Symantec by Broadcom

WHITE PAPER

AI, Automation, and Cybersecurity

Unease and Opportunities: A Broadcom Perspective

WHITE PAPER

AI, Automation, and Cybersecurity

The sudden emergence of ChatGPT caught the world by surprise. Is it an existential threat or the harbinger of the next great scientific revolution? Broadcom offers its perspectives.

TABLE OF CONTENTS

Introduction

The idea of artificial intelligence (AI) has long existed in our collective imagination. In virtually all legends, folklore, and popular culture, it has been considered both a benefit and a curse.

It's no surprise then, that AI is a subject that inspires both wonder and unease. This is particularly relevant now, as the emergence of ChatGPT, Bard, and other generative AI technologies has suddenly changed the expected timeline for widespread implementation.

Is AI an existential threat or a new scientific revolution? Is it good or bad? Is that even a legitimate question?

In this paper, we will dispel myths and discuss opportunities as we look at AI, specifically generative AI. We will focus on the aspects of AI that are most relevant to business organizations today, particularly its implications for cyber attacks, privacy, IP, and copyright issues.

The Broadcom Perspective

Broadcom has been using AI for a very long time. It is integral to our focus and our product solutions for protecting user and enterprise IT. In this white paper we will examine the following positions:

1. Generative AI enables valuable new tools for enterprise organizations.
2. Generative AI does not pose an existential threat if people control its uses and evolution.
3. The widespread adoption of generative AI may well be signaling the start of a fifth industrial revolution.
4. Broadcom cybersecurity solutions ensure that AI remains a force for good.

Let's begin by defining what AI is, and examining its implications for cybersecurity and the enterprise.

AN AI GLOSSARY

Artificial Intelligence
A term defined by Emeritus Stanford Professor John McCarthy in 1955: "The science and engineering of making intelligent machines."

Narrow AI
Intelligent systems for one thing; facial recognition, for example.

Generative AI
A type of AI that generates data based on patterns derived from training sets — images, videos, audio, text, and so on; the most well-known generative AI is ChatGPT.

Large Language Model (LLM)
A type of AI that uses massively large data sets to understand, summarize, generate, and predict new content.

Neural Network (NN)
An AI technique that trains computers to process data in a way like that of the human brain.

Adaptive AI
An emerging form of AI that adapts and improves itself in response to changes in the data or data environment.

What is AI?

Imagine explaining AI to people from the past. How would you do it? One way would be to tell them to imagine having the most intelligent people in the world, all together in one room. This group is so intelligent, they possess all the knowledge that humanity has ever generated; any question could be directed to this group. Having all the smartest people in human history at your command is one way to describe what AI is all about.

Another way is to consider how we think about something as simple as a pizza. Different people will have different ideas of what a pizza should be. Some will say that it should only have a specific shape. Others might say it should only have specific toppings. But, in general, a consensus will probably agree that a pizza is typically round and most often comes with cheese and a sauce. That determination will derive from the collective experience that people over time have had with pizza. The same way people describe pizza, because of what they have learned over time, is how AI works: by learning from experiences over time.

AI operates on neural networks powered by deep learning systems, just like the brain works. These systems are like the processes of human learning, but unlike human learning, the power of crowd-sourced data in machine learning (ML) or AI means that processing answers is a lot faster. What might take 30 years for an individual might take just the blink of an eye for AI.

AI and Automation

One of the most transformational use cases for AI involves how it pairs with automation. Two of the aspects of automation most affected by AI are automated decision making and or content generation.

Automated Decision Making

AI is intended to help simplify human tasks or perform tasks that a person would need to learn how to do. These could be jobs that make life easier, such as repetitive tasks, and can also include creative tasks that AI may perform even better than most people, such as writing a legal brief or explaining complicated technology.

Consider its value for the factory floor. Imagine a decision must be made regarding where to place a particular widget. Or consider a new car rolling off the assembly line. An inspector needs to decide if it has been put together correctly. Is a bolt missing? AI makes these types of decisions better and far faster, and removes the need for people to make these kinds of decisions. This is automated decision making.

