






Prepare. Protect. Prosper.

Cybersecurity White Paper

From Star Trek to Cognitive Computing: Machines that understand Security

Second in a series of cybersecurity white papers from Chameleon Integrated Services

In this issue:

-  *The rise of cognitive computing*
-  *Data, Data, Everywhere – how to use it*
-  *Cyber Defense in Depth*

Chameleon
Integrated Services



Executive Summary

Our first white-paper on cybersecurity, “Maskelyne and Morse Code: New Century, Same Threat” discussed the overwhelming nature of today’s cybersecurity threat, where cyberwarfare is the new norm, and we gave some historical perspective on information attacks in general. We also discussed preventative measures based on layered security. In our second in the series we delve into the need for the “super analyst” and the role of Artificial Intelligence (AI) in Cyber Defense and build on the concept of “defense in depth”. Defense in depth, describes an ancient military strategy of multiple layers of defense, which is the concept of multiple redundant defense mechanisms to deal with an overwhelming adversary. The application of AI to cybersecurity is the next frontier and represents another layer of defense from cyber-attack available to us today.

The original television series Star Trek, is a cultural icon that inspired significant elements of the evolution of technologies that we all take for granted today, from the original Motorola flip phone to tablet computers. Dr. Christopher Welty, a computer scientist and an original member of the IBM artificial intelligence group that was instrumental in bringing Watson to life was among those heavily influenced by the fictional technology he saw on Star Trek as a child (source startrek.com article)¹. Conversationally talking with a computer, and have it understand what you are saying seemed a natural direction in the evolution of the computer. While the connection between cybersecurity and machine natural language processing might not be readily apparent, AI and its understanding of language have opened up the world of computer analysis of log files and reports designed to be read and understood by real people. For instance, Threat reports are issued on a daily schedule and are intended to be read by Security Analysts, but the problem is that the information flows too fast for real people to deal with. This scenario, rather than Kirk conversing with the computer is what brings us to the new world of the AI “Super Analyst”.


Swimming in a Sea of Cyber Data

Security Information and Event Management (SIEM) platforms that combine information management and event management into a single platform are the minimum threshold these days for organizations serious about cybersecurity. These platforms were originally intended to take the massive amounts of machine generated log and security data from multiple sources and aggregate it into a single view in order to identify results that differed from normal. Most SIEMs are either rules based or use a statistical correlation engine of the sort that is common in network management tools. Many of these platforms are an outgrowth of network management tooling, which is logical given that network attacks are often the first line of cyber-attack.

As the number of sources that could be monitored increased, the sophistication of the tools available increased in reaction. As threat vectors have become more sophisticated, the information needed to counter that threat has increased exponentially. Organizations responded by increasing staff and apportioning the data across available analysts. Even in a perfect world, with unlimited analysts available, this approach simply is unable to scale as the volumes of data to analyze continues to increase faster than analysts can be trained and integrated. Industry stepped up to the challenge with more sophisticated SIEM platforms, and traditional anti-virus vendors used their respective knowledge of emerging threats (by virtue of receiving reports from millions of endpoints and creating new defenses for the threats as they emerge) to create security threat intelligence services for industry.

Threat Intelligence – the next line of cyber-defense.

Since we can't control the threat, we must look at how we respond to the various threats and the tools that we employ within our defense arsenal. As is the case of Nation States, that have depended for millennia on intelligence and spy craft to predict possible threats, the world of cyber defense has turned to cyber threat intelligence (CTI). This information source has become an increasingly higher priority for businesses and Government organizations, as proven by the US Department of Defense recently elevating US CYBER Command (USCYBERCOM) to a full and independent Unified Combatant Command working with the National Security Agency. CTI is an important tool in our arsenal but getting the most from it can be challenging. A recent survey of threat Intelligence users (source recordedfuture.com)ⁱⁱ showed the top use cases and the results were not very surprising.



