

Augmenting Cyber Assessment through Dynamic Malware Analyzer

Ambrose Kam

Lockheed Martin Rotary Mission Systems (RMS)
Moorestown, NJ 08057
ambrose.kam@lmco.com

Michael Nance

Lockheed Martin Rotary Mission Systems (RMS)
Collegeville, PA
michael.h.nance@lmco.com

Wenke Lee

Georgia Institute of Technology
Atlanta, GA
wenke@cc.gatech.edu

Kyuhong Park

Georgia Institute of Technology
Atlanta, GA
kpark 302@gatech.edu

Burak Sahin

Georgia Institute of Technology
Atlanta, GA
buraksahin@gatech.edu

Carter Yagemann

Georgia Institute of Technology
Atlanta, GA
yagemann@gatech.edu

ABSTRACT

Malware is one of the major threats in the cyber space today. It is estimated that over 400,000 new samples¹ are being introduced to the world every day. To defend against these evolving advanced persistent threats (APTs), there is a need for an automatic, scalable, and secure analysis capability. Lockheed Martin is leveraging Georgia Tech's technical expertise in malware behavior analysis capability, and coupling it with cyber simulation to support cyber resiliency analysis and risk assessment.

This paper will describe an approach on how a dynamic malware analysis capability can help design a system with better cyber resiliency against existing and emerging advanced persistent threats (APT). In general, malware will remain dormant until it is triggered by specific computing or network conditions or events. Hence, malware analysis framework, like Georgia Tech's, needs to offer a flexible environment to evaluate a plethora of malware. Most current analysis frameworks cannot segregate the malware traffic in the test network environment; the analysts must decide to isolate, or to let the malware traffic pass through the security controls while letting the malware exhibit their behavior. The dynamic malware analyzer solution from Georgia Tech leverages both static and dynamic code analyses, symbolic execution techniques to examine the interactions between the malware and its command & control (C2) server so that malware can exhibit its intended behaviors in the native environment. By coupling this capability with a cyber simulation, a company like Lockheed Martin's cyber team can develop specific remedies more quickly against the threats given the detailed nature of the malware analysis. Additionally, with a simulation environment, the cyber team can also test and evaluate their "what-if" cyber defensive postures relative to zero-day threats that have not been launched in the real world. Ultimately, this will enhance the system survivability and resiliency.

ABOUT THE AUTHORS

Ambrose Kam is a Fellow in cyber in Lockheed Martin; with over 25 years of experience in the Department of Defense (DoD) industry, he is currently leading teams of engineers and contractors to support cyber threat analysis and resiliency assessment through modeling and simulation. He regularly gives cyber lectures at public forums and conferences including, MIT, Georgia Tech and Military Operations Research Society (MORS). Ambrose received

¹ <https://www.av-test.org/en/statistics/malware/>

his Asian American Engineer of the Year (AAEOY) award in 2017 for his technical contributions, leadership and community services.

Charles Johnson-Bey is the Director of Cyber Innovations in Lockheed Martin Rotary & Mission Systems. He is the recipient of the 2018 Black Engineer of the Year Award (BEYA). He has spent his entire career solving complex engineering and technology problems. He has been a professor and a researcher, earned a doctorate in electrical engineering, worked at Motorola research labs and Corning's Research and Development Center, and is in a 15-year tenure at Lockheed Martin.

Wenke Lee is a Professor and John P. Imlay Jr. Chair in the School of Computer Science, College of Computing, the Georgia Institute of Technology. He is also the Co-Director of the Institute for Information Security & Privacy at Georgia Tech. His research interests include systems and network security, applied cryptography, and machine learning. Most recently, he has focused on botnet detection and malware analysis, security of mobile systems and apps, detection and mitigation of information manipulation on the Internet, and adversarial machine learning.

Carter Yagemann is Ph.D. student in College of Computing at Georgia Tech; he is a computer scientist and cybersecurity researcher. His interests include hacking, system design, and software engineering.

Kyuhong Park is Ph.D. student in College of Computing at Georgia Tech. He studies Systems Security, Operating Systems, Programming Languages; he is also a graduate research assistant at Georgia Tech.

