DARKTRACE

CISO's Guide to Buying AI

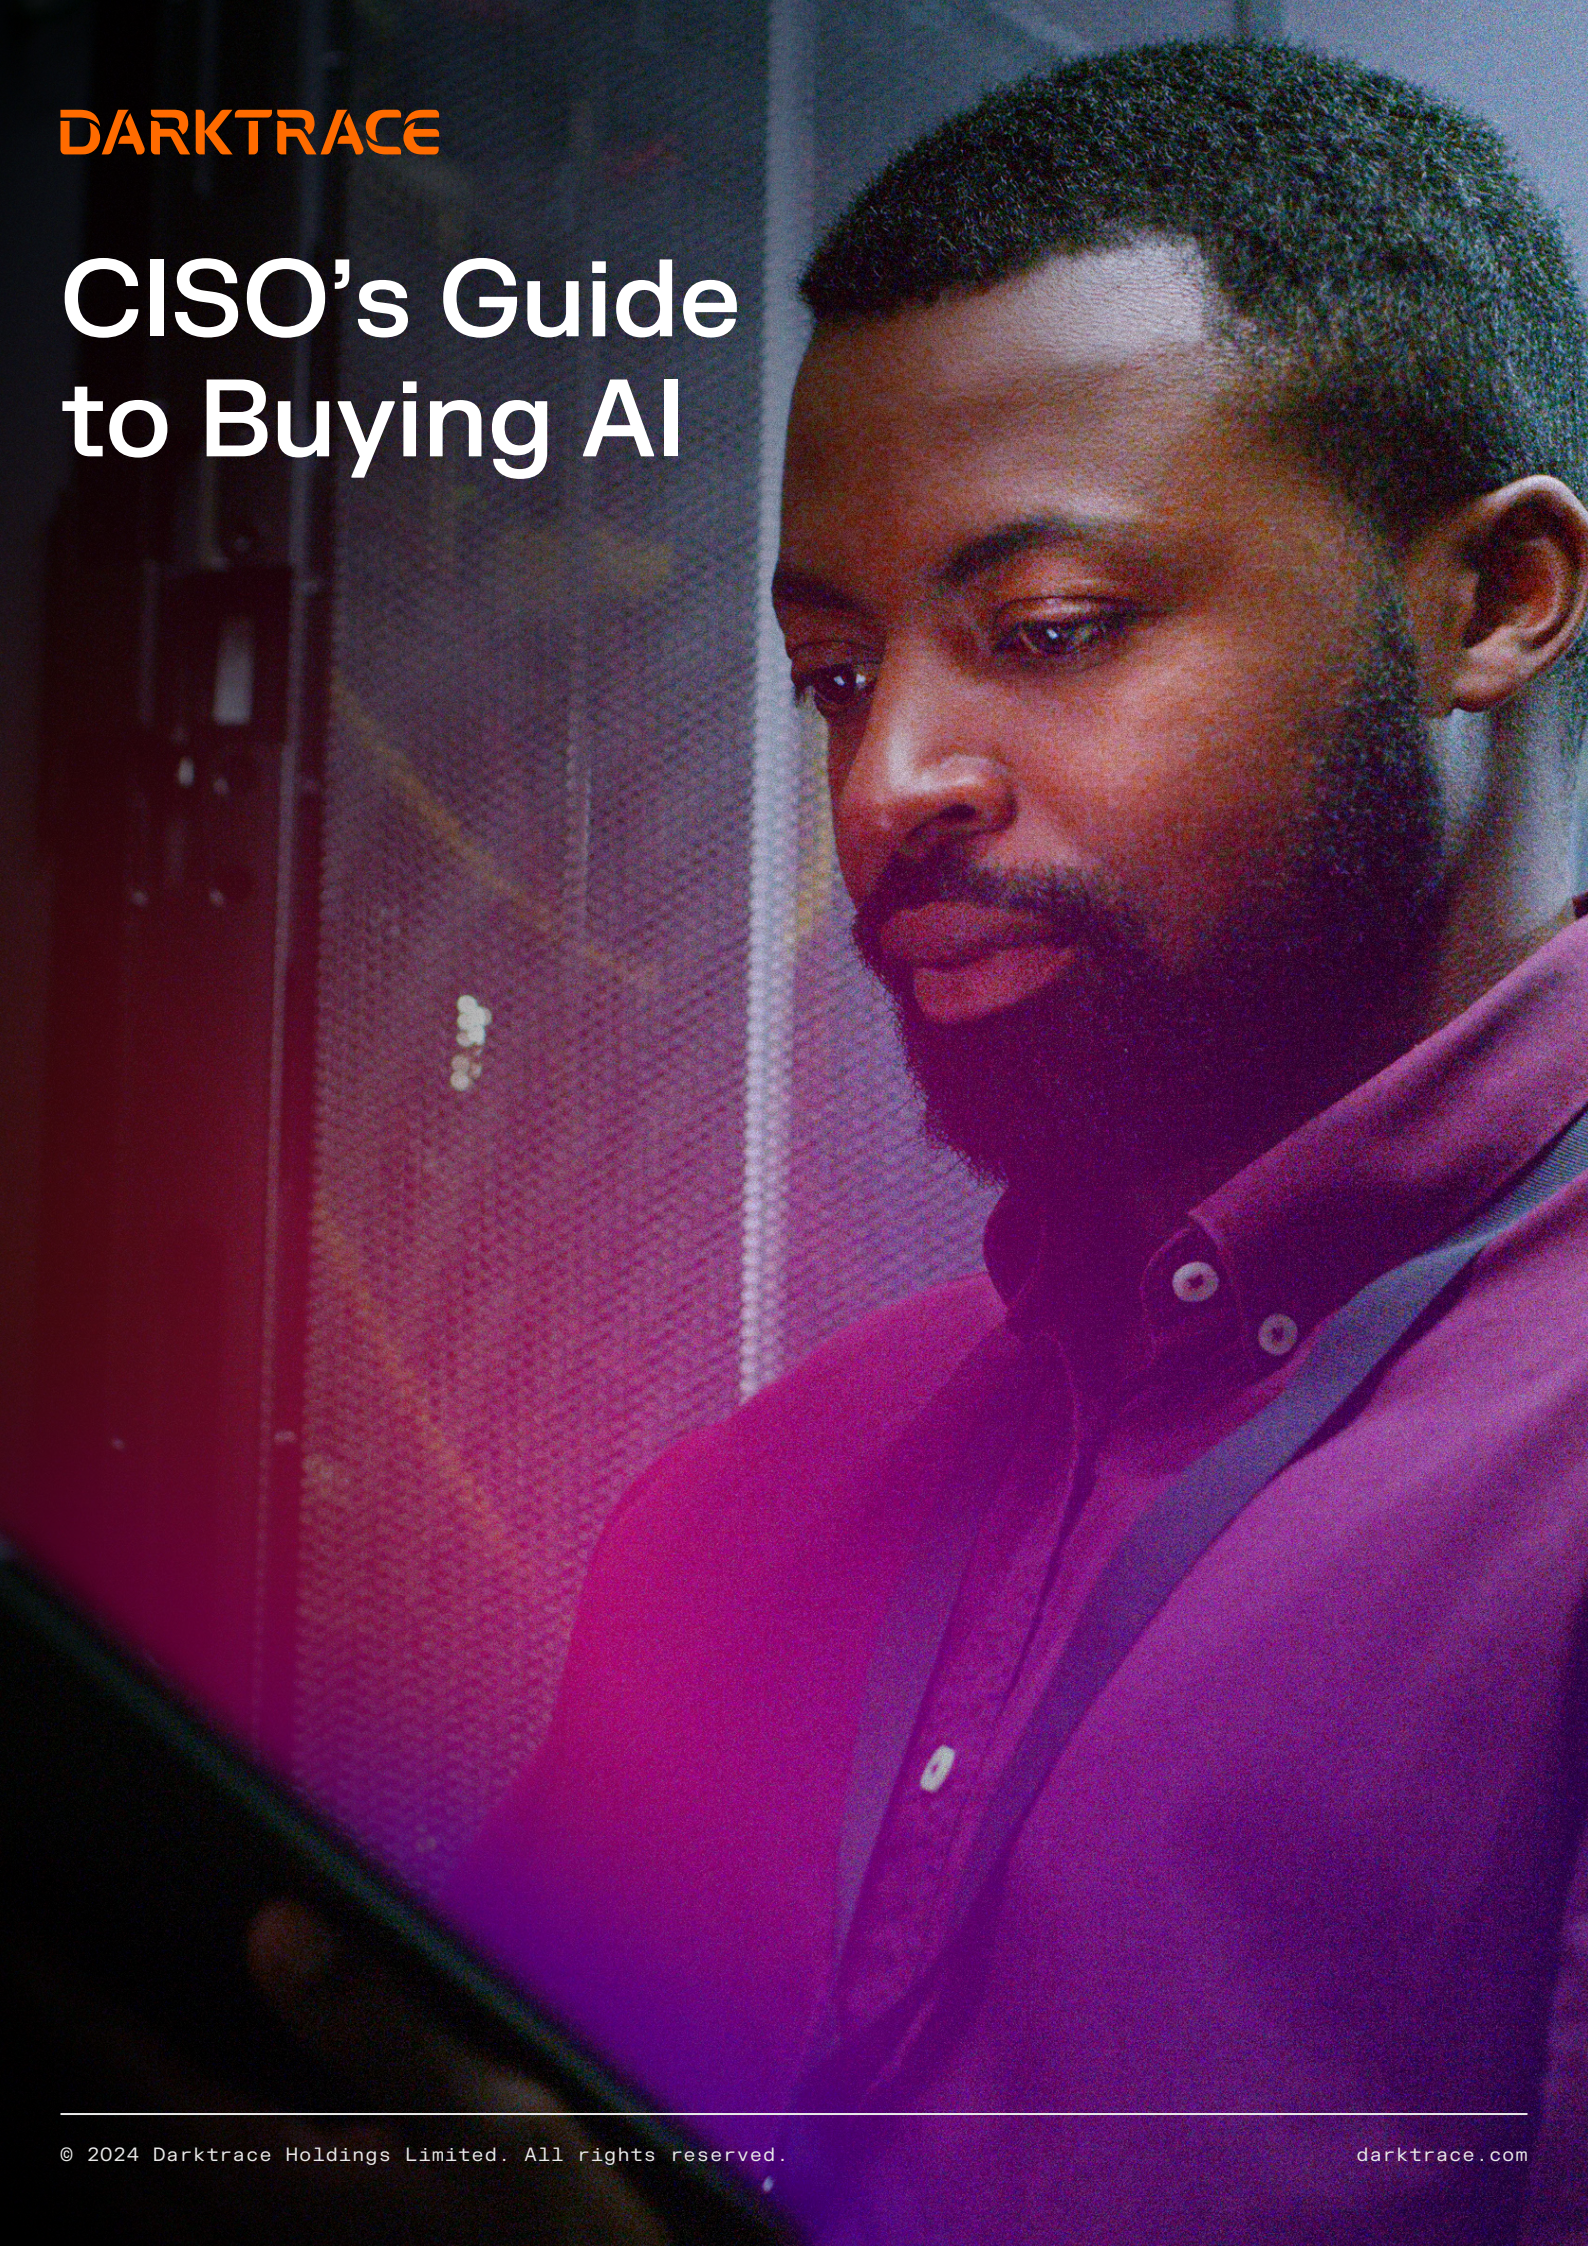

Contents

Abstract

First-time buyers of AI cybersecurity solutions may not know where to start. They must determine the type of tool they want, know the options available, and evaluate vendors. Research and understanding are critical to ensure purchases are worth the investment.

