents. For this to be true, the solution must be run as a VM. But, if it's a VM, a user will need to be responsible for feeding it data from their own network infrastructure. This implies numerous challenges, not the least of which is that most organizations do not have the networking capacity to feed their entire network traffic to a single device on their network. Even if they did, they would face severe architectural challenges getting the data to it on a network already burdened with a significant amount of traffic. So in reality, VM-only based solutions are really only capable of operating on metadata (flows, etc.), and would therefore be incapable of deep or even shallow packet inspection.

> ***For physical offices, passive sensors are deployed at the most logical points on your network. These sensors will see all traffic flows which shows us everything communicating on your network, including any IoT devices that may be 'hiding'. Sensors can also be deployed at co-lo or data center locations - providing an additional view into everything traversing your extended network. Even remote or global operations - including things like ATMs, point of sales, industrial control systems, etc. can be seen. Finally, we provide full visibility into the public cloud, allowing you to see things like microservices, virtual functions, virtual machines, etc.***

What is the scope of deployment?

Some NDR solutions are cloud-only. Others are marketed as 'SaaS only'. In either of these cases, the implication is that no sensor/ appliance is required. If only net flow data is being captured, all deep packet inspection information will be missed, leaving a shallow level of detection capability.

Others lack a packet mirroring solution for the cloud, requiring instead that you route all external traffic through their sensor, using it as a proxy gateway. This means they do not track traffic between individual VMs within your network.

> ***ThreatWarrior NDR is a true hybrid-cloud solution. Our solution has the ability to see and protect anything connected to the network - on-premise machines, IoT devices, cloud environments, industrial control systems and more. Further, ThreatWarrior NDR can be deployed on site behind the user's firewall, providing maximum security as data is stored and analyzed within the user's network.***

Detection / Analysis

Is pre-training used?

Some NDR vendors claim the use of deep neural networks where anomaly detection works 'out of the box' and requires no baselining. The claim is that this is completely automated with pre-trained AI models, requiring no human triage and no learning delay. In essence this means their AI was trained on someone else's network. By definition this introduces bias in training data - the biggest problem in machine learning. Do you want a solution coming in with preconceptions of what it should expect, or misconceptions that must be slowly corrected over time as it re-learns? And buyers should always remember that pre-training had to come from someone else's network. Is your network traffic being harvested to expand the vendor's commercial model?

> ***ThreatWarrior NDR excels at feature extraction. Feature extraction is the process of building complex hierarchies of meaning to express information from raw data. Apply this to cybersecurity, and rich information can be derived from raw traffic, e.g.,, "who talked to whom about what" to conceptualize higher-order patterns in the environment. Using unsupervised neural networks to perform deep learning allows you to observe significantly more detail, so what you see is a better, more accurate picture of your security environment. Antiquated solutions can require manual work for programmers to codify examples of what's normal into their platforms, taking up valuable time and resources. ThreatWarrior NDR does this with no supervision, no feature engineering, and no pre-training.***

How is DPI actually used?

Some NDR solutions use a very minimal level of DPI, e.g., inspecting simply on security certs and DNS. No threat feeds are utilized, nor is any AI applied. This is an example of simply capturing NetFlow data, and then leaving it to an analyst to cold review raw data with little alert enrichment.

> ***ThreatWarrior NDR uses continuous deep packet inspection - far different than most cybersecurity solutions - in any security category - that are only doing traditional packet capture. DPI evaluates the packet header and contents in order to identify malware and classify traffic by application protocols. This allows us to see any policy violations in real time and identify more than 30,000 known malware signatures.***
>
> ***Unlike traditional packet capture and plain packet filtering, continuous DPI examines more than just packet headers. It allows you to dissect and analyze the packet payload and all network data, identifying granular information from end to end. Continuous deep packet inspection enhances network security by helping you understand which application packets belong to, and by revealing network patterns and user behavior — and it provides this information 24/7 in real time.***
>
> ***Without utilizing continuous deep packet inspection, cybersecurity solutions can't constantly extract or filter any information beyond packet headers without costly full packet capture.***

Is the ML supervised, semi-unsupervised, or fully-unsupervised?

