1.0 CORPORATE INTRODUCTION

Founded in 2001 and headquartered in Rome, New York, Assured Information Security, Inc. (AIS) specializes in high-risk research and development for the Department of Defense and Intelligence Community. AIS expertise spans artificial intelligence, machine learning, biometrics, cyber security, cyber operations, signals intelligence, audio analysis and exploitation, and other areas. AIS plays a leading role supporting the Federal Government in technology development, from fundamental research through sustainment, as well as in rapid R&D in support of urgent operational requirements.

2.0 RELEVANT CORPORATE EXPERIENCE

For more than 22 years, AIS has led cyber operations research and development for users across the DoD and IC. Our experience and technology development span all stages of cyber operations as well as red teaming, vulnerability assessment, challenge curation, and voice-of-the-offense roles. AIS serves as a red team supporting the National Cyber Range and develops cyber operations tools from research through operational deployment. AIS subject matter expertise also includes AI/ML, having contributed fundamental advances to fundamental research in the areas of deep reinforcement learning, transfer learning, explainability, semantic modeling, natural language processing, and other areas. Our experience extends into the application of AI/ML to cyber operations and the behavioral monitoring of users. This includes the current state-of-the-art behavioral biometric user identification technique, user activity detection, passive cognitive load measurement via keystroke and mouse dynamics, and measurement of user trust and suspicion. Our team also leads research in the application of AI/ML to modeling and simulation environments and cyber operations planning technologies. Our scientist and engineers publish findings in leading academic journals (e.g., AAAI, ACL) and present findings within the cyber operations and cyber security research community (e.g., DefCon, BlackHat, CanSecWes).

Consisting of more than 200 research scientists, engineers, and support staff, AIS possess a deep subject matter expertise across AI/ML, biometrics, movement modeling, modeling and simulation, and numerous other areas. AIS also hosts facilities to conduct classified research and development, with support for storage and processing of data up to and including TS/SCI. As a performer on IARPA SCITE and CAUSE, as well as more than 35 relevant DARPA efforts and numerous related AFRL projects, AIS researchers understand high risk, high reward research and development and achieve significant advances to the current state of the art.

3.0 RELATED PAST PERFORMANCE

AIS has employed various technologies to create testbed environments for evaluating user behavior, cognitive load, and cognitive state. The environment is based on a combination of virtual machine introspection and kernel-based monitoring techniques, which are used to monitor system and operation state, introduce effects that induce cognitive state change, and collect data without impacting the test environment. The technologies used to actuate these effects include the Megatron cyber deception framework developed in support of AFRL/RI, the IntroVirt® introspective hypervisor employed on DARPA/I2O Cyber Genome and Active Cyber Defense (ACD), and technology developed under DARPA/I2O ACD, which is used to engage adversary cyber actors to expose their tactics techniques and procedures.

DARPA/I2O Cyber Genome. The technology developed under Cyber Genome leverages and extends concepts from biology and linguistics to support effective malware recognition, defense,

and response. The technology uses a variety of novel techniques to build a comprehensive phylogeny of known malware; it uses that phylogeny to rapidly assess the lineage of new digital artifacts, it uses characteristics of the malware, including its lineage, to characterize the malware attacker; it uses models of attackers and the phylogeny of current malware to predict properties of future attacks; and it provides a graphical interface to support effective use of these novel capabilities.

Under Cyber Genome, AIS served as cyber operations and malware analysis SMEs, as well as the developer of the analysis platform, semantic models, reasoners, and intermediate representation language (IRL) used to capture high-level malware functions expressed in system-level behaviors. The team participated in all phases of the Genome program and transitioned the resulting technology to the intelligence community.

DARPA/I2O Active Cyber Defense (ACD). AIS also leads hypervisor and virtual machine introspection technology development, which ACD employed for the monitoring of users and adversary behavior.

Under ACD, AIS developed both a virtual machine introspection and a kernel-based platform that dynamically contained and monitored advanced persistent threat actor behaviors. The technology seamlessly segregated nodes across a compromised network, creating a logical separation between compromised and uncompromised nodes. The technology then monitored legitimate users, collecting their activities (e.g., authoring documents, interacting with email clients) and, after sanitizing these events from sensitive data, replaying those events within the compromised environment. These interaction, exposed to adversaries, were used to elicit behavior that exposed their intentions (e.g., cyber operations objectives) as well as their tools, tactics, and procedures.

AFOSR Cyber Trust and Suspicion (CTS). Cyber Trust and Suspicion was an AFOSR-funded fundamenal research effort to investigate methods for measuring trust and suspicion in the cyber environment. AIS developed several cyber sensors, including keystroke and mouse dynamics sensors, for remotely and discreetly measuring trust and suspicion in operators. Analyzing data collected at Syracuse University, AIS found a statistically significant correlation between keystroke timings and suspicion. Gaze tracking, application logging, and context logging sensors were also designed to investigate suspicion attribution.