Structured or tagged data, the easiest for direct machine ingestion, was the most used, and while some of the other intelligence such as proactively hunting for registry indicators (provided by the intelligence reports) was among the least used, it could have proven to be the most effective in preventing a cyber incident. One of the clear conclusions of the recordedfuture.com SANS webinar behind the data in the table above, was “Relying purely on threat feeds is a recipe for information overload and “alert fatigue.” (source recordedfuture.com)ⁱⁱⁱ Currently most organizations rely on human interpretation of threat intelligence reports which puts a tremendous strain on their analyst resources. Data from reports is unstructured data that has traditionally been of little use to computers. Unstructured data tends to be the result of human intelligence applied in a human readable written format. It can be considered a form of human processed intelligence and in the intelligence spectrum is as archaic as writing itself. With unstructured data, a human brain first processes the data and arrives at insights and then records those insights. A good researcher rarely speaks in absolutes, so instead, they will generally present their confidence levels in the insights reached as a result of their analysis. Fundamentally, this is what cognitive computing does. It arrives at one or more, insights by reading text, examining the work of others, making sense of it, and then proposing conclusions, each accompanied by a confidence level.

Cyber Threat Intelligence Data Sources

While available CTI sources are too numerous to call out here, there are a number that are worth highlighting. The Department of Homeland Security (DHS) has set up a free source, the Automated Indicator Sharing (AIS) website for private companies to share cyber threat indicators with the federal government such as malicious ip addresses and of phishing emails senders addresses that fraudulently induce individuals to disclose sensitive information. The FBI hosts a similar site, the InfraGard Portal. There are also sites dedicated to a particular threat like the Swiss Ransomware Tracker, The international Spamhaus Project, and free antimalware and antivirus identification sites virustotal.com, malwr.com and VirusShare.com. There is also a class of CTI services that are tailored to work best with that providers specific SIEM platform like Alien Vault®, CA technologies, IBM X-Force®, splunk>®, and others. The point here is that there are a lot of CTI data sources available and it can be useful to have help on sorting out the best approach for your organization. As might be discerned from the title of this whitepaper, we favor tools and services that have, or can be integrated with AI to help reduce data and alert fatigue.

We need a “Super Analyst” to make sense of all the data

By now, you probably see the advantage in a machine that can actually understand natural language, like that produced in CTI reports. Just as the fictional Captain Kirk from Star Trek relied on the ships computer to answer his questions about what lies ahead, Cybersecurity officers would like to do the same with their computers today. The problem is that true natural language processing is still in its early stages. It is a far cry from Siri®, Alexa®, or Cortana® translating speech to text, and having that text understood and acted upon. Siri, Alexa®, Cortana®, and Google build complex algorithms to pull out focused meanings from the phrases it understands that are “purpose built”. This is somewhat analogous to the way Electronic Numerical Integrator and Computer (ENIAC), introduced to the world in 1946, revolutionized computing, and how purpose built AI is revolutionizing smart home devices and the Internet of Things (IOT) today. Alexa® learns “skills” that the user chooses in order to get the most out of the platform. Given the low cost of entry for these platforms, they are capable of amazing feats. To go a level deeper, where the AI has a much broader understanding of what constitutes knowledge is like the difference between ENIAC and the first general purpose computer capable of performing the full range of applications for computers from commercial to scientific. IBM introduced the System 360 (model 30) in 1964. It was IBM’s “moonshot” at the time and was the first system to make general purpose computing practical in industry ushering in the “age of the computer”. Interestingly, the System 360 also played a leading role in the Apollo Moonshot program. A Super Analyst that can ingest all the information available to cybersecurity analysts today is beyond what Siri, Alexa or Cortana are capable of now. It will require a machine that not only understands natural language, but can “think” – a cognitive computer.

Coming closer to the super analyst and the Star Trek connection


In 2011 IBM developed a question-answering computer system capable of answering questions posed in natural language in their DeepQA project by a research team led by principal investigator Dr. David Ferrucci (also a Star Trek fan). Watson was named after IBM’s first CEO, industrialist Thomas J. Watson. As a proof of technology, Watson was initially developed to answer

Fun Fact:
 You would need millions if not billions of ENIACs and even then they would run at under a millionth of the iPhone's speed. We have come a very long way since those old computer.
 Source: Quora

questions on the quiz show Jeopardy! and, in 2011, the Watson computer system competed on Jeopardy! against legendary champions Brad Rutter and Ken Jennings winning the first place prize of \$1 million. (source Wikipedia)^{iv} In a PBS documentary on the Jeopardy win by Watson, a number of the principal investigators talked about the influence of Star Trek and of watching Captain Kirk having full blown conversations with the ships computer. They thought that this was a goal worthy of aiming for. Dr. Ferrucci's goal of "talking" to a computer, like Captain Kirk did, and receiving intelligent answers was on its way to realization with the successful

Jeopardy challenge. Winning the Jeopardy! challenge was just the first step on a road of continuous advancement for Watson. Since that time, IBM has applied Watson's learning and interactive capabilities in diverse areas that include everything from cancer research to legal reviews and even more recently tax return preparation and very importantly, cybersecurity.