NDR solutions built strictly upon supervised machine learning require human assistance to classify anomalies as known malicious activity. What if new anomalous activity is malicious but doesn't match a known pattern? Supervision-based approaches simply cannot keep pace with malware mutations.

Some NDR solutions rely upon a semi-unsupervised machine learning system, where supervised models are trained based on analyst feedback driven by what the system surfaces as threats.

> ***ThreatWarrior NDR uses a primarily unsupervised system, with supervision provided by analyst feedback on threat validity, which feeds into data curation. The feedback loop defines if the alert is good or bad. As in, "Don't tell me about this again," or, "Good job!"***

Is Bayesian machine learning used?

Bayesian machine learning is an unsupervised approach without the use of deep neural networks. This means the solution relies upon engineers to anticipate precisely which aspects of the cyber environment should be tracked and measured to defend against a constantly changing threat surface.

> ***This is a far less capable, outdated technique that lacks the power, sophistication and results achievable through a series of deep neural layers - as employed by ThreatWarrior NDR - which constantly recalculate the most significant aspects of the environment to track on each network it monitors.***

Is detection based on a single engine?

Single-engine NDR approaches - systems that operate only on DPI, behavioral, rule and policy inspection, etc. - are likely to generate high false positive counts. There simply is not enough context, correlation and cross-checking to make a strong determination of intent. As a simple example, simply examining a packet flow for protocol across a well-known port is far less conclusively severe than examining port, protocol, and the packet payload as well.

> ***ThreatWarrior NDR uses multiple engines (DPI, Behavioral, Integration, Cloud Entity, Rules & Policies, and Insights) to make determinations on benign, suspicious, malicious findings - and the ultimate severity of high-alert conditions. As one example, the Insights engine provides long-term behavioral profiling of networks and the devices that live on them. By tying together raw network traffic, deep packet inspection results, parsed protocol data, known threats, and AI engine results, ThreatWarrior NDR offers full-context situational awareness and predictive analytics to keep security teams ahead of threats. The system also features coherent change detection, so it alerts to emerging patterns and slow changes made over time.***

How is DPI actually used?

Some NDR solutions rely upon feature engineering - where engineers pick and choose what data to feed to an AI algorithm. Instead of feeding raw data into a neural network, analysts must predetermine very specific features within the data set that are important.

Examples of feature engineering include:

- Number of concurrent flows
- Number of packets per second
- Flow duration
- Lots of bytes going over DNS
- Lots of data going off the LAN
- Number of requests in a minute

The rest of the data is discarded. This implies the vendor has 'experts' who know how to pick the right features to input - a dubious claim at best.

> ***ThreatWarrior NDR's deep learning approach ensures the solution itself determines the parameters it should care about. This is a superior learning approach given that human-defined parameters for one network could be entirely erroneous for the next.***
>
> ***By way of example, deep learning systems are known to excel at classifying pictures of anything without the need of 'domain experts' hand-crafting filters to pick out the image features. Prior to deep learning, a set of expert-crafted filters were required to pick out noses and eyes for facial recognition. Another set of filters would be needed to pick out wheels and grills for car detection. This presents clear scale issues. With deep learning, the same analysis pipeline is leveraged for both types of images, and the deep network builds the filters that matter for the kind of images fed to it.***
>
> ***The very same principle applies for network data. One network may need more filters for IP address subnets, while for the next network, protocols are more important. Unsupervised neural networks and deep learning will build what they need to meet an objective.***

Is the machine learning really just UEBA-based?

Some NDR solutions are effectively just extensions of User and Entity Behavior Analytics (UEBA) - an approach that involves using machine learning and other analytical techniques to identify unusual patterns of behavior among users and entities, e.g., devices or applications, on a network. By monitoring a wide range of activity data, such as login attempts, file access, and network traffic, UEBA purports to identify potential security threats that might otherwise go unnoticed.