DARPA/I2O Valuation of Information for Covert Collection Computation and Transmission from Offender Red Systems (VIC3TORS). VIC3TORS, a project under a larger DARPA/I2O program, involved the research and development of techniques to track adversary cyber actors across a wide range of devices (e.g., desktop systems, laptops, tablets, mobile platforms) and modalities (e.g., keystrokes, mouse, touchscreen interactions, accelerometer data).

The team analyzed data collected from students and Red Force operators to develop several different types of models: (1) persona, (2) identity, and (3) cyber-relevant actions. The persona models are traditional biometric signatures that distinguish between different users – they're capable of verifying a user, but do not provide information as to his/her identity (e.g., the typing information collected matches User A's signature), whereas the identity models provide information to help link a persona to a real-world identity (e.g., User A is likely a right-handed, native French speaker in the security industry). AIS's work on persona classification led to the creation of the Deep-Vectors framework, a general framework for biometric verification [1]. Finally, the cyber-relevant actions models represent the user's behavior and serves to infer intent (e.g., User A was trying to perform a denial of service attack). When used in conjunction, these

models provide sensors that can help link cyber-attacks to specific individuals or groups and provide the blue-force with a better understanding of their adversaries. Further, AIS developed a reinforcement learning model and simulation environment to enable intelligent C2 that could learn and optimize policies for collection, execution, and transmission towards maximizing efficiency (e.g., transmit as much as possible) while evading detection.

DARPA/I2O Configuration Identification Normalization & Enforcement (ConfINE). Funded under DARPA's ConSec program, the objective of the DARPA ConfINE effort is to infer formal specifications of required functionality for different operational contexts and to generate a configuration-aware model of the system's functionality. ConfINE constructs system-wide configuration settings for reducing attack surfaces and eliminate configuration-based vulnerabilities for each operational context and automatically deploy the settings and will monitor system configuration to ensure its integrity. ConfINE relies on AIS's cyber-operations mission-distribution framework, Metasponse, to deploy secure system configurations and to continuously monitor system integrity. Metasponse is an agentless, extensible, and scalable solution that is currently an integral component in two Air Force weapons systems (ACD and CVA/Hunt), as well as in large-scale cyber exercises such as PacSentry, Arctic Eagle, and CyberShield. The framework, which is at TRL-9, enables cyber defense personnel to rapidly collect, normalize, and analyze data from diverse remote systems, and facilitate enforcing and correcting configuration parameters. In the ConfINE project, AIS built techniques for (semi-) automatically generating Metasponse modules for managing and enforcing secure configurations across complex systems built from diverse devices.

DARPA SMART (MAXWELL). Under MAXWELL, AIS researches and develops modeling and simulation (M&S) technology for rapid and accurate creation of complex, multi-domain mission (i.e., cyber, air, land, sea, space) simulations. The technology addresses limitations in M&S technologies stemming from incomplete knowledge of model properties and the difficulty of identifying models best suited to a desired scenario.

MAXWELL employs reinforcement learning to discover model properties, exercising those models under varying conditions in simulation environments. It formalizes the discovered model properties through an underlying cyber-kinetic domain ontology.

MAXWELL allows users to define a desired mission scenario through free-form, text-based descriptions of desired actions, systems, and properties (e.g., an air interdiction mission to penetrate a region with adversary radar sites). It uses Natural Language Processing (NLP) techniques to interpret the scenario description and the cyber-kinetic domain ontology to formalize its meaning. It then queries the model repository and returns ranked candidate models (e.g., airframe models, radar models) best suited to the simulation requirements.

MAXWELL currently focuses on the AFSIM, ITASE, NGTS, and DARPA SAFE-SiM M&S environments, with the capability to extend to a wide range of other domains and M&S platforms. The completed solution will allow analysts or operators to quickly develop accurate simulations for complex, multi-domain missions.

National Cyber Range (NCR) and NCR II. The objective of this effort is to perform vulnerability assessment and penetration testing of capabilities fielded to support NCR exercises and experiments. AIS provides systems, software, and cyber evaluation engineering support to the NCR for this effort.

The NCR effort provides the opportunity to perform a broad spectrum of testing, with focus on vulnerability assessment and penetration testing of capabilities fielded to support NCR exercises and experiments. AIS is routinely called upon for red-team assessments and to support testing and evaluation throughout the development and operations lifecycle. Under the NCR, red-team and vulnerability assessment efforts, AIS brings architectural and source code audits expertise, as well as assessments and exploitation expertise against operating systems (OSs), networks, embedded systems, and multi-tiered web applications.