Optimizing Decision Making

One of the best and most widespread automation uses cases for AI involves optimizing decision making. Applied to design construction, AI is already much better at determining optimization than human engineers. AI can determine, for instance, where a curve might be needed, or recommend an odd shape that no one has ever considered. The emergence of generative AI may well be signaling the start of a fifth industrial revolution: the beginning of a new era when technology can make better, smarter, and faster decisions on behalf of human beings.

When the Symantec team talks about using AI in our products, we are almost always talking about using it to optimize decision making.

AI IS BEING USED TO INCREASE THE VOLUME AND FREQUENCY OF MALWARE AND SOCIAL MEDIA ATTACKS.

THIS DOES, HOWEVER, PROVIDE A BENEFIT TO CYBERSECURITY PROFESSIONALS: THEY CAN LEVERAGE CHATGPT AND AI TO BETTER DEFEND AGAINST THE ATTACKERS WHO USE THESE SAME TOOLS.

Content Generation

Whether summarizing existing information, writing, creating pictures, videos, music, content generation is the second most important automation use case for AI.

AI and Code Generation

The same way that AI can help generate text, it can also help generate code. This can be dangerous, as it can be used for good or for bad purposes. If AI can write code, that means that a barrier has decreased when it comes to the capability of being able to write code, including malicious, very sophisticated code.

AI is already being used to generate phishing sites. On a global scale, Symantec Threat Researchers are already seeing AI being used to translate phishing or other social media attacks into multiple languages, greatly expanding the target field. ChatGPT works extremely well as a translator, especially translating English into other languages. The program tends to get everything right: the words, the grammar, the syntax, and even colloquial expressions.

Not Dark Yet

AI's ability to help write spam, scam phishing email and social media posts, and malware is increasing the volume and frequency of these cyber attacks. Fortunately, AI is not yet sophisticated enough to make these attacks much different from what already exists. AI enables more variations on existing code samples or malicious code samples, but not more sophisticated versions of these code samples. For a generative AI program to come up with a more sophisticated attack, someone must have already thought about it and published context for it to the Internet.

It is important to note that AI cannot intentionally create a more sophisticated attack strategy on its own. Not yet. While attacks using existing strategies will become more widespread, they will not become more sophisticated. The odds are good that this will come at a later point, once people really know how to extract the right information, condense it faster, and make it more sophisticated.

Enhance Protection Against Cyber Attacks

A benefit of AI being used to increase the volume and frequency of malware and social media attacks is that, conversely, it allows cybersecurity professionals to use ChatGPT or other AI tools to better defend against the same attackers who use them. Enterprises can leverage the capabilities of these tools more broadly within their organizations. It allows enterprises to add to both the expertise of their security operations center (SOC) teams and the speed at which they can respond to potential threats.

AI and automation go together: both good and bad. AI enables the automation of attack chains and the automation defending those same attack chains. It is also a key aspect to adaptive AI, a form of AI that automatically adjusts to changing conditions. Automation makes adaptive AI possible.

HOW IS AI DIFFERENT FROM SEARCH?

AI is different from search in two major ways.

- The first way AI differs from search is in decision making. Imagine planning a cross-country vacation. The goal is to make the trip using the most fuel-efficient route possible. Search engines will offer a map to guide the travelers. They may offer a list of places of interest and could even provide a list of gas prices at stations along the route. AI, however, takes that same information and makes independent decisions. It will decide which route is fastest, where the most inexpensive gas stations are located, and chart the best route for the traveler.
- The second area where AI differs from search is that it always provides a single or summarized answer. Suppose, for example a researcher wants to learn about Henry VIII. Using Google, as of May 2023, the researcher will be offered 105,000,000 results. It's up to the researcher to decide which of those results to follow. AI does the opposite: it gives the researcher a direct answer to a question. Search provides data. It's up to the person doing the search to read and analyze that data. AI summarizes the data, and that summary may be completely wrong or may leave out information that would be important to the analyst or researcher.

Cybersecurity Aspects: Now Is the Time to Learn

When Bad Things Happen to Good People

AI can directly or indirectly affect cybersecurity in a multitude of ways. There are already several examples of ChatGPT exposing corporate data, and that could be just the tip of the iceberg. Generative AI programs are being adopted at a feverish rate. A recent survey reveals that nearly 60% of enterprises polled have purchased or plan to purchase a generative AI tool by the end of 2023.