This technology is quite limited, especially with respect to modern cyberattacks. Sophisticated attacks do not target a single machine. They execute a host of small, anomalous events across many devices. To catch this type of event, organizations need a holistic view of network activity, not a collection of independent behavior profiles. The later can lead to:

- Heavy false-positive alerting, which wastes precious analyst time and resources
- Limited visibility causing it to miss important patterns of behavior
- Weak threat identification owing to lack of context. While UEBA can identify unusual behavior, it often cannot explain why the behavior is occurring. Without additional context, it can be difficult to determine whether a behavior is actually a threat.
- Constant and expense rule updates - UEBA systems are designed to identify specific types of threats based on predetermined rules and models. If a new threat type emerges, the system may not be able to adapt quickly enough to detect it.

> ***As covered previously, ThreatWarrior NDR uses multiple engines (DPI, Behavioral, Integration, Cloud Entity, Rules & Policies, and Insights) to make determinations on benign, suspicious, malicious findings - and the ultimate severity of high-alert conditions. The Insights engine provides long-term behavioral profiling of networks and the devices that live on them - but is but one engine performing analysis. By tying together analyses from multiple engines, ThreatWarrior NDR offers full-context situational awareness and predictive analytics to keep security teams ahead of threats - and avoids the limitations of simple UEBA.***

Is the AI really just a rule-based algorithm?

Rule-based NDR solutions use pre-configured rules to identify known threats on a network. Rules are based on specific patterns of behavior, or signatures associated with known malware or attacks. Rule-based NDR solutions do not leverage unsupervised neural networking or deep learning, and will have the following weaknesses:

- Simple statistical models only work with data feeds provided, e.g., authentication and DB logs. There is no actual network monitoring.
- Only able to detect known threats; cannot identify new or emerging threats that do not match pre-configured rules
- Generate a large number of alerts, many of which will be false positives, resulting in alert fatigue and wasted analyst time
- With limited context about the detected threat, security personnel must investigate each alert to determine whether it is a real threat or a false positive
- Rule-based systems require constant rule updates in order to keep up-to-date with the latest threats

> ***ThreatWarrior NDR uses true AI and unsupervised machine learning to protect against complex cyberattacks.***

If deep learning is claimed, is it supervised or unsupervised?

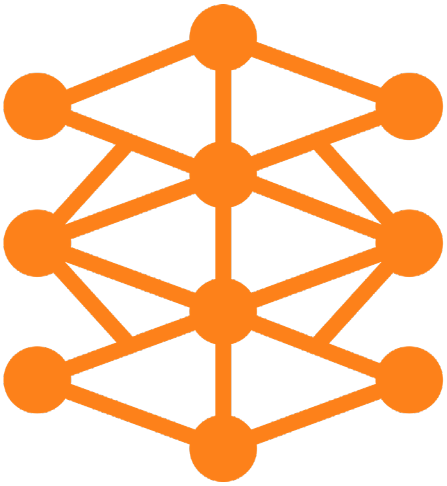

Supervised learning of any kind requires that the system be presented with a large number of legitimate and illegitimate traffic examples in order to classify one from the other with accuracy. As a result, supervised learning is a long, arduous and expensive process. To make matters worse, supervision-generated data sets are not customized to specific client networks. Given the highly-nuanced world of network designs and security policies, assuming one network's normal is the same, or even close, to the next is dangerous.

> ***ThreatWarrior NDR uses unsupervised neural networking which allows it to build complex hierarchies of meaning to express information from raw data. This allows analysts to observe significantly more detail and a more accurate picture of the network environment. With ThreatWarrior NDR, no supervision, no feature engineering, and no training period is required - saving valuable time and resources.***

Is the ML based on a set of signature-less techniques?