Prediction and Analysis of Cyber Scenarios (PACCS). The Prediction and Analysis of Cognition in Cyber Scenarios (PACCS) effort sought to develop passive techniques to measure cognitive load through behavioral biometrics. Under this effort, AIS collected data from 14 employee volunteers and analyzed the resulting EEG and keystroke data. The EEG signal was used to validate the experimental design and label the keystroke data with levels of cognitive load. The model sought to differentiate between four states: control, 4-digits, 5-digits, and 6-digits, where the last three represent the subject attempting to remember a number while typing. The generalized model that applied across all users was not predictive; however, the models trained for individual users were, though there was significant deviation both within and between users. In distinguishing between all users, the predictive accuracy was 0.69 (+/- 0.15), whereas distinguishing between low (control) and high (5- and 6-digits) workload was significantly more accurate and less varied at 0.88 (+/- 0.08).

AIS – Megatron. AIS has developed a mature (TRL-6), Air Force Research Laboratory (AFRL)-funded deception prototype called Megatron, which is built upon experience gained through the DARPA ACD effort. Megatron is an R&D defensive cyber deception capability enabled to manipulate and mislead adversaries, reducing their confidence and increasing their cost.

Megatron is an extensible deception framework for integrating and managing API conforming cyber deceptions, including third-party developed deceptions. Megatron includes deception techniques targeting multiple elements of the cyber-attack kill chain. Network-based deceptions target the reconnaissance and weaponization stages of the kill chain and are used against attackers as they scan networks for targets. These deceptions include erecting fake hosts and running fake services. Host-based deceptions leverage IntroVirt®, AIS's introspective hypervisor, to target the delivery and exploitation stages of the kill chain and to cause effects on hosts infected by the attacker. Existing host-based deceptions include modifying the list of running processes or altering the view of directories and files on the system. Such modifications can hide real files or applications, change the attack surface, or display misinformation to redirect attackers to less critical assets or prompt them to develop ineffective exploits.

AIS – Introspective Hypervisor Framework and Library (IntroVirt®). IntroVirt® (Introspective Virtualization) is an AIS developed technology that provides a library for robust virtual machine introspection using a modified Xen hypervisor. IntroVirt® assumes zero cooperation from the running guest virtual machine and does not require software running in the guest operating system. IntroVirt® is applicable to the areas of behavioral monitoring, user manipulation, system state analysis, task monitoring, and a wide range of other cyber security and cyber operations uses. IntroVirt® provides a robust API and tools can be quickly and easily created to suit the individual introspection needs of the user. As such, IntroVirt® is currently in use in numerous projects and spin-off technologies. IntroVirt® exists as an add-on for virtualization solutions, which include Xen, XenServer, and KVM-based solutions (e.g., ProxMox).

4.0 KEY PERSONNEL

Mr. Michael Sieffert is a Chief Engineer at AIS with extensive experience in cyber security, virtualization, reverse engineering, program analysis, and cyber operations.

Mr. Sieffert has served as a team leader for the Agile Cyber Solutions (ACS) group from 2013 to present, as a team leader of the Systems Analysis and Exploitation (SAE) group from 2010 to 2013, as a senior engineer leading and developing on research and development from 2007 to 2010, and as a junior engineer and researcher from 2002 to 2008. Mr. Sieffert has managed teams of up to 20 researchers, engineers, writers, and developers.

Mr. Sieffert supported the IARPA SCITE and CAUSE programs. He also serves as the principal investigator for numerous DARPA programs. This includes DARPA ACD and ConfINE/ConSec. He also oversees research and development of the Megatron deception framework and the IntroVirt® introspection environment. As part of AFRL's Adversarial Sciences Lab, he contributed to numerous cyber operations technologies, including the Selective Cyber Operations Technology Integration (SCOTI) and led multiple programs (e.g., DARKROOM, ICESTORM) aimed at developing exploits and/or effects for full-spectrum cyber operations.

Mr. Sieffert has more than 20 years of experience in research, development, and leadership in the field of government-focused cyber and information security. In those years other minor roles include technical lead, technical writer, software developer, researcher, reverse engineer, penetration-tester, red-teamer, and proposal lead.

5.0 NEXT STEPS

The AIS team possesses a broad skillset applicable to the ReSCIND program. This includes expertise in cybersecurity, Artificial Intelligence and automated decision making for cyber defense, defensive cyber operations, penetration testing/red teaming, cognitive psychology and behavioral science, cyber attack modeling, software development and integration, and semantic modeling. AIS is interested in identifying collaboration opportunities with partners who bring complementary skillsets, as a subcontractor or a prime. The team can be contacted by reaching out to Mr. Michael Sieffert, Chief Engineer, at sieffertm@ainfosec.com.

6.0 References

[1] J. Baldwin, R. Burnham, A. Meyer, R. Dora and R. Wright, "Beyond Speech: Generalizing D-Vectors for Biometric Verification," in Proceedings of AAAI 2019, Honolulu, HI, 2019.