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. When evaluating AI cybersecurity tools, buyers should ensure they address themes of security, accuracy, control, ethics, data privacy, and interpretability.

Section 01

How can you make the most of your AI investment?

In this dawning Age of AI, CISOs are increasingly exploring investments in AI security tools to enhance their organizations' capabilities. AI can help achieve productivity gains by saving time and resources, mining intelligence and insights from valuable data, and increasing knowledge sharing and collaboration.

Investing in AI can bring immense benefits to your organization, but how do you make the most of your budget?

A muddied marketplace

Key challenges in AI purchasing come from consumer doubt and lack of vendor transparency. The AI software market is buzzing with hype and flashy promises, which are not necessarily going to be realized immediately. This has fostered uncertainty among potential buyers, especially in the AI cybersecurity space.

As Gartner writes, "There is a general lack of transparency and understanding about how AI-enhanced security solutions leverage AI and the effectiveness of those solutions within real-world SecOps. This leads to trust issues among security leaders and practitioners, resulting in slower adoption of AI features."[1]

Given this widespread uncertainty generated through vague hype, buyers must take extra care when considering new AI tools to adopt.

1 Gartner, April 17, 2024, "Emerging Tech: Navigating the Impact of AI on SecOps Solution Development."

Investigation before investment

Before investing in AI tools, buyers should ask questions pertaining to each stage of the adoption journey. The answers to these questions will not only help buyers gauge if a tool could be worth the investment, but also plan how the new tool will practically fit into the organization's existing technology and workflows.

Table 1: Initial questions to consider when starting to shop for AI.[2]

01 Problem identification	02 Product assessment	03 Implementation and policy	04 Procurement and deployment
What is the problem to be solved with a new AI tool? What are the machine learning models within the solution? What are the strengths and limitations of those machine learning techniques for this use case? Ex. Does this machine learning technique use data analysis, time series analysis, semantic analysis, probabilistic models, generative capabilities, static or pre-trained models, or continuous learning?	Are the performance claims valid? Does it comply with regulatory standards? Is the tool interpretable? Can I determine how the results are generated? How accurate are the results? Does the security team have control on productizing the output? Does the product provide the relevant data in order to facilitate workflow optimization or data-driven decision making	Does it work in practice? What are my company's acceptable use policies, necessary access rights, and data classification? What do my end-users need to do to support the new tool? How do they productize the output or optimize their workflow? What new policies do we need to manage trust, risk, and security? Will the data be protected and private? What ongoing monitoring, testing, validation, and governance are necessary to ensure functionality? What are the security vulnerabilities that need to be mitigated with the deployment of this product?	Is the procurement process compliant? Is the deployment healthy? How do I optimize or configure the workflow? Does the tool have the appropriate integrations with the rest of my technical or security stack? Does the tool provide the needed visibility, audit logs, control, and security measures to implement securely? Does the tool allow for fine-tuning or reinforcement learning and what is that process?

2 Inspired by Gartner, May 14, 2024, "Presentation Slides: AI Survey Reveals AI Security and Privacy Leads to Improved ROI" and NHS England, September, 18, 2020, "A Buyer's Guide to AI in Health and Care," Available at: https://transform.england.nhs.uk/ai-lab/explore-all-resources/adopt-ai/a-buyers-guide-to-ai-in-health-and-care/

Goals of AI adoption

Buyers should always start their journeys with objectives in mind, and a universal goal is to achieve return on investment. **When organizations adopt AI, there are key aspects that will signal strong payoff. These include:**

- Wide-ranging application across operations and areas of the business
- Actual, enthusiastic adoption and application by the human security team
- Integration with the rest of the security stack and existing workflows
- **Business and operational benefits, including but not limited to:**
  - Reduced risk
  - Reduced time to response
  - Reduced potential downtime, damage, and disruption
  - Increased visibility and coverage
  - Improved SecOps workflows
  - Decreased burden on teams so they can take on more strategic tasks

Ideally, most or all these measurements will be fulfilled. It is not enough for AI tools to benefit productivity and workflows in theory, but they must be practically implemented to provide return on investment.

Section 02

What types of AI can you choose between?

Components of AI

Before investing in AI tools, it will help to understand the basics of AI. In the simplest terms, AI is composed of three elements:

Data
meaning the data storage in any way, shape, or form

AI algorithms and models
which are sets of mathematical equations, technical computations, and instructions on the data itself