The Jeopardy challenge was a gamechanger for what most people refer to as AI (which is a very broad category) since it had implications well beyond the demonstration on television. IBM is very particular about how it describes Watson and AI. IBM describes Watson as a Deep Question and Answer machine, machine learning platform, or more commonly as cognitive computing. They use this stylized Venn diagram to illustrate AI technologies. We bring this up because there are so many technologies called AI technologies today, however not all



of them can be applied to Cybersecurity in the way described in this paper. The broad category of AI encompasses Cognitive computing, which in-turn encompasses both Machine Learning and Deep Learning. Given the confusing landscape of what is considered in the definition of AI today, IBM's Head of Research John Kelly describes cognitive computing and what Watson does: Watson's analysis capability requires the computer to understand the meaning of language

"Cognitive computing refers to systems that learn at scale, reason with purpose, and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment. They are made possible by advances in a number of scientific fields over the past half century, and are different in important ways from the information systems that preceded them. Those systems have been deterministic; cognitive systems are probabilistic. They generate not just answers to numerical problems, but hypotheses, reasoned arguments and recommendations about more complex – and meaningful – bodies of data. What's more, cognitive systems can make sense of the 80 percent of the world's data that computer scientists call "unstructured." This enables them to keep pace with the volume, complexity and unpredictability of information and systems in the modern world." (source

based on its usage and context from vast inputs of unstructured data (the kind of data represented within typical intelligence reports). The application of any AI to the problem of Cybersecurity is a game changer, however, the ability for a machine to ingest the massive

amounts of CTI on the internet and apply it to what it sees in the SIEM is truly a force multiplier. When you consider the vast amounts of data on cyber activity available to organizations today, and the new tools available to work with that data, cybersecurity can now move to the next level – proactive protection through predictive analytics and threat hunting.

AI other than Watson in the cybersecurity space

The industry has taken notice of the utility of applying AI to cybersecurity and new partnerships are spinning up on a regular basis. The pairing of AI with existing solutions is geared toward providing additional analytics horsepower to existing technologies, driving greater efficacy, efficiency, and value. This tends to happen in one of two ways. In some cases, machine learning technologies are applied to existing security defenses as helper apps. For example, Bay Dynamics and Symantec have formed a partnership that applies Bay's AI engine behind Symantec DLP to help reduce the noise associated with DLP alerts. Fortscale does similar things by back-ending endpoint detection and response (EDR), identity and access management (IAM), cloud access security brokers (CASB), etc. (source: csoonline.com)^v Other uses of AI in this space involve the analysis of massive amounts of user behavior data to find anomalous behavior that could be indicative of an insider threat. This kind of pattern recognition requires sifting through mountains of structured information to find the needles in the haystack and is exceptionally useful. While we are on the subject of looking for insider threats, let's take a look at another kind of threat hunting and its relevance to cyber defense.

Cyber Defense, threat hunting, and Osama Bin Laden

Threat hunting, is basically searching through huge volumes of data to identify bad actors and threats to an organization's IT infrastructure. The ultimate goal is to prevent attacks before they happen and eliminate or at least to minimize their effects. Threat hunting tools may ingest threat intelligence feeds, vulnerability analysis reports, risk assessments, malware analysis, HR employee records, security event data, system logs, social media feeds and more. While threat hunting leverages tools, much of the work is manually driven by an investigator who researches the answers to questions that have been proactively developed. For instance, the investigator may choose to look for people in a particular organization with access to highly sensitive resources, who have recently expressed a negative sentiment toward the organization, as they could be potential insider threats. There are, of course, limitations to this sort of activity since it involves a great deal of factors to try to detect anomalous activity and then predict future events (source: ibm.com)^{vi}. As you might imagine, this is where the tools used in Big Data and intelligent analytics can help. There is a new class of analytical tools called "insight analytics" that has evolved out of Business Intelligence and Predictive Analytics. Insight Analytics software, like that developed by i2 (later acquired by IBM), can take data from multiple sources and create visualizations to present non-obvious connections and display patterns of behavior that might not otherwise have been clear. One kind of software, called social network analysis, determines the key players in a network and examines the effect that each has within that system – this is the technique that was used to find Osama Bin Laden (source: huffingtonpost.com)^{vii}. A similar technique was used in the capture of Saddam Hussein.