Burak Sahin is a Ph.D. student in School of College of Computing at Georgia Institute of Technology. His research interests are focused on malware analysis, software security and program analysis

Michael Nance is a Senior Fellow for Lockheed Martin. He oversees engineering teams in the classified design, implementation, and testing of trusted ground, airborne, and space-based systems. He has over three decades of experience in senior leadership roles, including as a CISO, focused in Cyber and RF based secure information technologies, and is an expert in computer network infrastructure and cyber space operations. Dr. Nance is a Certified Information System Security Professional (CISSP) and an Information Systems Security Management and Architecture Professional (ISSMP / ISSAP). He received his Doctorate of Science (DSc) in Information Assurance from Capitol Technical University (Capitol College). He is active in STEM and in the information technology community, including with the InfraGard organization and various related councils. Additionally, he is a Board member of several organizations and universities, including serving as an Adjunct Professor on weekends at the University of Maryland University College. You can find him on various social media and broadcast networks and systems as a public speaker for Lockheed Martin especially within the Wounded Warrior community

Augmenting Cyber Assessment through Dynamic Malware Analyzer

Ambrose Kam

**Lockheed Martin Rotary Mission Systems (RMS)
Moorestown, NJ 08057
ambrose.kam@lmco.com**

Michael Nance

**Lockheed Martin Rotary Mission Systems (RMS)
Collegeville, PA
michael.h.nance@lmco.com**

Wenke Lee

**Georgia Institute of Technology
Atlanta, GA
wenke@cc.gatech.edu**

Kyuhong Park

**Georgia Institute of Technology
Atlanta, GA
kpark 302@gatech.edu**

Burak Sahin

**Georgia Institute of Technology
Atlanta, GA
buraksahin@gatech.edu**

Carter Yagemann

**Georgia Institute of Technology
Atlanta, GA
yagemann@gatech.edu**

INTRODUCTION

Cyber defenders are always in a disadvantage position. There are more malware today than remedies to address them at the moment; and this is only to get worse as the rate in which malware is being released simply outpaces the rate of malware addressed. To understand why malware analysts are behind the curve, one needs to recognize a few challenges with respect to the malware analysis process itself. First, the effectiveness of existing malware analysis frameworks relies on specific conditions that need to be met during the analysis phase. That is, malware will only unfold their intended malicious behaviors only when the appropriate triggering conditions (computing or networking environments) are present. For instance, malware may rely on a command and control (C&C) server, which may be dead at the time of analysis; or an advanced persistent threat (APT) may target specific environments, such as industrial control systems (ICS) or Internet of things (IoT) devices, and will not trigger otherwise. In general, it is important to automatically collect and examine triggering conditions with cutting-edge analysis techniques to deal with the high volume of new malware. This automated analysis should assist human experts in precisely and efficiently understanding hidden malware behaviors. Given the current triggering conditions in modern malware, existing analysis frameworks struggle to analyze the hidden behaviors of malware.

The second challenge to malware analysis is the lack of methods that researchers could collaborate and share data. As the number and complexity of malware increase, the cost of manual analysis increases exponentially. Not only manual analysis is time consuming, but it also requires trained experts and the process itself is prone to human error. Although anti-virus (AV) companies release analysis and threat reports, the data they share is either artifacts (e.g., malware sample files) or open-ended reports that analysts cannot directly utilize for further analysis. Even when the reports are publicly released, analysts have to go through the daunting effort of reconstructing the analysis environment before delving further.

Third, existing frameworks suffer from inherent compromises between safety and data richness. For example, executing the malware in a real device without proper confinement (e.g., sandboxing) will fail to contain the attack, while monitoring or emulating the execution environment can trigger anti-analysis countermeasures (e.g., virtualization detection) in the sample. On the network side, the analysis framework cannot simply block the traffic from the malware because it might need network connectivity to unfold its malicious behavior. The analyst must decide whether to isolate or allow traffic to pass through the security perimeter to contain the attack while letting the malware exhibit its behaviors.

To overcome these limitations, a scalable, shareable and secure dynamic malware analysis framework (DMAF) was developed to support a collaborative experimentation. In the following sections, we describe the components and the properties of this framework:

- An automated system to examine and collect triggering conditions for malware using a cutting-edge analysis engine: a) Static and dynamic analysis to understand malware, and b) hybrid (symbolic and concolic execution) to automatically collect and examine triggering conditions.
- A reproducible experiment environment for malware analysis: a) Based on triggering conditions, generating masqueraded C&C server logic, and b) setting up network topology and file system requirements.
- Neutralization of the attacks in the malware using: a) binary re-writing to cut-off the core malicious logic and anti-analysis logic, b) automated and active containment of the network stack, and c) the careful isolation of the lab environment from the Internet.
- Shareable malware analysis environment: a) The entire analysis environment can be shared including not only pcap file, but also system call/API traces, and instruction traces for fine-grained analysis and b) The shared data allows analysts to efficiently resolve issues with others.
- A convenient and accessible interface between the framework and the malware analyst: a) empowering the malware analyst using automated tools, b) boosting the analysis with useful automation tools.

In this white paper, we explore the architecture of this DMAF framework, its modules and explain the benefits of this framework through a malware analysis use case.

