

COORDINATED SCIENCE LABORATORY


ENGINEERING AT ILLINOIS

ITILLINOIS.EDU

INFORMATION TRUST INSTITUTE

Making Sound Design Decisions Using Quantitative Security Metrics



Bill Sanders



INFORMATION TRUST INSTITUTE

The Problem: Assessing Security and Resilience

- **Systems operate in adversarial environments**
 - Adversaries seek to degrade system operation by affecting the confidentiality, integrity, and/or availability of the system information and services
 - “Resilient” systems aim to meet their ongoing operational objectives despite attack attempts by adversaries
- **System security is not absolute**
 - No real system is perfectly secure
 - Some systems are more secure than others
 - *But which ones are more secure?*
 - *And how much more secure are they?*

Practical Applications of Security Metrics

Organizational-level Metrics

Questions the CIO cannot answer:

- How much risk am I carrying?
- Am I better off now than I was this time last year?
- Am I spending the right amount of money on the right things?
- How do I compare to my peers?
- What risk transfer options do I have?

(From CRA, Four Grand Challenges in Trustworthy Computing, 2003)

A Question neither can answer:

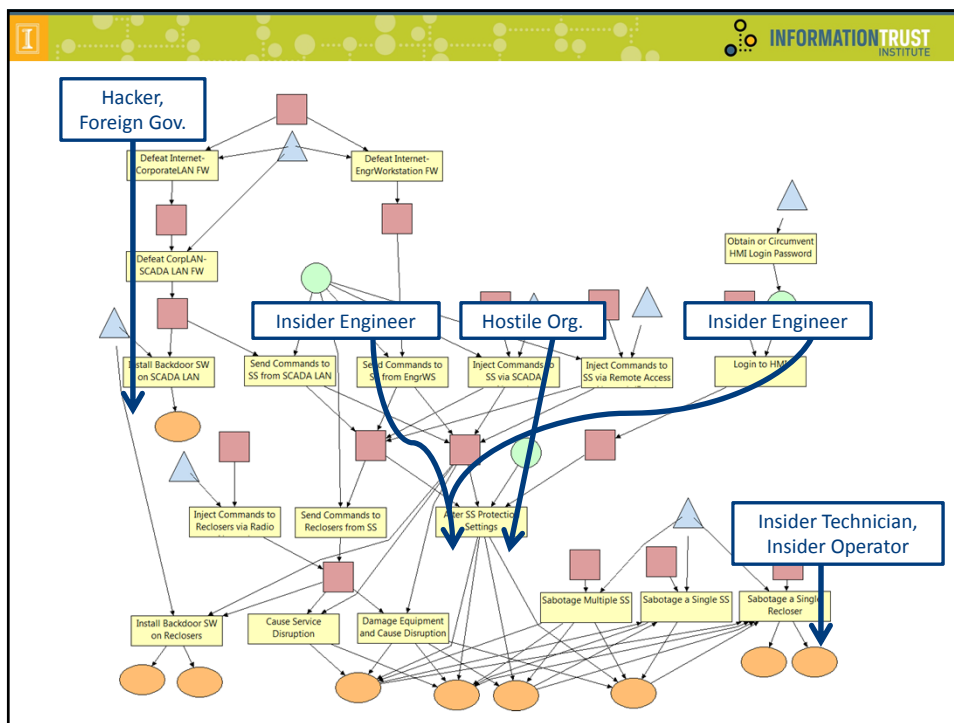
- How do the technical metrics impact the organizational-level security metrics?



Technical Metrics

Questions the design engineer cannot answer:

- Is design A or B more secure (confidentially, integrity, availability, privacy)?
- Have I made the appropriate design trade off between timeliness, security, and cost?
- How will the system, as implemented, respond to a specific attack scenario?
- What is the most critical part of the system to test, from a security point of view?