Bad Information: The Ghost Bug in the Machine

One of the biggest threats that AI poses to business today is that businesses rely on AI for truthful information. The problem with relying on AI is that there is no validation for its accuracy. When someone uses a search engine, the answers it returns provide attribution for their sources. Attribution allows the user to validate the truth or falsehood of data. With AI, there is no attribution because it is correlating information from thousands of different places. There is no current way for enterprises to trust AI-supplied data. The lack of validation is like a ghost bug in the machine system.

The Trust Issue

The fundamental problem with AI is that AI could be providing users incorrect information due to its model and how it was trained or fine-tuned. Given the speed of its adoption and spread across the enterprise, that is the biggest threat that AI poses to businesses today. This may not be the case 5 years from now if it is not accounted for. This incorrect information can create an existential crisis for organizations today.

Looking forward, it is easy to picture a time in which organizations have a tenant of ChatGPT, or some other AI trained by their own dataset for their own use. The questions and answers will be pertinent to their business and their dataset. AI is sure to move toward that future. But today, there is one common dataset for everyone; that is becoming a big problem for enterprises worldwide.

GENERATIVE AI VS. NEURAL NETS VS. LARGE LANGUAGE MODELS

When thinking of the potential benefits of ChatGPT, we should really consider large language models (LLMs) in general, rather than ChatGPT specifically. There are a variety of use cases today where machine learning is used, but where LLMs may prove to be more accurate.

At Symantec, we use ML analytics to correlate disparate and seemingly innocuous events happening across machines in an organization. We leverage this to identify critical breaches, such as situations in which attackers laterally traverse: each action on a single machine isn't enough to raise concern, but together as a whole the behavior is recognized as breach activity. These types of use cases may benefit from LLMs. To understand why, think about each of the attacking events as a word: the single word alone is meaningless, but combine them together and the LLM can 'understand' the meaning of those words in context.

A LLM can also assist is as a knowledge and recommendation source, increasing the efficiency of your security operations center (SOC). Analysts can ask virtually any knowledge question, such as guidance on remediation, or what the LLM knows about a particular IOC. ChatGPT has some of this knowledge today, but our testing discovered that it is not yet sufficient and would require additional training data.

It's easy to imagine, given the right training data, that LLMs could provide strategic recommendations across the board and not just in tactical situations, such as, "Given the following assets or constraints, provide the network security design architecture I should employ." LLMs are still in their infancy, especially in production use, but organizations should expect more adoption in the years ahead.

Cyber Risks Of Generative AI

There are three primary considerations when considering the cyber risks of using AI: attackers, privacy, and IP and copyrights.

Cyber Attackers

For cyber criminals, ChatGPT, GitHub Copilot, and other generative AI solutions do not provide much additional gain. There are some areas where generative AI tools can help launch an attack, for example by constructing better or dynamic phishing emails, or writing code. ChatGPT, like the other generative AI tools, is an information content development tool and not a self-conscious entity. It can be asked to "Tell me all the common ways to infect a machine." But it cannot be asked to "Infect these machines in a way never thought of before."

Malware code or a the text in a phishing email's message body represents only 1% of the entire effort required to execute a successful attack. While AI helps multiply the number of attacks and common implementation in certain areas of an attack chain, it doesn't automate a cyber attacker's end-to-end needs.

An additional consideration is accidental attacks. Code received from a generative AI tool and put into production unvalidated can inadvertently introduce a new attack surface or cause a business disruption.

Privacy, Data Loss, and Risk

When it comes to privacy, the principal consideration is the way ChatGPT can be used by employees or staff. Not only could they be uploading sensitive documents or asking queries that leak sensitive corporate information, but the information and queries they make can also be integrated back into the ChatGPT app.

Employees using AI must guard against inadvertently giving away sensitive corporate information. Providing sensitive information to generative AI programs has the same effect as giving that information away to a third party. Information fed into AI programs, like ChatGPT, becomes part of its pool of knowledge. Any subscriber to ChatGPT has access to that common dataset. This means any data uploaded or asked about can then be replayed, within certain app guardrails, back to other third parties who ask similar questions.

Learning to measure an AI system to protect it against bias and the inadvertent mixing of critical enterprise data is critical. The information the AI system provides must be verified. Misinformation and bias exist across the Internet. When it comes to AI, those factors must be tested and corrected before any of the data received by the AI is used. This is even more important if the enterprise is using ChatGPT for code writing. It is important to understand that AI systems are measurable in terms of their bias and outcomes, and consequently they are predictable to that extent.