ARCHITECTURE OVERVIEW

The DMAF framework provides an environment that transform the malware into an analyzable format so that researchers can concisely analyze the malware, and extract the result data of analysis in a shareable format. By leveraging cutting edge technologies, researchers can now investigate the core logic of the malware to find triggering conditions. With this crucial information, one can replicate a computing environment that makes the malware into thinking that it is communicating with a real Command & Control (C2) server. As a result, the malware's intended behaviors would be exhibited and studied. The Figure below shows how a malware being spoofed in a lab environment that matches the trigger conditions.

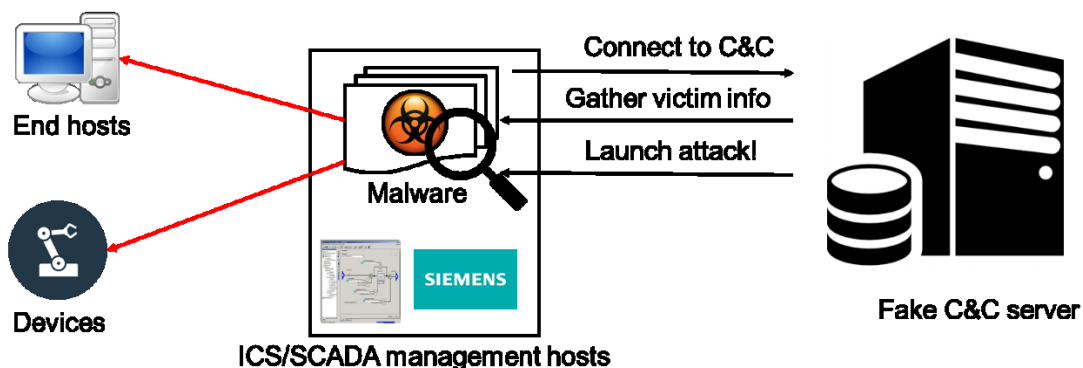


Fig. 1. A Malware Being Spoofed into Communicating with a Fake C&C Server

The DMAF consists of three main phases: preprocessing, analysis, and postprocessing. The preprocessing phase handles detection of features that thwart in-depth analysis of the malware and elimination of the obstacles to transform the given malware into an analyzable format. Also, the fuzzing feeder component automatically searches the network pcap files of the given malware in a malware analysis database, investigates the pcap files to extract C&C related

packets in a best-effort manner. In the analysis phase, each component performs automated analysis to trigger conditions. As such, both static analysis and dynamic analysis generate required outputs to run the automatic symbolic execution with hybrid engines. Each component in this process iteratively interact with each other and examine possible triggering conditions. In the postprocessing phase, a secure runtime environment is constructed by leveraging the known triggering conditions from the previous steps. With this environment, malware researchers can perform repeatable, shareable experiment to concisely investigate malware behaviors, and generates data in network level, file I/O level, and even instruction-level that supported by the latest hardware feature. All data is generated in a shareable format, and the entire experiment process is repeatable—based on the verified trigger conditions. At the end of three phases, the analyst can utilize the outputs from DMAF to build reliable and highly effective protection mechanism for his/her infrastructure with intrusion detection system (IDS)/intrusion protection system (IPS) rules

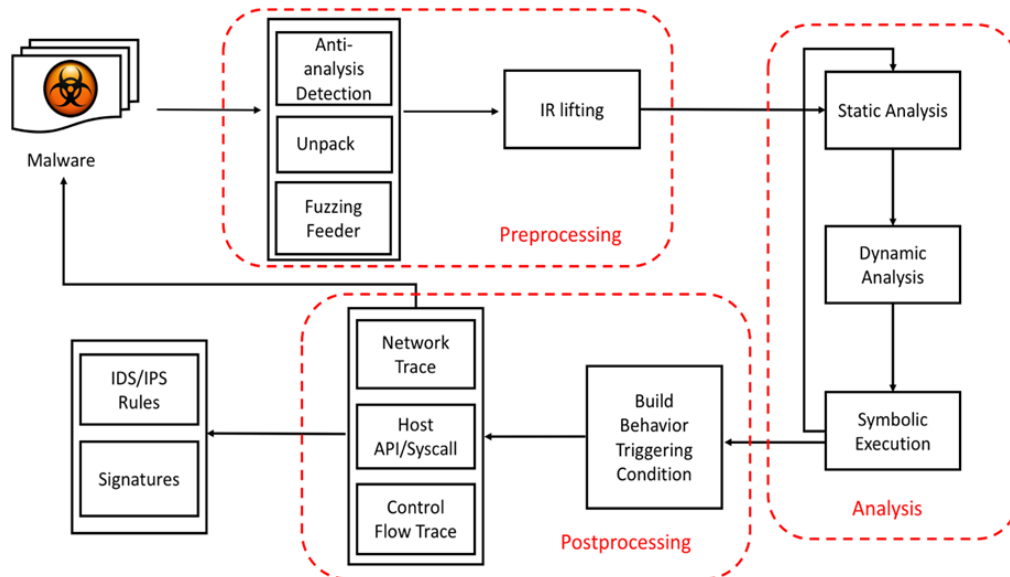


Fig. 2. Dynamic Malware Analysis Framework (DMAF)