3







Related Work Motivating ADVISE

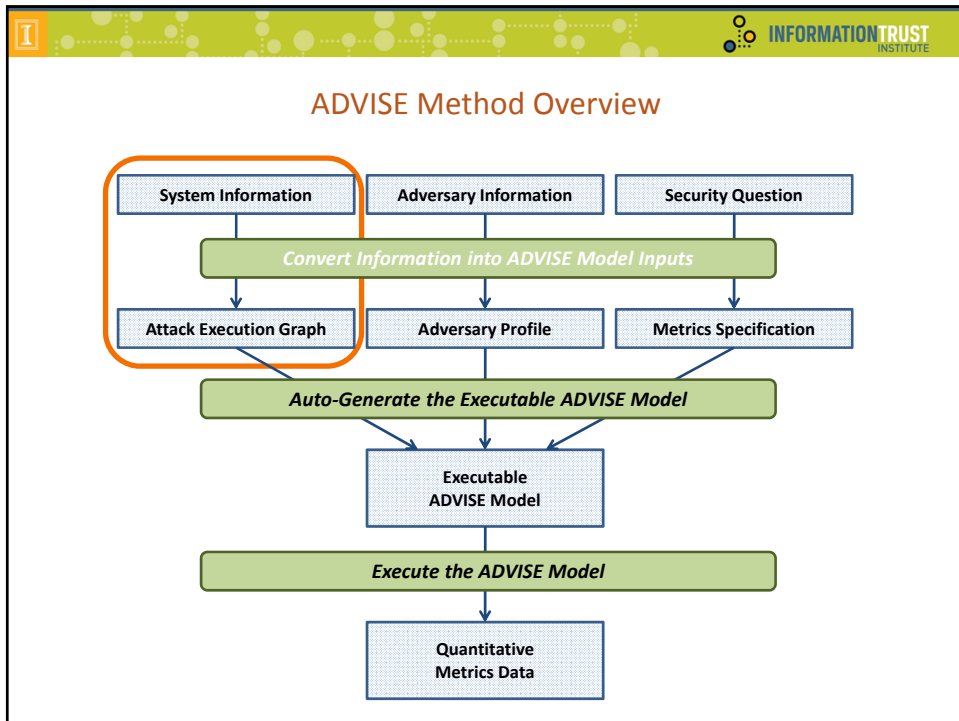

- **Model-based security analysis**
 - Attack Trees
 - Attack Graphs and Privilege Graphs
- **Adversary-based security analysis**
 - MORDA (Mission-Oriented Risk and Design Analysis)
 - NRAT (Network Risk Assessment Tool)


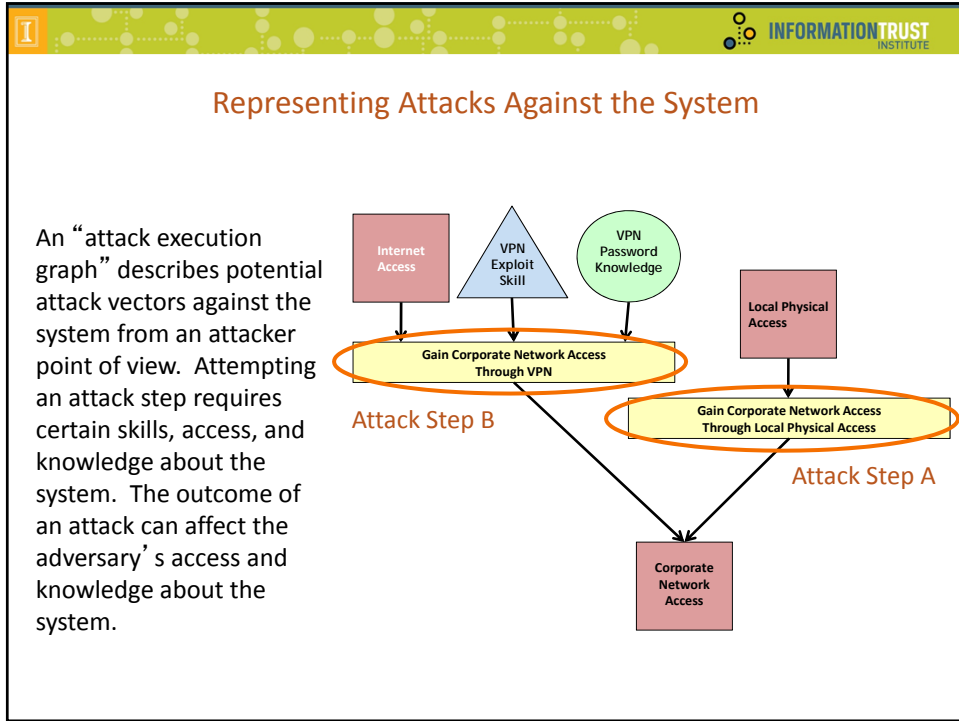
ADVISE integrates the benefits of both model-based and adversary-based security analysis





Adversary View Security Evaluation (ADVISE) approach

- **Adversary-driven analysis**
 - Considers characteristics and capabilities of adversaries
- **State-based analysis**
 - Considers multi-step attacks
- **Quantitative metrics**
 - Enables trade-off comparisons among alternatives
- **Mission-relevant metrics**
 - Measures the aspects of security important to owners/operators of the system





Attack Step Definition

An attack step a_i is a tuple:
 $a_i = \langle B_i, T_i, C_i, O_i, P_i, D_i, E_i \rangle$

Note: X is the set of all states in the model.