APIs or connections between the above

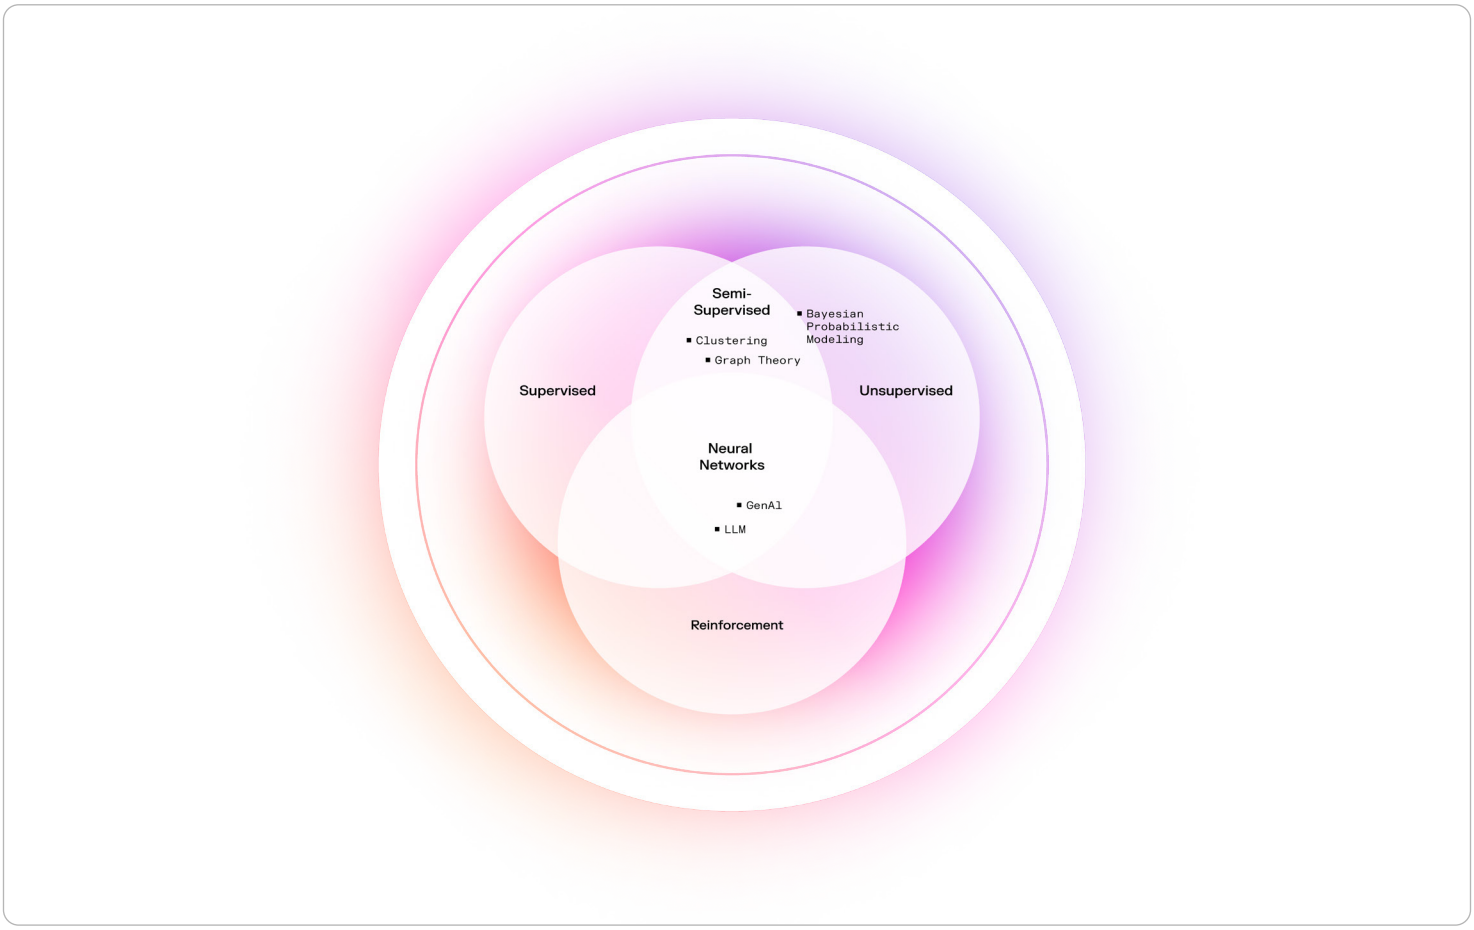

Figure 01: The AI techniques commonly mentioned by AI cybersecurity vendors can be sorted into the four types of machine learning.

Types of machine learning

Machine learning is a subset of AI that enables machines to develop problem-solving models by identifying patterns in data instead of leveraging explicit programming. The learning refers to the training process.

There are four main ways machine learning can be categorized. The main distinctions are the objective of that machine learning output and what data the system needs to get there.

01 Supervised Learning

feeds known, labeled data into algorithms that can create outputs as closely aligned with the desired result as possible. This technique is used for classification and regression.

Example: threat detection based on large corpuses of reported malicious activity; email spam detection

02 Unsupervised Learning

uses unknown, unlabeled data with predictive models that can discover patterns. This technique's training is not led by humans, the model identifies patterns based on data ingested. In contrast to supervised learning, which requires labeled data to train the model, unsupervised learning models discover patterns, structures, or relationships within the data on their own.

Example: image or video analysis;
clustering similar behaviors, patterns, or assets

03 Semi-Supervised Learning

builds a model through a mix of labeled data and unlabeled data, using unsupervised learning to predict data labels to self-train the AI. This technique is useful when limited labeled training data is available and/or to combine human-led training with unsupervised learning to identify previously unknown patterns. Natural Language Processing (NLP) and Large Language Models (LLMs) can fit into this category.

Example: improving recommendation systems
or improving labeling data and images

04 Reinforcement Learning

is based on rewarding desired behaviors or punishing undesired ones. Instead of one input producing one output, the algorithm produces a variety of outputs and is trained to select the right one based on certain variables. So, it can use unlabeled data.

Example: Human feedback of output of a LLM

Applying AI to cybersecurity

These machine learning techniques can be applied to cybersecurity to improve efficiency and workflows. While applications can vary greatly, there are three general methods that are most common.

Table 2: There are three common ways to apply AI to cybersecurity.