Security is paramount in any given cyber lab environment; any possibility of contamination on the core infrastructure should be suppressed and eliminated. This proposed framework provides a compartmented environment to ensure safe handling of malware during an experiment. For example, the virtual hosts such as the virtual machine (VM) that runs malware, the virtual DNS server, and the virtual C&C servers are isolated as separate VMs. Additionally, the network is also isolated per virtual host. An autonomous firewall proxy monitors the network traffic among the virtual hosts and determines whether to allow the traffic to the Internet, block the traffic, or redirect the network traffic to internal virtual hosts to keep the malware live in the runtime environment.

PRE-PROCESSING

Over the past couple of decades, researchers have invented numerous forms of static analysis to identify malicious software. In response, malware developers obfuscate their code to make static analysis difficult and ineffective. Specifically, most malware samples rely on a technique known as *packing* to hide the malware's core functionality.

In the DMAF system, the team employs the industry best practices for “unpacking” malware samples so the code could be revealed. The team runs both specialized and generic “unpackers”; specialized unpackers are tailored made for handling commonly reused packers whereas generic packers tend to be slower and less accurate². On the other hand, a generic “unpacker” typically works on all packers (although less accurate); it is also harder for malware to detect the presence of this type of “unpacker” because the monitoring and code extraction happens in a hypervisor running at a higher privilege level than the malware's sandbox.

² For example, if the packer never unpacks a portion of code, the generic unpacker will not recover it

The other important step in pre-processing is the packet capture (pcap) analysis. Here, the team leverages a large malware corpus to generate input for the fuzzing machine. The goal is for every malware family, to have a reference (which we refer to as “dictionary”) about how the payloads of the network traffic that it receives and send, how the structure looks like, which are the stable parts (remains the same across binaries of the same family, and which are the variable parts (keep changing).

In this DMAF environment, pcaps are run on a daily basis and packets are pre-processed to extract the payloads via BRO and encrypted payload are filtered. The team would need to examine the encrypted payload for the randomness of the payload characters. For those payloads that are deemed valid, the team further processes them, to extract a “dictionary” like input to the Fuzzing Feeder.

The last step of Pre-Processing is Intermediate Representation (IR) lifting. IR is the data structure or code used internally by a compiler or virtual machine to represent source code. IR is appealing for architecture interoperability reasons; a given malware might be native to specific computing architectures, so translating the malware binary code to lower level virtual machine (LLVM) bytecode would be essential. The following figure below is the process of IR lifting for a malware sample.




Fig. 3. Malware Code Translation via IR Process

ANALYSIS

Static analysis is a technique which gathers information about the binary without execution. Once binary code is lifted to an IR, one can use it on different architectures. In this phase, light-weight static analysis methods are applied, and all the complex analyses are handled during concolic and hybrid executions.

First, heuristics are applied to intraprocedural Control Flow Graph (CFG) to determine control flow information and data dependency information. Using that information, it is possible to create a set of pairs of possible start points and target malicious activities. For example, some candidate variables for the command inputs can be the network-related functions where the malware initiates the first dialog by reading an input from C&C server or sending a message to C&C server.

“Concolic” execution is, by definition, a combination of CONcrete and symbolIC executions. Concrete execution generally means running a program on a specific set of inputs, and hence, it is insufficient to rely on concrete execution alone. Symbolic execution allows us to execute a program through all possible execution paths, thus achieving all possible path conditions³. The limitation of symbolic execution is scalability. For a simple program, symbolic execution is often preferred. However, with a large program, the number of possible paths grows exponentially.

In concrete execution, since the inputs are set to specific values, the behavior we observe is valid only for those particular values. Unlike concrete execution, concolic execution would yield a more generalized input space. Concolic execution can represent all the possible runs for a path. On the contrary, concrete execution can only produce exact values for that specific runs. While performing concolic execution, instructions are emulated, and their resulting effects reflect the symbolic environment to imitate concrete execution.

Hybrid analysis involves orchestrating static analysis, concolic execution, and dynamically instrumentation execution. In other words, the hybrid analysis method supports concolic execution while running the binaries in the real

³ A path condition is the set of logical constraints that takes us to a specific point in the execution

environment. Concolic execution engines cannot model the entire system. Also, lifting binary into IR is not always possible. For example, for the packed binaries, we cannot use any static analysis or concolic execution techniques. Even if malware is not packed, the lifted IR may not represent the binary with 100% accuracy. The inaccuracy in lifting affects the accuracy of CFG, instruction, function boundary, etc. Moreover, malware has additional drawbacks on concolic execution. They can be packed, obfuscated or perform cryptographic computations. Modeling the environment and trying to emulate the concrete execution are not enough to overcome the drawbacks of aforementioned symbolic execution.