$B_i: X \rightarrow \{True, False\}$ is a **Boolean precondition**,
 e.g., (Internet Access) AND ((VPN account info) OR (VPN exploit skill)).

$T_i: X \times R^+ \rightarrow [0, 1]$ is the **distribution of the time to attempt the attack step**,
 e.g., normally distributed with mean 5 hours and variance 1 hour.



$C_i: X \rightarrow R^{20}$ is the **cost of attempting the attack step**, e.g., \$1000.

O_i is a finite set of **outcomes**, e.g., {Success, Failure}.

$P_i: X \times O_i \rightarrow [0, 1]$ is the **probability of outcome $o \in O_i$ occurring**,
 e.g., if (VPN exploit skill > 0.8) {0.9, 0.1} else {0.5, 0.5}.


$D_i: X \times O_i \rightarrow [0, 1]$ is the **probability of the attack being detected when outcome $o \in O_i$ occurs**, e.g., {0.01, 0.2}.

$E_i: X \times O_i \rightarrow X$ is the **next-state that results when outcome $o \in O_i$ occurs**,
 e.g., {gain Network Access, no effect}.

The “Do-Nothing” Attack Step

- Contained in every attack execution graph
- Represents the option of an adversary to refrain from attempting any active attack
 - The precondition $B_{DoNothing}$ is always true.
- For most attack execution graphs,
 - the cost $C_{DoNothing}$ is zero,
 - the detection probability $D_{DoNothing}$ is zero, and
 - the next-state is the same as the current state.
- The existence of the “do-nothing” attack step means that, regardless of the model state, there is always at least one attack step in the attack execution graph whose precondition is satisfied





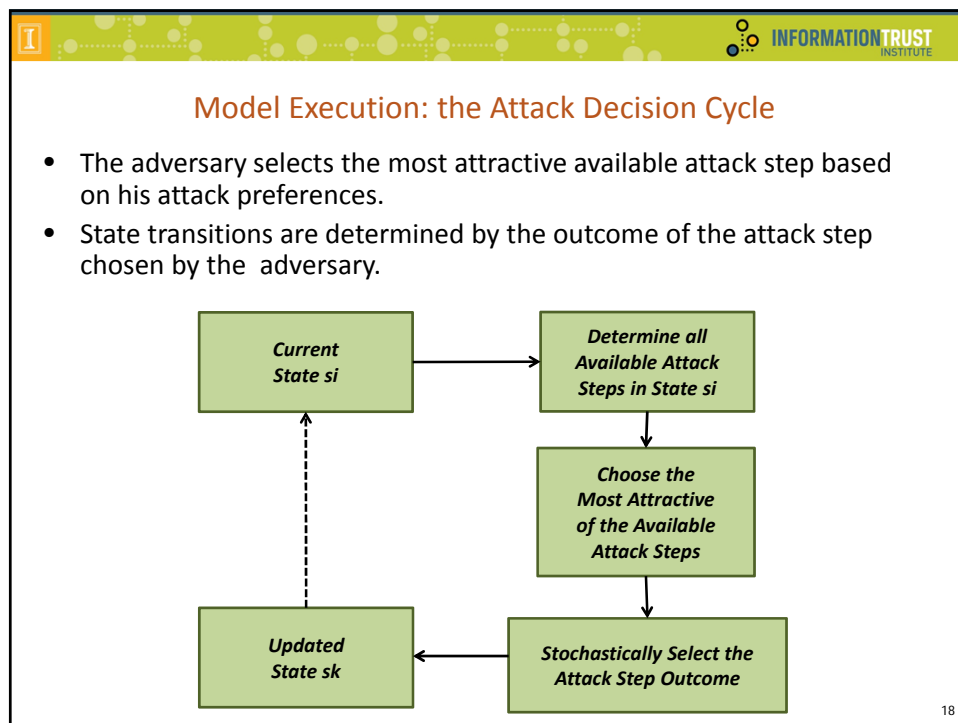
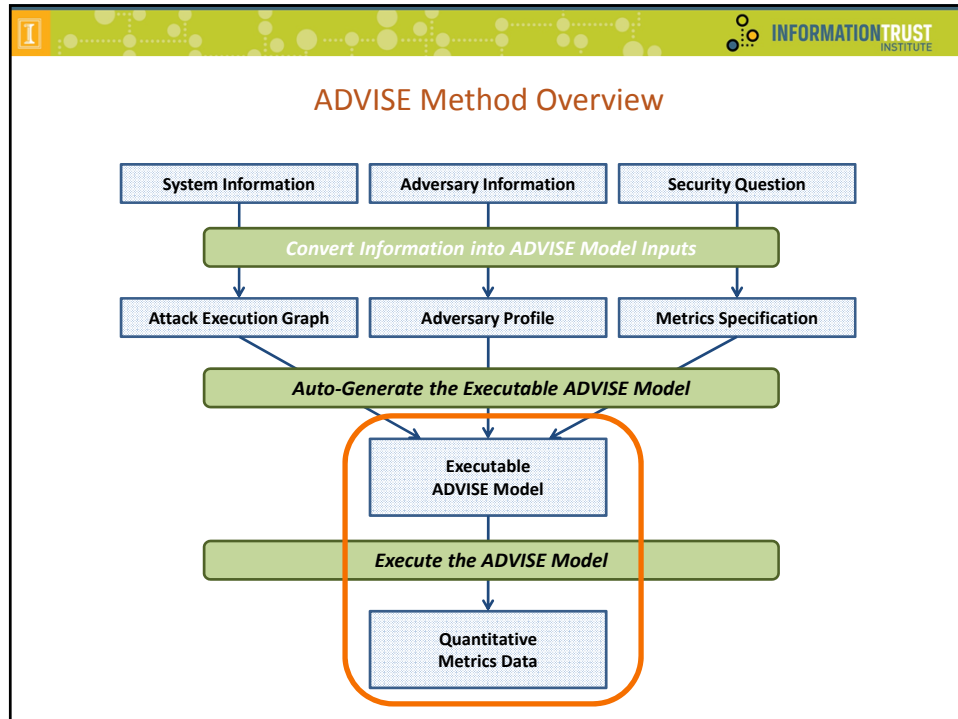
ADVICE Adversary Information: Adversary Profile