	Training Data	Output Type	Models	Common Use Case	Main Limitation
Attacker-centric supervised machine learning	Pre-trained or static data and consistently retrained	Known attack patterns	▪ Identify identical or similar attacks ▪ Attack simulations ▪ Supervised ML classifiers trained for specific tasks	▪ Detects known attacks ▪ Detects known attack variants	▪ Cannot detect un-known, novel attacks, or insider threats ▪ Exhaustive data integrity process
Business-centric unsupervised machine learning	Continuously learning	Organization-specific behavioral data and patterns	▪ Anomaly detection ▪ Probabilistic modeling ▪ Relationship analysis ▪ Attack simulations ▪ Unsupervised ML to perform learning, data correlation and analysis	▪ Detect unusual behavior ▪ Identifies known, unknown, novel threats ▪ Identifies abuses, misuses, misconfigurations	▪ Requires many ML techniques and modeling to translate anomaly detection into attack identification
Open-source generative AI and LLMs	Pre-trained and feed-back from users	Internet or data lake	▪ Language, visual content, or audio ▪ Text-based LLMs ▪ Unsupervised machine learning with semantic analysis	▪ Data retrieval optimization ▪ Content summarization ▪ Content generation	▪ Detection and response based on semantic analysis ▪ Open to confirmation bias ▪ Accuracy optimization ▪ Latency

Attacker-centric supervised machine learning

One of the most common types of AI in security today is supervised machine learning. These models are commonly found in eXtended Detection and Response (XDR) solutions. They are trained on massive volumes of structured, labeled attack data and threat intelligence, and they perform extremely well at stopping those known attacks. This makes them a good starting point for any security stack.

However, these models can fall short when they encounter something they haven't seen before. If the model hasn't been trained on a specific pattern, it can easily miss it. Additionally, this is a big data solution which causes data integrity issues. Co-mingled benign or legitimate data (syslog, network traffic, etc.) can cause big problems in the efficacy and accuracy of this AI's performance. These classifiers need to be re-trained as soon as new attack patterns have been reported. This is computationally expensive. Also, data typically needs to leave the environment to be run through these large classifiers. And, just like most AI, testing and validation is crucial to ensuring accurate outcomes.

Given the volume of data, this large corpus of signal can provide external context intelligence to assets outside of an organization.

Business-centric unsupervised machine learning

This type of AI continuously learns what is normal for an organization, device, account, user, and/or cluster. By understanding normal, this AI can identify abnormal behavior, patterns, and activity, which may indicate a cyber-attack.

Many products state that they are in behavior analysis or anomaly detection. Anomaly detection in isolation may not be that accurate. Harnessing unsupervised machine learning to get insight into the pattern of life of every asset within an environment enables more effective anomaly detection. Also, a layered machine learning model approach helps drive accuracy and efficacy.

This anomaly detection approach can provide insights that traditional cybersecurity tools, which only use supervised machine learning, tend to miss. Since business-centric unsupervised machine learning learns the organization, its people, and its devices, it can detect every kind of attack. That includes ones that are unknown and novel, like attacks exploiting zero days, attacks that exploit legitimate credentials, account takeovers, insider threats, lateral movement, and cross domain attacks.

Open-source generative AI and LLMs

Publicly available, text-based generative AI systems, which are powered by LLMs, are pre-trained on massive volumes of internet data and can be applied to human language, machine language, and more. For consumers, LLMs are game-changing, as we are already seeing. However, many LLMs will likely remain the domain of existing Big Tech players, as training these systems requires access to massive volumes of data and computing, as well as significant manual labor for validating, cleansing, and preparing the data for processing.

Generative AI and LLM tools promise increases in productivity and new ways of augmenting human creativity. Generative AI performs semantic analysis on large corpuses on data and it excels at learning the fluidity of language and impersonating humans. This will be a large net gain in the future for translation abilities.

However, employee adoption of these tools introduces risks around privacy, especially through the lenses of data exfiltration, giving away business strategies or competitive advantages, and carries legal implications. For example, an employee could input proprietary information as part of a prompt to ChatGPT.

In security, generative AI can optimize data retrieval for defenders by converting natural language queries to power queries across tools and different data sets. It can also be used to summarize in the reporting process. But, since this is semantic analysis, LLMs are not the best machine learning model for threat detection.

The right machine learning model should consider the type of data it is analyzing and how to drive accuracy and optimization.

■ Tip for Buyers

You cannot only invest in one AI technique. Every AI method has its own strengths and weaknesses, so one machine learning model or technique alone is not sufficient. The more layered an AI approach, the more diversity of perspective it has, the better outcomes and corroboration it will yield.

Generative AI and many LLMs are trained on billions of parameters of static data scraped from the internet. This makes them ideal for tasks like emulating sophisticated phishing attacks for preventative security or creating simple to use querying mechanisms for better human interaction. Another ideal scenario is noticing shifts or changes in language (semantic analysis) which can facilitate anomaly detection with other machine learning models within email.

However, if not applied responsibly, generative AI can cause confusion by "hallucinating," where it references invented data, or by providing conflicting responses due to confirmation bias in the prompts written by different security team members.

In addition, prompt-based models are insecure by design. Prompt injection risks are known, prevalent, and exhaustive attack vectors for LLMs. This vulnerability needs to be mitigated as well as exhaustively tested.

■ Learn more

About what industry professionals think about AI in cybersecurity by reading our report "State of AI Cybersecurity 2024."

How can you evaluate AI vendors?

Validation

An easy starting place to evaluate AI vendors would be to look for proof points and customer testimonials. These will provide evidence that the AI tool works, and that it can also work for your organization. Especially look for insight from customers in similar situations as your own, whether that be industry, organization size, region, or other attributes.

Asking deeper questions

Earlier in this white paper, we shared questions to ask at the beginning of your AI acquisition journey. Those questions are good to imagine how a tool will fit into your organization and determine if a vendor is worth further evaluation. Once you decide a tool has potential use and feasibility in your organization, it is time to dive deeper and learn more.

Ask vendors questions about their technology.

This information will most likely not be on their websites, and since it involves intellectual property, it may require a NDA.

AI techniques

- What are the strengths and limitations of this specific AI approach?
- Everyone uses AI, how are this vendor different?
- How much computational power is required for the AI tool, and how will it adapt to increasing data volume or area of coverage?