Finding Osama

American intelligence officials discovered the whereabouts of Osama bin Laden by tracking one of his couriers. Information was collected from Guantánamo Bay detainees, who gave

intelligence officers the courier's pseudonym as Abu Ahmed al-Kuwaiti and said that he was a protégé of Khalid Sheikh Mohammed. In 2007, U.S. officials discovered the courier's real name and, in 2009, that he lived in Abbottābad, Pakistan. Using satellite photos and intelligence reports, the CIA inferred the identities of the inhabitants of the mansion. In September 2010, the CIA concluded that the compound was "custom built to hide someone of significance" and that bin Laden's residence was very likely there. Officials surmised that he was living there with his youngest wife. (source: Wikipedia)^{viii} The investigative technique of finding someone based on their connections to others that they know (their "network") is not dissimilar from that needed to connect cyber bad actors and the techniques they use. It can also help the Cybersecurity Information Officer (CISO) identify potential insider threats and possible methods. This is why a "Super Analyst" that can ingest news and social media feeds from around the world and can make use of Insight Analytics, can go a long way toward "predicting" the next threat vector and giving CISOs a fighting chance to protect, before an attack reaches their firewall. Just as insight analysis was used to help find Osama Bin Laden, i2 Enterprise Insight Analysis can be used to perform threat hunting and investigations. IBM's i2 QRadar Offense Investigator integrates the QRadar SIEM capabilities with i2 in order to improve the efficiency and effectiveness of the investigations. Intelligence to hopefully get ahead of the next cyber-attack.

Defense in Depth – an ancient strategy relevant today

Defense in depth is an ancient military strategy first known in history to be used by Hannibal at the battle of Cannae in 216 BC. The idea that multiple layers of defense, when used properly can mitigate the advantage of an attacker with superior numbers is as relevant to cybersecurity as it is to military defense planning. The United States strategic nuclear triad is a defense in depth strategy based on redundancy that has helped make the threat of nuclear war unthinkable to an adversary given the difficulty of knocking out all of our missile fleet (easiest to find), and strategic bomber fleet (harder to find), and finally our nuclear submarine fleet (the hardest to find). This led to a doctrine of mutually assured destruction that leaves no advantage to an attacker that also is assured of being destroyed.



Attributed to:
https://commons.wikimedia.org/wiki/File:Defense_In_Depth_-_Onion_Model.svg

This concept as applied to IT today puts the most valuable resource at the center protected by the security layers associated with the network, then the host computer, then the application and finally the data itself. As in the military example, the more redundant protective layers you can add to the model, the harder it is for an attacker to get to their ultimate goal of your data. Adding in CTI information is like adding a "world" layer outside of your network which essentially adds another protective layer to your defense in depth. Understanding and responding to threats identified in this "world" layer before they have reached your organization by taking proactive protective measures (closing ports, installing patches, educating users, etc.) is obviously preferable to repair actions after the fact.

Watson for Cybersecurity is an instance of IBM's cognitive computing capability Focused specifically on the cybersecurity space. It consumes structured security information from threat intelligence feeds, your organizations' security events and related data, as well as unstructured sources such as research papers, security blogs, websites and advisories. It then stores this as a massive corpus (i.e., knowledge base) consisting of more than 10 billion elements and

refreshes its understanding at the rate of 4 million more elements each hour. In a sense, Watson for Cybersecurity is like a security expert who reads the web 24x7, never forgets, formulates hypotheses about attacks based upon this highly dynamic knowledge base and gets smarter over time. It can function like a “Super Analyst” in your Cyber Defense arsenal and effectively adds another “layer” to your defense.