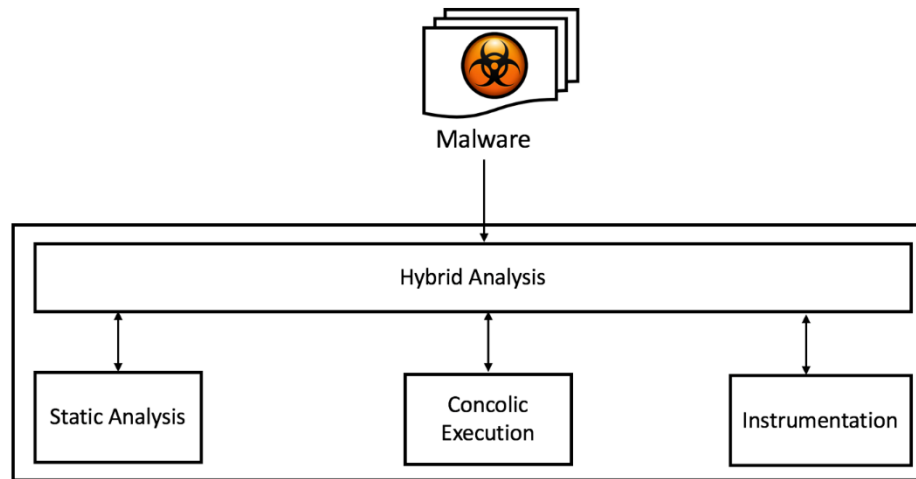


Fig. 4. DMAF Hybrid Analysis Architecture Diagram

While the malware binary is being executed, we should monitor, introspect and even change the information held by the application to manipulate the malware runtime behavior. To achieve this, some of industry Dynamic Binary Instrumentation (DBI) standard tools like DynamoRIO, Intel PIN or Frida are being used to dynamically execute malware and change their behavior. These tools provide the ability to trace instructions, memory, and CPU flags and registers which all have concrete values. As a result, we can get the information via DBI and feed these concrete values into the concolic execution engines, make only the point of interest as symbolic variables such as a network input which could carry the triggering value for a given behavior. This way, the number of the symbolic variables in the concolic execution engine can be reduced. This may reduce the path explosion and constraint solving problems. Another benefit is that we can always run heavy computational functions in concrete execution which will be faster than constraint solving. As shown in the Figure 1, the analysis module includes static analysis, dynamic analysis and symbolic execution iteratively feeding each other.

The above figure shows orchestrating analysis and execution methods. First, for any given malware, the DMAF framework tries to gather static information. Upon detecting the possible malicious paths or any pair of the start location and the target location, the information will be transferred into the concolic execution engine. In the presence of control flow information, the DMAF's concolic execution framework will use directed symbolic execution. Otherwise, the concolic execution engine will follow the path exploration technique.

Second, when the concolic execution techniques find feasible paths and cannot determine what the input would be due to lacking enough constraints, then hybrid analysis will run the malware using dynamic instrumentation tools powered by the possible set of inputs that should be fed to the running malware by the instrumentation tool.

Third, an iterative execution approach is applied to the malware sample when the instrumentation tool cannot proceed further with the information gathered from concolic execution module. In this phase, the concretized memory and the context are transferred into the concrete execution environment to make symbolic space smaller. Hence, one can gain some accuracy and performance due to reduced symbolic space which makes constraint solving faster. Again, whenever the concolic execution module cannot determine a result, we follow the second and third steps iteratively. This iterative process will help both instrumentation and concolic execution. First, concrete execution doesn't have to be performed to increase code coverage. Second, concolic execution will suffer less because of reduced space of symbolic formulae and captured unsupported environment variables/functions during concrete execution.

POST-PROCESSING

The primary objective of the postprocessing phase is to guarantee a secure, repeatable and shareable dynamic malware analysis while not only letting malware run live on the environment, but also feeding the malware sample with triggering conditions that are collected from the analysis phase. After the analyst performs deep dynamic malware analysis, DMAF generates execution traces in network-level and system-level in shareable format.

The postprocessing phase relies on the secure runtime environment using Docker, which is the latest operating-system-level virtualization and containerization technology. With triggering conditions, DMAF generates a fake C&C server, a virtual DNS server and a C&C client that is running the given malware sample. All instances are generated as virtual host by leveraging Docker image and Docker container. Docker system allows an efficient way to build virtual hosts and network topology and to ensure isolated network environment and an isolation layer on the host machine. With the Docker system, DMAF provides the secure and isolated environment aimed for repeatable malware analysis purposes.