Models

- Are the models pre-trained or continuously learning? Why was that approach chosen?
- If pre-trained, how often do they retrain or update?
  - How does the AI deal with environment changes (e.g., migrations, new sites acquired, etc.)?
  - How does the AI deal with threat landscape changes?
- How long are behavioral models retained for?
- How are false positives and false negatives handled? Is there a process in place to learn from them?
- What is the Test, Evaluation, Validation, and Verification (TEVV) process to measure the accuracy of the AI as well as test for Adversarial Machine Learning (AML)? Is this continual?
- What is the security around permissions for model access?

Training data

- Where does the tool get training data, and what is that data's volume and variety?
- What is the data integrity process? Is it continual?
- What measures have been used to prevent bias and data poisoning?
  - How does my organization avoid poisoning our baseline if there is already compromise or hygiene issues internally?
- How is my organization's data used and stored? How is privacy maintained?

Interpretability

- Does the vendor explain the AI with confidence scores, feature importance, or explainability?
- Can a human interact with the AI to build a trusting relationship?
- What level of control does a human have over the AI's decision making?

AI governance

- What are the security vulnerabilities with this machine learning approach? How are the mitigated?
- Does the AI tool comply with NIST's AI Risk Management Framework?

Interpreting the answers

There are no correct answers to those questions. The answers will vary drastically based on the vendor. Additionally, in some cases, you will hope to hear different answers depending on what you want your AI investment to accomplish.

There are, however, some key themes to look for as you learn more about possible AI vendors. No matter what, you want an AI tool that is reliable, which means it must be:

- Accurate
- Secure
- Interpretable
- Controlled
- Ethical
- Privacy preserving

■ Tip for Buyers

A reliable AI tool is a safe AI tool. Without security to ensure trustworthy AI, organizations' investments in new technologies could be wasted. AI security needs to be embedded across every step of an AI system's creation and deployment.

To learn more about AI security, read our whitepaper "Best Security Practices for Implementing AI in the Enterprise."

Let's look at how those themes can show up in the answers to the deeper questions:

AI techniques

As previously mentioned, there is no correct type of AI. This will depend on the vendor and your use case.

Models

Accurate:
Models should continuously learn or frequently update if data is prone to change. Models should be consistently tested and measured for efficacy.

Secure:
Look for strong protections with selective permissions and access. Also, evaluate potential security vulnerability mitigations.

Data

Accurate:
Data integrity should be consistent for reliable AI outputs. Look for an existing TEVV plan.

Ethical:
AI tools should consider AI ethics to reduce potential biases.

Privacy preserving:
Data should be secured to avoid data leaks.

■ Tip for Buyers

Data is essential for reliable AI outputs. Make sure you have a clear plan for sourcing, updating, and protecting the AI tool's training data.

Interpretability

Interpretable:
The output of models needs to be accurate and be able to be productized. This comes with confidence scores, explanations, and feature importance.

Controlled:
There should be trust between humans and AI tools. To build that trust, the human users should have visibility and control over AI decision making. This can be achieved with AI that shows its work with reporting and by having customizable guidelines for the AI.

AI governance

Ethical:
Tools you adopt must comply with regulations and best practices.

■ Tip for Buyers

For security teams to make the most out of new AI tools, they must trust the AI. Interpretability and control help teams build a trusting relationship with new AI tools.

Darktrace as an AI cybersecurity vendor

With this baseline understanding of machine learning techniques, how they fit into cybersecurity, and the themes to look for in their applications, let's see how Darktrace functions.

Our approach

Darktrace has been using AI technology in cybersecurity for over 10 years. As a pioneer in the space, we have made innovation part of our process.

The Darktrace ActiveAI Security Platform uses multi-layered AI that trains on your unique business operations data for tailored security across the enterprise. This approach ensures that the strengths of one AI technique make up for the shortcomings of another, providing well-rounded and reliable coverage. Our models are always on and always learning, allowing your team to stop attacks in real time. Additionally, since Darktrace focuses on using the customer's data across its entire digital estate, it brings a range of advantages in data privacy, interpretability, and data transfer costs.