Adding Intelligence to the depth of your defense

As we have tried to make clear in this paper, there is far too much cybersecurity intelligence data for a single person, or even a team of people, to digest. Statistics at IBM point to the fact that cybersecurity analysts are able to keep up with only 8 percent of newly published information. Without machine assistance, the battle against network intruders will be lost. In our last paper we talked about the fact that we can’t change the threat, only the way we react to it, however, that is not entirely the case. Since cybersecurity is an ever-evolving game of “cat and mouse”, our actions can have an impact on the threat. Every time industry comes up with a counter, the “bad guys” need to make a change on their side. Just as is the case with conventional kinetic threats, raising the sophistication of our defenses, raises the cost of action to our adversaries, altering their behavior. Cognitive computing and AI represent a technological advantage that the average hacker will find difficult to match. State sponsored adversaries have the resources to match our defenses, however, adding AI helps to level the playing field and adding a technology like Watson can provide an edge.

Watson’s Edge

Watson is able to analyze billions of data points gathered from network security analysis programs and correlate it against all known structured and unstructured articles, threat feeds, books, blog posts, and other sources that provide cybersecurity intelligence. What Watson provides is insight into possible threats, and it can do that up to 60 times faster than attempting the same response without it. The report includes a list of possible threats and their ranking based on the likelihood that this is the threat being encountered. The synthesis includes collaboration with human engineers to perform analyses of root causes, particularly those encountered by the organization being attacked, and the vast amount of structured signature data, and lastly, insights gained by analyzing and learning from the cybersecurity corpus of information. The end result is ten times more actionable data than the analyst would otherwise have at his fingertips (source: IBM)^{ix}. Watson helps create the “Super Analyst”.

Summary: AI can transform Cybersecurity

Traditional information processing only sees 20% of the data that is available, and is blind to the 80% of unstructured data that contains processed human intelligence. Cognitive computing is the key to unlocking that information and using it as part of a cybersecurity defense in depth. The simple insight that “wouldn’t it be great if computers really understood us like on Star Trek” has led to a world poised on the edge of new revolution in computing and has begun to unleash the real potential of Artificial Intelligence.

While the science fiction of Star Trek has a long way to go to being realized, by striving to get there, science fact has come much farther than the fiction writers could have imagined.

At Chameleon Integrated Services, we look forward to helping you take your cybersecurity posture to the next level whether that starts with basic services to lock down your infrastructure,

or advisory services to find the best way to increase your defense in depth posture. We hope you have found this paper to be informative, interesting, and thought provoking. Chameleon presently works with our customers to help them “lock down” their infrastructure and comply with STIGS. Moving forward, we are partnering with IBM to help our customers get to the next level of cybersecurity beyond a lockdown mentality.

William J Kapes, Chief Technologist, Chameleon Integrated Services

About




With You from Strategic Vision to Functional Delivery

Enhancing mission effectiveness by reducing cybersecurity risks. Helping you successfully capture transformational opportunities in IT modernization, cloud computing and building the workforce of the twenty-first century. Chameleon is a proven and trusted solutions partner that delivers transformational results and successful mission outcomes to federal agencies, state and local governments, and commercial companies. Our unique approach is built around one set of goals—to help our clients: Prepare. Protect. Prosper.

We're an SBA-certified small disadvantaged business headquartered in St. Louis, Missouri, with offices in Belleville, Illinois; Bowie, Maryland; and Montgomery, Alabama.

Contact us

Chameleon Integrated Services

 <https://www.chameleonis.com>
 (314) 773-7200
 inquiries@chameleonis.com

Endnotes and Bibliography

ⁱ <http://www.startrek.com/article/watson-vs-jeopardy-champs-the-trek-connection>

ⁱⁱ Recordedfuture.com SANS webinar co-hosted with Dave Shackelford.

ⁱⁱⁱ <https://www.recordedfuture.com/threat-intelligence-sources/>

^{iv} Wikipedia, deepQA project by IBM

^v <https://www.csoonline.com/article/3250850/security/artificial-intelligence-and-cybersecurity-the-real-deal.html>

^{vi} IBM Security and Artificial Intelligence FAQ

^{vii} https://www.huffingtonpost.com/2011/05/02/osama-bin-laden-killing_n_856633.html

^{viii} Wikipedia "Manhunt for Osama Bin Laden"

^{ix} Artificial Intelligence and Cybersecurity for Dummies IBM limited edition