The figures below show the results of the host and network level tracing of a given malware sample. The artifacts (including commands/API calls, host/network tracing) produced can be shared with other cyber analysts so the community can better understand a given malware and its behaviors within a computing environment. Armed with this information, cyber analysts can test different mitigation techniques to identify, suppress or eliminate the malware through experimentation. Figure 5 and Figure 6 clearly demonstrate the validity of this process and the value of DMAF.

June 4, 2017, 9:36 p.m. NtCreateMutant	desired_access: 0x001f0001 (STANDARD_RIGHTS_ALL STANDARD_RIGHTS_REQUIRED DELETE READ_CONTROL WRITE_DAC WRITE_OWNER SYNCHRONIZE) mutant_name: WindowsResilienceServiceMutex2 initial_owner: 0 mutant_handle: 0x00000094	SUCCESS	0	0
June 4, 2017, 9:36 p.m. LdrGetDllHandle	module_name: kernel32.dll module_address: 0x7c800000	SUCCESS	0	0
June 4, 2017, 9:36 p.m. LdrGetProcedureAddress	ordinal: 0 function_address: 0x7c8305b1 function_name: IsWow64Process module: kernel32 module_address: 0x7c800000	SUCCESS	0	0
June 4, 2017, 9:36 p.m. CreateProcessInternalW	thread_identifier: 1808 thread_handle: 0x000000ac process_identifier: 1872 current_directory: filepath: C:\WINDOWS\system32\mspaint.exe track: 1 command_line: filepath_r: C:\Windows\System32\mspaint.exe creation_flags: 4 (CREATE_SUSPENDED) inherit_handles: 1 process_handle: 0x000000a8	SUCCESS	1	0

Fig. 5. Host Tracing (Sample Results)

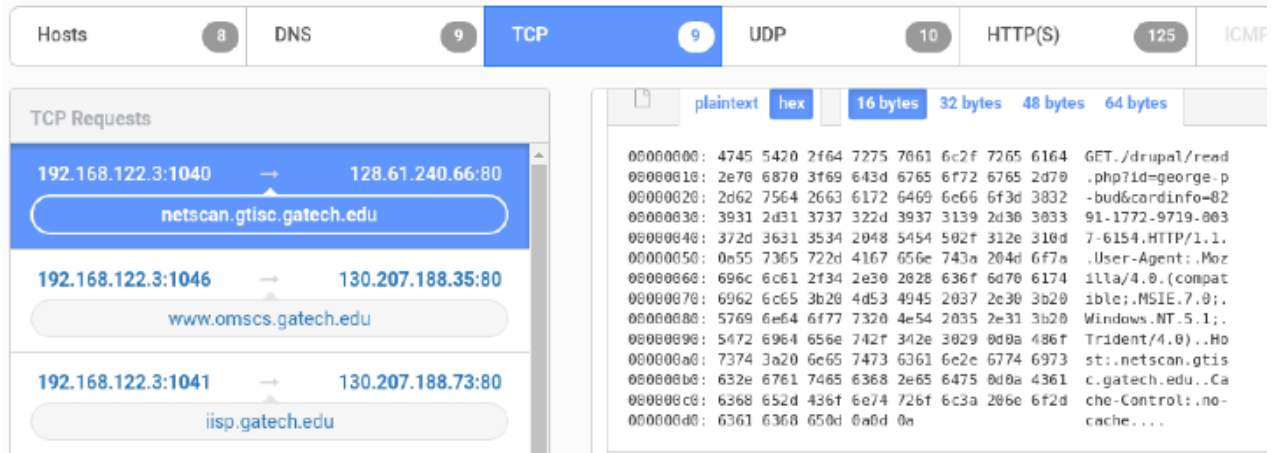


Fig. 6. Network Tracing (Sample Results)

CYBER SIMULATION

One of the main challenges in applying traditional simulation techniques to cyber is threat modeling. In general, it is not hard to model the information system network you are trying to protect. Network topology data and user/application traffic profiles should be available from the network administrators; the adversary and their chosen attack vectors might be a little more difficult to model. However, Red Teaming or Cyber Table Top (CTT) can typically provide some insights, and artifacts can then be fed into a cyber simulation environment. Threat modeling might be the longest pole in the tent since cyber threats, by nature, are very dynamic. This is especially the case for new malware. Zero Day threats are, by definition, new threats that have not been released to the wild, so there is no data on them. As a result, anti-virus and anti-malware solutions are not effective to identify and deter these Zero-Day threats as threat definition database are not populated with their information yet. Without having an analyzer framework, it is difficult to understand what the malware are doing to the information systems; hence, it would be almost impossible to model these new cyber threats and evaluate cyber defensive solutions against them.