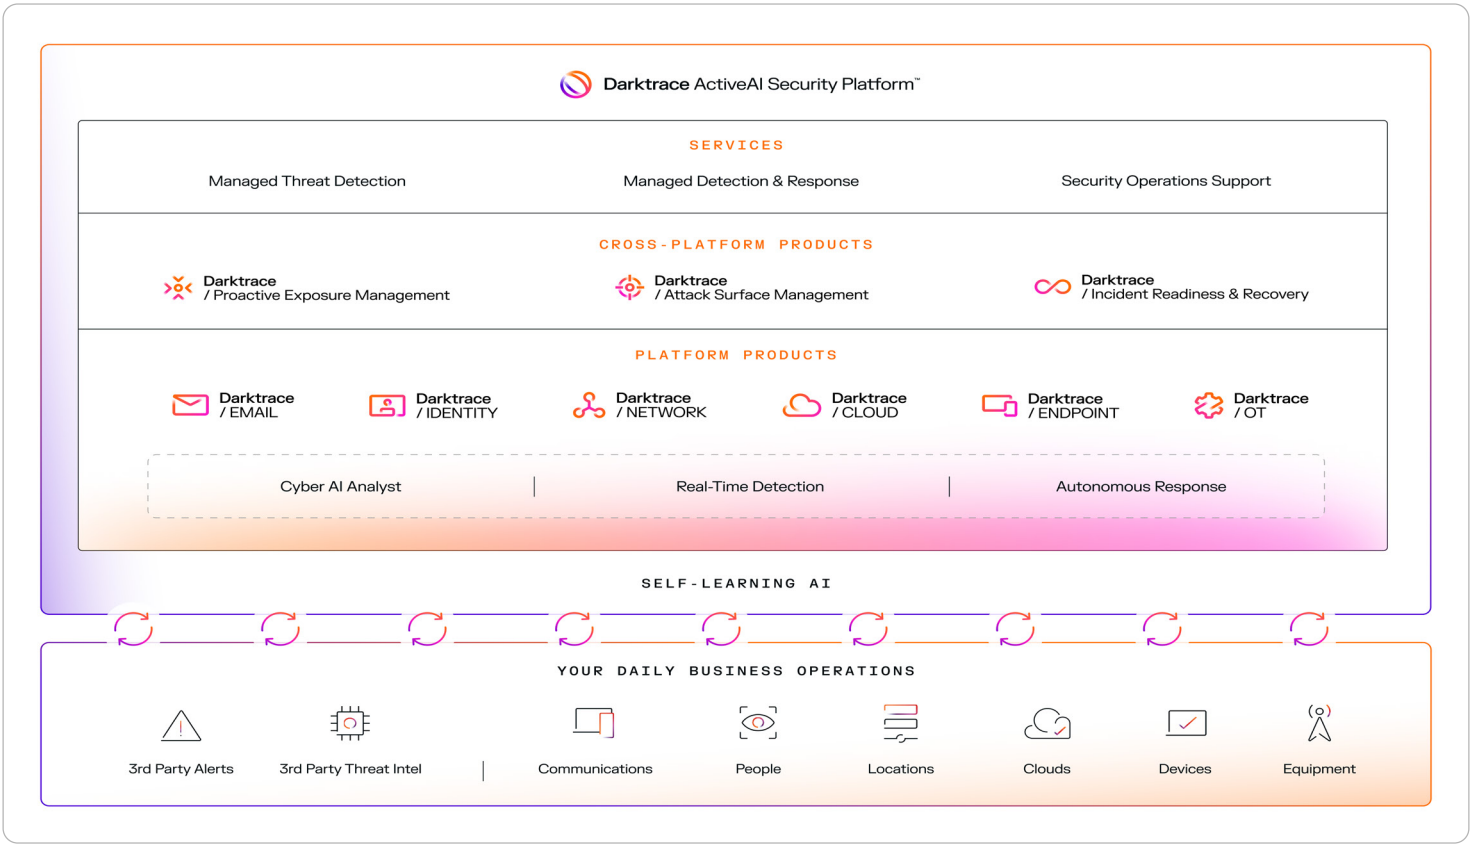

Figure 02: The Darktrace ActiveAI Security Platform uses AI to empower security teams with the speed and scale needed to handle contemporary challenges.

■ Tip for Buyers

Strong ROI often comes from tools that can be applied across several processes and areas of the business. Look for platform-based approaches, integrations with existing investments, and coverage across the entire incident lifecycle.

We use various types of AI

Self-Learning AI

Darktrace's self-learning technology is unique in its ability to understand and detect novel attacks. It contains several types of machine learning algorithms together, including pairing business-centric unsupervised machine learning with other models to translate anomaly detection into attack identification. This makes it an absolute must-have in any future-facing security stack. For example, traditional email security tools that rely on knowledge of past threats take an average of 13 days from an attack's being launched to detect it. In contrast, Darktrace / EMAIL is capable of spotting and stopping threats as soon as they are launched. [3]

Today, we are applying our Self-Learning AI to augment human security teams throughout each stage of cyber resilience:

- **Prevention:** With Darktrace / Attack Surface Management and Darktrace / Proactive Exposure Management, we are applying AI to help security practitioners harden security inside and out. The AI continuously monitors the attack surface for risks, high-impact vulnerabilities, and external threats. It also looks inside the environment to expose potentially vulnerable attack paths and high value targets. Darktrace / Proactive Exposure Management further supports resiliency with tabletop exercises and augmented pentests.
- **Detection:** Darktrace Real-Time Detection uses AI to uncover threats by analyzing thousands of metrics at machine speed and revealing subtle deviations which a human wouldn't see but that may signal an evolving threat. This includes even unknown techniques and novel malware that may bypass all other security controls.
- **Response:** Our AI enables Darktrace Autonomous Response to take action against attacks, potentially bringing response times down from hours and days to mere seconds. Autonomous Response is a surgical action against critical anomalous or risky behavior, containing the incident, stopping potential spread, reducing the amount of damage caused, and buying humans time to perform an investigation and remediation.
- **Recovery:** Darktrace / Incident Recovery and Response can help organizations assess their readiness for an attack and practice with real-world scenarios. During an attack, AI helps to prioritize remediation actions to augment human teams. Post attack, our AI allows businesses to recover and get back to full operations faster and more confidently than a human team can alone.

Darktrace's Self-Learning AI was built around four core principles:

- **It learns 'on the job'** – it does not depend upon knowledge of previous attacks.
- **It learns on real business data** and thrives on complexity and diversity of modern businesses.
- **It constantly revises assumptions** about behavior, using probabilistic mathematics.
- I**t is always up to date** and not reliant on human input.

Applied supervised machine learning

When using supervised machine learning in raw threat detection, the AI will fail to spot novel, never-before-seen threats because it is restricted by the historical attack data on which it is trained. But in another application of cybersecurity, supervised machine learning can excel. Specifically, automating the initial analyst triage of anomalies on the network, establishing which are dead ends, and which may be part of a wider and more significant security incident.

Darktrace has incorporated supervised machine learning into its product stack through its Cyber AI Analyst. This investigation tool combines a set of machine learning techniques, including unsupervised machine learning, supervised machine learning, and models trained on expert cyber analysts to significantly expedite triage and reduce time-to-meaning for security teams.

Cyber AI Analyst takes individual anomalous events discovered by Darktrace's core Self-Learning AI, and then applies a second layer of AI to these findings, using supervised machine learning to correlate alerts into incidents that are then surfaced for human review, allowing analysts to focus on a few prioritized, critical incidents instead of thousands of alerts. It generates incident summaries that highlight every stage of the attack and includes a natural language summary that even a non-technical person can understand.

■ **Over a 90% reduction**
in triage time reported by security teams using Darktrace.

