Enterprise Security Metrics with the ADVISE Meta Model Formalism

Brett Feddersen, Ken Keefe,
William H. Sanders
Information Trust Institute
University of Illinois
Urbana, Illinois, USA
{bfeddrsn, kjkeefe, whs}@illinois.edu

Carol Muehrcke, Donald Parks
Cyber Defense Agency
Wisconsin Rapids, Wisconsin, USA
{cmuehrcke, cparks}@cyberdefenseagency.com

Andrew Crapo, Alfredo Gabaldon,
Ravi Palla
General Electric Global Research
Niskayuna, New York, USA
{crapo, alfredo.gabaldon, palla}@ge.com

Abstract—**Building secure, complex systems is a daunting task. The ADversary VIew Security Evaluation (ADVISE) formalism was designed to offer a model of an adversary attacking a system. As currently implemented in Möbius, ADVISE provides a rich and flexible system security model that, with the other features of Möbius, offers quantitative security metrics. For large systems, constructing realistic ADVISE models can be tedious and impractical. To remedy this issue, we propose the ADVISE meta modeling formalism. An ADVISE meta model is used, with the Möbius framework, to generate ADVISE models and other Möbius components from a higher level model constructed from components, adversaries, and metrics provided by associated Web Ontology Language libraries. This paper briefly reviews Möbius and ADVISE, then introduces the ADVISE meta modeling formalism.**

Keywords - Quantitative Security Analysis; State-based Security Model; Discrete Event Simulation; Adversary Behavior Model

I. INTRODUCTION

Enterprise security metrics are a key component of any system design analysis. For over a decade, the Möbius framework has offered quantitative system performance and reliability metrics on Möbius models defined with formalisms such as Stochastic Activity Networks (SANs) and fault trees. Recently, the ADversary VIew Security Evaluation (ADVISE) formalism was added to the Möbius framework to provide a way to model attacks on a system by a variety of adversaries. With the ADVISE formalism and Möbius, enterprise-level, quantitative, security metrics can be measured on models of existing systems or systems still being designed.

While ADVISE has proven to be a useful approach, building large models can be difficult. Furthermore, when a system changes, reflecting those changes in the representative ADVISE model can be very time intensive. To alleviate this problem, we propose the ADVISE meta modeling formalism. An ADVISE meta model contains a higher level system diagram composed of component objects connected by relationship arcs. An ADVISE meta model also contains a set of adversaries, a set of security metrics, and a set of experimental configurations. We leverage the power of Web Ontology Language ontology descriptions to provide components, relationships, adversaries, and metrics to be used in ADVISE meta models.

The remainder of this extended abstract explains our approach and how we intend to validate it. In Section II, the Möbius framework is introduced. In Section III, we provide an overview of the current ADVISE modeling formalism. The new ADVISE Meta modeling formalism is explained is Section IV. We conclude with information about an Alpha trial we will be conducting in the near future to validate our approach in Section V.

II. THE MÖBIUS FRAMEWORK

The Möbius framework is a mature, extensible modeling and solution framework for discrete event systems. The Möbius tool combines the modeling formalisms and solution methods currently defined in the Möbius framework to offer a user-friendly graphical interface for defining complex system models, useful measures of the system, and a set of experiments. With the Möbius tool, these components are used by analytical solution techniques or the discrete event simulator to find values for the defined metrics.

In the Möbius framework, *atomic models* define the smallest pieces of the system being modeled. Several atomic models, using a mixture of modeling formalisms, can be defined to handle the necessary components of the system. Using one of the *composed modeling* formalisms implemented in Möbius, the defined atomic models, or several instances of an atomic model, can be joined together to build a complete, executable system model. With the *performance variables* formalism, various metrics based on time or system events can be created to quantitatively measure the behavior of the system model. A set of experiments are defined in a *study* to study the impact initial model parameters have on the behavior of the system. Finally, one of the analytical solvers can be used for certain classes of models or the discrete-event simulator for all models to generate results for the defined metrics.

III. THE ADVISE FORMALISM

The ADVISE atomic model formalism is composed of two parts: the attack execution graph and the adversary profile. These parts are necessary to defined the executable behavior of an adversary attacking a system and the effect those attacks have.