The other challenge is security. As a policy, defense contractors are not allowed to introduce real malware to the tactical systems for fear of contamination. Hence, there is considerable difficulties to assess how tactical systems are going to perform in a contested and congested environment where cyber threats are prevalent in the operational environment. Predicting their cyber resiliency would also be near impossible if malware and their impacts are not clearly understood. This drives the need for cyber simulations and their capability to, at least, provide insights on how tactical system can perform in these challenging hostile environments. Cyber Attack Network Simulation (CANS) is an example of such tool. The objective of CANS is to enable the study of cyber events against a model of a planned or operational information system. For this malware analysis case study, CANS is used to simulate the execution and effect of cyber incidents in a modeled network environment, mimicking the behaviors of an intrusion (or an exploit) while demonstrating cyber impacts on the mission system, hardware, and software applications. The adversary effects on mission functionality are represented in the simulation environment, and remediation actions could be tested and evaluated by cyber defenders. Fig 7 below shows a high-level system architecture overview of CANS.

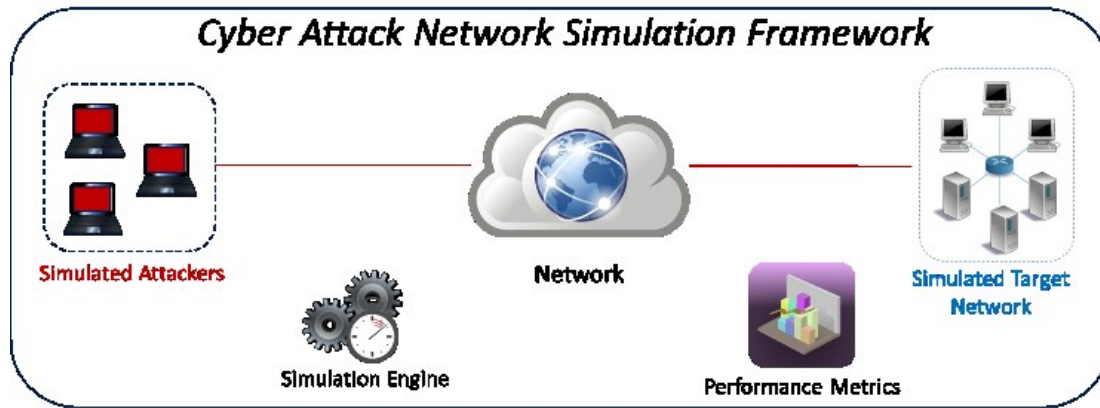



Fig. 7. CANS System Architecture Overview

To support playback and post-processing analysis, a cyber simulation needs to have attack scripting capability. An attack script might be a common vulnerability exposure (CVE) from a vulnerability database, or a malware (or malicious code) downloaded onto a network host node during a phishing attack launched by the cyber attacker. Once the CVEs or malware are defined in the attack script, the cyber simulation can be executed to produce the outcome. To support experimentation, end users can come up with the defensive strategies to mitigate and minimize the cyber risks associated with the chosen attack patterns. Because each scenario run is tied to a random number, a specific iteration out of a set of Monte Carlo runs can be replicated at a later time if the run requires additional investigation.

While modeling a cyber threat through the use of CVE is useful to a cyber simulation like CANS at a high level, it is preferable to model a specific malware with great details. The aforementioned DMAF fills this void; not only a cyber simulation can now model a real-world malware without introducing live malicious source code to a tactical environment, it allows decision makers to assess cyber impacts at a mission level. Operators can be trained in realistic operational environment, so they know how to respond during a cyber event. Cyber defenders can also use this simulation environment to plan ahead, so the overall cyber resiliency level can be improved.

As a case study in 2018, the team leveraged DMAF capability and extracted a new zero-day threat so that cyber analysts could use CANS to come up with a better strategy to mitigate the effects of a malicious malware. At the time, there was a little-known malware called Mnutbot⁴. The team obtained a sample of Mnutbot from a threat database library and started to analyze its behaviors by using DMAF. The malware went through the unpacking and IR Lifting pre-processes outlined in Figure 2. The team then executed both static and dynamic analysis on the malware sample. Concolic execution provided insights on the Mnutbot source code so that the cyber analysts can better understand the triggering conditions. For most malware, a C&C server is typically used to communicate with the malware and send specific commands to be performed. What is unique about Mnutbot is that it does not interact with a C&C server directly. Instead, it relies on a Microsoft SQL server and uses it as a conduit. By wrapping malware commands in seemingly normal SQL traffic makes it hard to detect because anti-virus and anti-malware solutions simply cannot distinguish it from legitimate activities. Figure 8 & 9 below show snapshots of Mnutbot and its SQL server communications captured by the aforementioned pcap and network/host tracing process. Like any other Remote Access Trojan (RAT), Mnutbot needs to receive commands from the server. To do so, it has to query the Microsoft SQL database server for a new command periodically. For example, the highlighted blocks in Figure 8 shows various SQL function calls to be executed on the infected machine. As seen in Figure 9, Mnutbot provides SQL server details—including server address, port, username and a password, all of which are hardcoded inside the sample. By examining the API calls and system commands, DMAF provides a detailed play-by-play description of the malware behaviors through control flow analyses. With this insight, the CAN team replicated the effects of these commands in a simulation environment by mapping them to the common vulnerability exposures (CVEs) and attack patterns. This was all done without the actual Mnutbot malware being deployed to the tactical environment and risk contamination (or worse, violate any information system security policy).