Some NDR solutions are marketed as leveraging unsupervised and supervised machine learning, but in reality are built upon signature-less techniques like event rarity and min-max clustering algorithms. These data analysis techniques fall short of true deep learning in several ways:

- Struggle to scale up and handle larger datasets since they rely on the identification of patterns and anomalies in the data. As the size of the data increases, the process of identifying these patterns can become more complex, resulting in slower processing times and reduced accuracy.
- Difficulty generalizing new and unknown threats. They are often designed to detect specific patterns or anomalies and may not be effective in detecting more complex attacks that have not been previously identified.
- Signature-less techniques are easier to evade. Attackers can modify their behavior to avoid triggering simplistic detection algorithms.
- These techniques struggle to contextualize the data they are analyzing, which can result in false positives or missed detections. Without a deeper understanding of the underlying patterns and relationships in the data, accurately distinguishing between normal and anomalous behavior becomes error-prone.

> ***ThreatWarrior NDR uses deep neural networks for its unsupervised machine learning. It looks at the network as a whole, instead of independent and isolated devices. The advantages of the more holistic approach include:***
>
> - ***Increased accuracy: more comprehensive and nuanced relationships between devices and users within the network, which can lead to higher accuracy in detecting anomalous behavior. By considering the network as a whole, the deep neural network can identify patterns of behavior that may be missed by traditional approaches that only focus on individual devices.***
> - ***Better context: Better context around which devices are operating leads to more accurate identification of anomalous behavior. Deep neural networking takes into account factors such as time of day, location, and user behavior, all of which can impact how devices communicate with each other.***
> - ***Improved scalability: Deep neural networking can handle larger and more complex networks with many interconnected devices, without becoming overwhelmed. Viewing the network as a whole, as opposed to a collection of individual devices, leads to more efficient processing of data and faster detection of anomalies.***

What is the level of policy detection granularity?

Some NDR solutions use micro-segmentation principles to define groups of devices and set policies for them.

> ***ThreatWarrior NDR goes much further than traditional micro-segmentation.The ThreatWarrior AI engine detects ad-hoc applications on your network and will notice deviations in the behavior of any device. It's more nano-segmentation than micro-segmentation in nature - having separate policies for each and every device on your network, all defined by observations of normal behavior without the analyst having to set up anything.***

How is threat severity established and managed?

Some vendors rank threat severity in terms of specific exploits/vulnerabilities, e.g., 'CVE-XXX-XXXX Print Bomb: Certain versions of Windows 10 are vulnerable to a local exploit that can allow privileged access...'

> ***In contrast, ThreatWarrior NDR ranks threat severity by event. Each signature gets ranked differently. Some signatures are too basic to have an exploit, e.g., 'this is running on the wrong port, that's odd'. Other threats are attached to specific exploits or vulnerabilities, but there can be many attached to the same one. Some might be stronger indicators than others, which will drive a higher score. Others might be weaker and get scored lower.***

Does the solution use client data pooling?

Some NDR vendors do not isolate data on a client-by-client basis. At some level their system will pool data across customers, enabling broad alerting when a given company has been exploited.

> ***ThreatWarrior NDR rolls up alerts by alert type, and on a per customer basis only.***

Visualization

Is visualization 2D and static or 3D and dynamic?

Many NDR visualization tools are two dimensional and static. They only present a flat, point-in-time snapshot of devices with lines between them for alerts. While better than a traditional table view of activity, it provides very limited value in actual detection and response work performed by analysts.

> ***ThreatWarrior 3D Universe provides a spherical view of your network - including all devices (each labeled with the manufacturer logo), connections, protocols in use, alerts, and the ability to drill into extensive detail associated with any of these elements.***

Are contextual alerts, notification and actions provided?

NDR solutions routinely provide alerts and notifications. But when they are not contextual, it creates more work for analysts as they now need to manually investigate and analyze the alert to determine if it is a true positive or a false positive, and subsequently assess its potential impact on the network. For example, if an alert is triggered for a large amount of traffic being sent from one device to another, without contextual information, the analyst may not be able to determine if this is normal behavior or indicative of a security breach.

> ***With ThreatWarrior, analysts can navigate around a 3D view of the network and click exactly where they want for detailed information. Alerts, notifications, and suggested actions are presented when something relevant happens worthy of attention - reducing the amount of manual investigation needed and helping analysts to more quickly identify and respond to real security threats.***

Are analysis-defined views supported?