3 Thirteen days mean average of phishing payloads active in the wild between the response of Darktrace / EMAIL compared to the earliest of 16 independent feeds submitted by other email security technologies.

NLP

In addition to generating coherent, easily-understandable reports on significant cybersecurity incidents, Darktrace's Cyber AI Analyst makes its investigation process transparent with explainable AI powered by NLP – showing what questions it asked before arriving at its conclusions. This is in stark contrast to the 'black box' model that risks eroding trust between humans and AI, which can present problems with compliance and audit requirements.

Darktrace uses NLP elsewhere in its product suite to intelligently map Advanced Persistent Threat (APT) capabilities – to identify threat actors likely to target an organization and assess the organization's susceptibility to their known approaches and methods, allowing the organization to take the appropriate preventative actions. Darktrace has a history of using techniques from the NLP space, which have formed the basis from which LLMs have emerged. We use the best technique that has often been borrowed from NLP and adjacent technologies and continue experimenting with techniques such as n-grams, Long Short-Term Memory (LSTM) networks, and transformers for novel use cases.

LLMs

When it comes to LLMs, the integrity of their outputs relies on the data on which they are trained. That is why the Darktrace models have been trained on our proprietary security data and are then applied to each customer's specific environment to build contextual datasets. This ensures both quality and relevance.

Darktrace first applied LLMs to our product set with a feature that looks for emerging attacks targeting our customers and shares behavioral signals with other participating customers. LLMs were used in this case to categorize malicious communications based on textual properties. Subsequently, we began using LLMs in Cyber AI Analyst to try and understand what the purpose of a certain hostname is in a more heuristic way, by trying to classify the known internet.

This improves the precision of detections.

Further improvements and quality-of-life changes for Darktrace customers arise from Cyber AI Analyst's proprietary LLM classifier to categorize malicious communications based on textual properties. Crucially, these potential security incidents do not just relate to malicious activity by attackers' attempting to breach networks from outside – they can help to detect data leaks by employees too.

In addition, Darktrace now uses LLMs to augment attack engagements designed to proactively harden defenses while providing security awareness. Bringing LLM-generated attack emulation capabilities alongside existing NLP-derived attack engagements allows attacks to be emulated in a wider range of sophistication and control the level of complexity to meet customer needs.

Table 3: Darktrace's multi-layered AI includes many different machine learning techniques.

Unsupervised machine learning	Establishes pattern of life and identifies rare and risky behavior (when paired with neural networks, clustering methods, regularization, and probabilistic and anomaly detection)
Bayesian probabilistic methods	Allows models to be efficiently updated and controlled in real time
Generative and applied AI	Runs simulated phishing campaigns, tabletop exercises, and realistic drills
Deep learning engine	Replicates the thought processes of humans
Graph theory	Understands the incredibly complex relationships between people, systems, organizations, and supply chains
NLP and LLMs	Interprets and produces human consumable output
Post-processing the output of our AI engine	Exposes risk scores with transparency and explainability and offers control of operationalizing the output (assigning value, priorities, thresholds, and actions)
Supervised machine learning and neural networks	Cyber AI Analyst performs autonomous investigations, incident containment, and intelligence for AI-generated playbooks

Human teams maintain visibility and control

Darktrace supports the human security team's adoption of our technology by building trust. To do that, we designed our platform to give your team visibility and control over the AI.

Darktrace focuses on interpretability and sharing confidence levels. This includes specifying the threshold of what triggered a certain alert and the details of the AI's investigations to see how it reached its conclusions. The interpretability of our AI uplevels and upskills the human security team with more information to drive investigations and remediation actions.

The human security team can also modify all the detection and response thresholds for our model alerts to customize them to fit specific business preferences.

AI security measures

Darktrace recognizes that different organizations have different privacy cultures and regulations. The Darktrace ActiveAI Security Platform is therefore adaptable to differing levels of privacy needs.

Customer data is stored either on site on a physical appliance, or on a cloud appliance which is accessible to customers. A core principle of Darktrace's approach is to perform all analysis in a non-invasive manner, analyzing metadata rather than content. Examples of metadata analyzed include the time that an asset was created, the byte size of a document, and the type of file created.

In addition, every action taken in the Darktrace interface requires granular permissions. If desired, user visibility can also be restricted to specific subnets. Audit logs allow our customers to see all actions taken from the interface, therefore boosting accountability, interpretability, compliance, security, and troubleshooting.

Conclusion

CISO's are increasingly considering investing in AI cybersecurity tools, but in this rapidly growing field, it's not always clear what to look for.

Buyers should first determine their goals for a new AI tool, then research possible vendors by reviewing validation and asking deeper questions. This will reveal if a tool's security, accuracy, control, ethics, and privacy measures are a good match for the organization to move forward with investment and adoption.

As leaders in the AI cybersecurity industry, Darktrace is always ready to help you on your AI journey.

Learn more about how Darktrace uses AI to protect your network.

Read the solution brief →

Discover the influence of AI on email security in the webinar featuring Forrester: Exploring the Evolving Email Security Landscape.

Read the solution brief →

Contact Darktrace for a personalized demo.

Get started →

About Darktrace

Darktrace (DARK.L), a global leader in cybersecurity artificial intelligence, is on a mission to free the world of cyber disruption. Its technology continuously learns and updates its knowledge of 'you' for an organization and applies that understanding to help transform security operations and improve cyber resilience. Breakthrough innovations from its R&D Centers have resulted in more than 200 patent applications filed. Darktrace employs 2,400 people around the world and protects over 9,700 organizations globally from known, unknown and novel cyber-threats.

North America: +1 (415) 229 9100 Europe: +44 (0) 1223 394 100 Asia-Pacific: +65 6804 5010 Latin America: +55 11 4949 7696

darktrace.com | info@darktrace.com