⁴ Discovered by IBM researchers in 2018, Mnutbot was one of the zero-day threats captured and identified.



Armed with this knowledge, the cyber team can now try different IDS/IPS techniques to mitigate the effects of the commands executed by this Mnutbot malware while ensuring the original intended operations can still be performed by the tactical mission systems.

FUTURE WORK

This paper covers current progress in cyber modeling & simulation (M&S) to incorporate effects of the real malware threats. Our minimum viable prototype (MVP) demonstrates the value of this system by integrating and testing different cyber solutions on a potential tactical system environment. Thus far, this process is still very labor intensive as automation cannot be applied to expert system yet. The goal is to leverage Machine Learning (ML) and Deep

Learning (DL) techniques to quickly find the most optimum cyber solutions without this lengthy operations analysis process. Lockheed Martin is currently partnering with Massachusetts Institute of Technology (MIT) and experimenting the use of convolutional neural network (CNN) based Deep Learning techniques. MIT's expertise in ML and DL helps apply these cutting-edge technologies to cyber defenses. Collaboratively, the team has studied Google's TensorFlow and Facebook's PyTorch and determined a set of features that should be used to drive such optimum solution(s) in response to a given set of threats.

SUMMARY

This paper demonstrated how DMAF framework has been utilized to perform malware analysis on a real-world zero-day malware; by leveraging various cutting-edge techniques, the team has unveiled intricate malware behaviors by efficiently analyzing the malware (through static and dynamic code analysis) and its C&C server (through a SQL server). Ultimately, these behaviors are captured and eventually modeled in a cyber simulation environment for cyber risk assessment and system resiliency analysis. The progress described in this report indicates the approach is effective and robust for understanding malware. By coupling this malware analysis capability with a cyber simulation like CANS, cyber analysts can use this capability to assess cyber risks and boost cyber resiliency. Mission planning and cyber operator training activities are also getting benefits from this approach as CMF and CPT can now anticipate certain threat scenario(s) and continue their mission execution while a cyber event is taking place.

REFERENCES

- [1] Barnum, S., *An Introduction to Attack Patterns as a Software Assurance Knowledge Resource*, OMG Software Assurance Workshop 2007
- [2] Bodeau, D., Graubart, R., *Characterizing Effects on the Cyber Adversary: A Vocabulary for Analysis and Assessment* MITRE Technical Report MTR130432, Bedford, MA Nov 2013
- [3] Committee on Technology for Future Naval Forces, *Becoming 21st Century Force, Modeling and Simulation*, Volume 9, U.S. Naval Studies Board National Research Council, 1997
- [4] Department of Defense Instruction (DoDI) 8500.2, *Information Assurance (IA) Implementation* Secretary of the Navy Instruction (SECNAVINST) 5239.3B, *Department of the Navy (DoN) Information Assurance (IA) Policy*
- [5] Jaquith, A., *Security Metrics: Replacing Fear, Uncertainty, and Doubt*. ISBN 9780321349989
- [6] Powell, S., *Cyber Effects Prediction, Black Hat Briefing, DC, 2010*
- [7] Ostermiller, Voula, *TI12 Network Modeling Report for BL9.C2*, Lockheed Martin MST, Nov, 2014
- [8] Security Technical Implementation Guide - Application Security and Development, Defense Information Systems Agency (DISA), 23 January 2014
- [9] Sergio Herrero-Lopez, S., Williams, J., Sanchez, A., *Large-Scale Simulator for Global Data Infrastructure Optimization*, 2011 IEEE International Conference on Cluster Computing, MIT, 2011
- [10] Simmons, C., et. al., *AVOIDIT: A Cyber Attack Taxonomy*, Dept. of Computer Science, University of Memphis, IEEE Mag, Apr, 2010
- [11] Uma, M., Padmavathi, G., *A Survey on Various Cyber Attacks and Their Classification*, International Journal of Network Security, Vol.15, No.5, PP.390-396, Sept. 2013