Most NDR solutions provide some sort of network visualization - albeit typically two-dimensional and static.

> ***ThreatWarrior NDR allows analysts to dynamically view either the entire network in the aggregate, or broken down into analyst-defined subviews, i.e., "orbs". Orbs can be defined by geography or organizational constructs, e.g., departments. Orbs can be saved and later accessed for instant drill down.***

Can traffic views be segmented?

Most NDR solutions provide traffic visualization. But traffic may only be node-path-node in nature - unable to isolate traffic destined only for addresses on a given network subsegment, vs cross-segment or cross public-private boundary.

> ***ThreatWarrior NDR comprehends how packets get rewritten on a network. As a result, ThreatWarrior 3D Universe can represent views of traffic concealed to interior networks, across differing network boundaries, or destined for public network transit, including connections to a different part of the world, offering a more encompassing view of the network connection environment.***

Can queries be contextual?

Typical NDR solutions only provide for queues and filtering in the aggregate.

> ***ThreatWarrior NDR goes deeper and provides not only a single global view, but also statistically grouped, and customizable orbs.***

Can network communications be replayed?

Some NDR solutions have the ability to display a series of two-dimensional network graph diagrams with a 'time slice' of communications at once, i.e., a time-lapse photo where cars blur into streaks of light.

> ***ThreatWarrior's three dimensional, animated visualizer provides a dynamic and interactive display of the network environment. With ThreatWarrior NDR, network communications can be recreated, step-by-step, and then replayed.***

Automation / Integration

Does the solution integrate with network enforecement points?

NDR solutions should be able to integrate with firewalls to block access.

> ***ThreatWarrior's three dimensional, animated visualizer provides a dynamic and interactive display of the network environment. With ThreatWarrior NDR, network communications can be recreated, step-by-step, and then replayed.***

Can network communications be replayed?

NDR solutions should be able to integrate with endpoint agents, providing access to endpoint telemetry.

> ***ThreatWarrior supports out-of-the-box integrations with leading EDR vendors including CrowdStrike and CarbonBlack.***

Can you automate and scale analyst workflows?

Clearly AI in its many forms can inspect network traffic and make a variety of anomaly and IOC detections. But, more advanced NDR solutions go beyond just learning from network traffic itself.

> ***ThreatWarrior NDR learns not only from observing network traffic, but also from how analysts respond to threats. It correlates network events with actual analyst behavior, then applies advanced heuristics and behavioral modelings to anticipate which security incidents are most relevant to your team. This powerful***

How is threat intelligence integrated?

Some vendors' threat intelligence integration is merely a raw text dump into a search repository.

> ***When ThreatWarrior NDR integrates open source intelligence, rules and signatures are gathered from the wider security community, enriched with our own context, and automatically deployed in the field. If a new vulnerability drops, we don't simply ping customers with a news article. Our sensors know to watch out for it on their networks and alert if seen.***

Summary

We hope these questions and explanations will prove informative and helpful as you navigate the process of selecting an NDR solution that meets your organization's unique requirements. We believe ThreatWarrior NDR offers the most comprehensive protection against potential network threats. If you're interested in learning more, contact us for a short demo and to discuss how our platform can help safeguard your network.

About ThreatWarrior

ThreatWarrior is a leader in network detection and response (NDR) for on-premises, hybrid and multi-cloud enterprises. Our AI-powered platform helps organizations identify threats in real time and stop them before they cause breaches. By combining complete visibility, deep packet inspection, behavioral anomaly detection, forensics and threat hunting, ThreatWarrior delivers the context security teams need to detect, investigate, and respond to threats. Leading organizations use ThreatWarrior to defend against APTs, zero-day exploits, ransomware, and more.

Copyright © 2023 ThreatWarrior, Inc.

THREAT WARRIOR

info@threatwarrior.com
threatwarrior.com
844.463.9440