A. Attack Execution Graph

An attack execution graph details the attack surface of the modeled system. The graph consists of a set of attack steps, which are atomic actions that an adversary can choose to attempt. Upon completion of an attack step attempt, one of a set of outcomes defined on each attack step is stochastically selected and updates the model state. Model state is defined by the set of access, knowledge, skill, and goal elements defined in the attack execution graph. Each of the state elements represent whether or not the access or knowledge has been obtained, what degree of skill the adversary possesses, and whether or not a goal has been achieved.

B. Adversary Profile

The adversary profile of an ADVISE model details the initial model state, as well as the decision making ability and preferences of an adversary. A subset of access and knowledge elements are selected from the attack execution graph to indicate that the adversary possesses these at the beginning of the model's execution. A subset of skills, with a proficiency level for each skill, is defined to model how effective the adversary is at using these skills. A subset of the goals, with an associated payoff value for each, is also defined in the profile.

Part of the ADVISE method is the evaluation of the attack execution graph by the adversary in order to determine the attack path that will be attempted. This is done with a game-theoretic approach that evaluates the relative attractiveness of all possible attack paths to a defined depth. This depth is called that planning horizon and is an essential part of the adversary profile definition. A low planning horizon will result in a very fast model execution, but may limit the possible goals that can be achieved because the adversary may not explore deep enough to see the payoff from those deep goals. A large planning horizon can yield a slow model execution time.

In order for an adversary to evaluate the attractiveness of attack paths in the attack execution graph, three components of an attack path are considered: risk of detection, cost, and payoff. In the adversary profile, relative preference weights must be defined to model the adversary's interest in or aversion to one of those components. For example, a teenage hacker may not care about the costs of an attack the most and not very much about the risk of detection.

IV. THE ADVISE META FORMALISM

The ADVISE formalism has been used to develop useful models of real systems and is currently being used for system security research by several organizations. However, the approach is not without its drawbacks. Attack execution graphs and defined manually by the modeler and can be quite time consuming for realistic definitions of large systems. Moreover, making changes to complex ADVISE models are also tedious. Also, past users frequently express uncertainty about their modeling decisions in various adversary profile parameters or whether or not their attack execution graph is complete enough to cover all potential attacks on a system.

To address these issues and more, we propose the ADVISE meta modeling formalism. The ADVISE meta model is a higher level model that can be used to generate ADVISE models. The ADVISE meta model consists of a *system diagram*, *adversary profile set*, *metric set*, and a set of *configurations*. From these components, one or more ADVISE models, performance variables models, studies, and discrete event simulators are generated in a Möbius project.

A. System Diagram

The system diagram is a graphical representation of the system being studied. The graph consists of blocks on a canvas that represent components of the system. The blocks are connected by relationship arcs. For example, if a server room is being modeled, an uninterruptable power supply and server may be blocks in the diagram and may be connected by a *poweredBy* relationship. Each component also has a set of attributes that can be defined on the component instance. For example, a server component may have an attribute that defines the number of processors possessed by the server. An ADVISE meta model may contain multiple system diagrams in order to easily study several structural variations of the same system.

B. Adversary Profiles

A set of adversary profiles are defined in a similar way to ADVISE models. These adversary profiles are later paired with a system diagram in the set of configurations. The key difference in an ADVISE meta model is that adversary profiles are no longer defined from scratch, but are instead selected from a structured library of adversaries included with the tool. Once an adversary template is selected, the profile is added to the adversary set and the user can make changes to the profile instance in the meta model.

C. Metrics

Similar to the adversary profile set, system security metrics do not need to be defined from scratch, but are rather selected from a library of metrics. Each metric has required attributed that must be defined by the user. For example, a metric that studies the average time until a specific server is compromised will need to choose a server instance from the system diagram to investigate.

D. Configurations

The final step in defining an ADVISE meta model is the set of configurations. A configuration matches a system diagram, an adversary profile, and a subset of the set of metrics. Each configuration will be used to generate an ADVISE atomic model and a performance variables model (with a performance variable for each of the metrics selected in the configuration).

V. ALPHA TRIAL

As part of our ongoing research, we will be conducting an alpha trial of the ADVISE meta modeling formalism in late September, 2015. We are actively seeking interested participants from academia, government, and industry. Alpha participants will be given access to an alpha version of Möbius with the ADVISE meta formalism, as well as useful documentation for learning the tool. We will conduct regular conference phone calls with participants and work directly with organizations to make their experience a smooth one. We will provide a community exchange with a mailing list and wiki. We are hoping to receive useful feedback from participants to improve on the tool and method.

If you are interested in participating in the alpha trial, please contact us at advise@mobius.illinois.edu.