IP and Copyright Issues

One of the thorniest issues relating to generative AI is its implications for intellectual property (IP) rights, especially regarding code or creative work using ChatGPT. The issues are not as cut-and-dried as some may assume. If an enterprise developer creates code using ChatGPT, does that become the organization's IP or is it considered the IP of the Internet? What if the code ChatGPT created came from two sources? Is it a 50/50 attribution or an abrogation of another developer's IP?

Enterprises using ChatGPT, or a similar AI solution like GitHub Copilot, to construct text or code need to understand the origination of that text or code may not be copyright free. They could be integrating code into applications they produce that are not licensed properly. They could be republishing text on their websites in or in documentation that is already copyright protected.

Bias, IP, and Data Leaks

Bias is an important attribute of decision making in both AI and our world. Enterprises need help with their AI to identify sources, create guardrails, and in providing additional context to verify the accuracy of their AI information.

The two biggest concerns of organizations today regarding AI involve IP:

- Is my data leaking?
- How do I control what my system is learning from the data?

THE ORIGINALITY QUESTION

The capability of AI to perform creative tasks leads to one of the most contentious of AI-related issues:

Is what AI creates in any sense original, or is it always just a good copy or fake?

The answer is complicated. Although it may sound counter-intuitive, there is a fine line when determining what is and what isn't original. Most AI needs to be trained to learn, just like people are trained and learn over time. Eventually, like people, AI software makes decisions based on its previous experiences. So, is there a difference?

Consider asking an AI program to create a picture that doesn't already exist anywhere in the world today. Let's say it's a cat eating a pizza while sitting on top of a basketball. In a very real sense, that's an original picture. Are any of the picture's constituent elements original? No. But would this picture be considered original if an artist painted it? If one questions whether AI can create something original, isn't it valid to ask if people can do the same?

In 2020, a team created an algorithm that generated 68 billion unique melodies. They then released those melodies into the public domain. The intent was to counteract the increasingly common copyright suits filed by artists in the music industry against other artists for allegedly stealing their melodies to create hit records. Was that effort valid or just another infringement on artistic freedom? The issue remains unsettled because the project itself can be a violation of copyright laws as some of the melodies produced by the AI had already been recorded by other artists.

This question surrounding AI and originality is one that's sure to loom even larger in the years ahead.

Let's look at two traditional use cases, one for DLP and one for compliance, to see the role generative AI plays in each.

Traditional DLP Use Case: Phishing

The insider threat becomes significant with AI. Insiders with intimate knowledge of their enterprise can use ChatGPT to create very realistic email. They can duplicate another's style, typos, everything. Moreover, attackers can also duplicate websites exactly.

Enterprise data loss protection (DLP) solutions like Symantec® DLP can help block these types of attacks. The code is not sophisticated enough right now, but in five years it may be a different story. Among other benefits, DLP can use AI to speed incident prioritization, helping senior analysts triage the most significant threats and recognize those that are not a critical threat to the enterprise.

ChatGPT, Bard, and similar generative AI programs are not as mature a tool as some may assume. The real security problem today is how AI programs improve the quality and the volume of phishing email and websites. Social engineering attacks rely on users being careless or in a hurry to click on a spear phish email link. That is not an issue intrinsic to AI, but more human nature. In terms of security, defending against these kinds of attacks is not really any different than deploying security for Facebook and other social media phishing attacks. The greatest threat is that unlike social media, where an individual posts single posts, with ChatGPT, bad actors can post literally gigs of data at once.

Traditional Compliance Use Case

Generative AI can easily lead to a classic compliance use case when it is used to base decisions on wrong information. To avoid compliance issues, organizations must stress the need for better data hygiene and assess for bias when testing AI solutions. Enterprises will need new resources to manage AI, to handle bias and situations we don't even know about yet.

Benefits for Security Professionals

When we discussed the threat landscape in terms of the increased number of attacks, we mentioned that volume can also be a positive development for security professionals. If SOCs can increase the sample size, they have a bigger set of data to use to tune their own AI, which then allows them to protect against those types of attacks. The combination of greater computing power with larger data sets also allows for faster response to attacks. It is a potential game changer as the Mean time to Response (MTTR) could potentially shrink from weeks to days to even milliseconds.

The true power of these tools will be realized when they are trained in reliable, known-good, and relevant data sets. The Symantec Security Threat and Response (STAR) team has been using this technology and method to generate credible and accurate threat intelligence.

AI DRAMATICALLY CHANGES THE PLAYBOOK FOR INCIDENT PRIORITIZATION: IT CAN HELP SECURITY ANALYSTS PRIORITIZE SECURITY INCIDENTS MORE EFFICIENTLY AND MORE CONSISTENTLY.

Incident Prioritization

AI dramatically changes the playbook for incident prioritization. It can help security analysts prioritize security incidents more efficiently and more consistently. That provides a significant benefit for both threat protection and data protection. When a target breach happens, often the problem is not that the security analysts didn't know the breach happened, but that there were so many incidents, responding to them all was physically impossible.

The same holds true from a DLP perspective: which DLP incident should be prioritized? Does the analyst check the potentially 4,000 incidents that are already blocked, or the one which might have passed through where sensitive data was involved for a use case not already covered? AI can help accelerate that decision-making process and help triage and direct senior analysts to those cases first.

Broadcom And Generative AI

Adaptive and Predictive Capabilities

Broadcom has an enormous amount of data on which we run our AI/ML engines. The Symantec division has been innovating AI for at least a decade. This has provided us with the opportunity to run ML and AI engines on our own data lakes. This supplies us with the great opportunity to become better at identifying threats with both the data we have and the events happening elsewhere. It allows Symantec solutions to identify threats quicker.

ChatGPT is an example of how ML — the basic model behind AI — has essentially taken a lot of the complexity out of security by painting a picture for analysts from many different data points. We can take any behavior and see how it could be exploited by cyber attackers. We can then say, even though there is no evidence of this specific attack, how it could happen, so an attack can be shut down before it happens.

Symantec Advantages

We have a unique approach in using AI and automation to bring data protection and threat protection together. Broadcom combines both data schemes with our Symantec Enterprise Cloud (SEC) security solution, building upon our history of integrating AI in our cybersecurity products and solutions.

Protecting user and enterprise IP is key to our focus. Organizations who already have our DLP solution can feel especially confident when it comes to the threat posed by generative AI systems. It is not true that ChatGPT is

creating malware we can't catch. In this use case, we are not starting from zero. A decade ago, we began using ML with signatures. This experience has positioned us with antivirus protection. We have been in the forefront of innovation with AI ever since. Most recently, this innovation has been leveraged in how we use adaptive protection AI technology to protect enterprise environments from the shift in the threat landscape toward more sophisticated and targeted attacks.

BROADCOM CYBERSECURITY SOLUTIONS WILL ENSURE THAT AI REMAINS A FORCE FOR GOOD. FOR OUR CUSTOMERS, PARTNERS, AND VENDORS WORLDWIDE, KNOW THAT WHEN IT COMES TO AI AND THE FUTURE OF SECURITY, BROADCOM WILL ALWAYS HAVE YOUR BACK.

Conclusions: A Force for Good

Generative AI systems provide a valuable new tool for enterprise organizations:

- Allows security professionals to provide better security.
- Addresses the shortage of skill sets in cybersecurity and other analytics-driven professions.
- Shortens the time to access actionable knowledge by reducing the time to detection, incident response, and remediation.
- Reduces the time for new cyber talent to become productive.

Generative AI does not pose an existential threat if people control its uses and that same evolution:

- It is not intrinsically evil. Don't be afraid of it.
- It will change how we live and work, but we must exert caution.
- We are confident — and see the evidence — that companies are putting guardrails on it moving forward.

The emergence of generative AI may well be signaling the start of a fifth industrial revolution: The beginning of a new era when technology can start making better, smarter, and faster decisions on behalf of human beings.

Broadcom cybersecurity solutions will ensure that AI remains a force for good. For our customers, partners, and vendors worldwide, know that when it comes to AI and the future of security, Broadcom will always have your back.

To learn more, visit symantec.com

BROADCOM connecting everything®

For more information, visit our website at: www.broadcom.com

SEC-AI-WP100 June 27, 2023