The adversary profile is defined by the tuple
 $\langle s_0, L, V, w_C, w_P, w_D, U_C, U_P, U_D, N \rangle$,

where

- $s_0 \in X$ is the **initial model state**, e.g., has Internet Access & VPN password,
- L is the **attack skill level function**, e.g. has VPN exploit skill level = 0.3,
- V is the **attack goal value function**, e.g., values “View contents of network” at \$5000,
- w_C , w_P , and w_D are the **attack preference weights for cost, payoff, and detection probability**, e.g., $w_C = 0.7$, $w_P = 0.2$, and $w_D = 0.1$,
- U_C , U_P and U_D are the **utility functions for cost, payoff, and detection probability**, e.g., $U_C(c) = 1 - c/10000$, $U_P(p) = p/10000$, $U_D(d) = 1 - d$, and
- N is the **planning horizon**, e.g., $N = 4$.







ADVISE Model Execution Algorithm

- 1: Time $\leftarrow 0$ Simulation time and model state initialization
- 2: State $\leftarrow s_0$
- 3: **while** Time < EndTime **do**
- 4: Attack_i $\leftarrow \beta^N(\text{State})$ Adversary attack decision
- 5: Outcome $\leftarrow o$, where $o \sim \text{Prob}_i(\text{State})$ Stochastic outcome
- 6: Time $\leftarrow \text{Time} + t$, where $t \sim T_i(\text{State})$ Time update
- 7: State $\leftarrow E_i(\text{State}, \text{Outcome})$ State update
- 8: **end while**

$\beta^N(s)$ selects the most attractive available attack step
in model state s using a planning horizon of N

Goal-driven Adversary Decision Function

When the planning horizon N is greater than 1, the
attractiveness of an available next step
is a function of
the payoff in the expected states
 N attack steps from the current state
(the **expected horizon payoff**)
and
the expected cost and detection probability
of those N attack steps
(the **expected path cost** and **expected path detection**).

INFORMATION TRUST INSTITUTE

Goal-driven Adversary Decision Function

Attractiveness of an attack step a_i
to an adversary with planning horizon $N =$
 $UC(E[C]) * w_c + UP(E[P]) * w_p + UD(E[D]) * w_d$

$E[C]$ = Expected Path Cost to get to a state N attack steps away via attack step a_i .

$E[P]$ = Expected Horizon Payoff in a state N attack steps away via attack step a_i .

$E[D]$ = Expected Path Detection to get to a state N attack steps away via attack step a_i .

$E[C]$, $E[P]$, and $E[D]$ are computed using a State Look-Ahead Tree.

INFORMATION TRUST INSTITUTE

Consider an adversary attack decision in state s with $N = 1$

Attractiveness of attack step $a_i =$
 $U_C(\text{cost of } a_i) * w_c +$
 $U_P(E[\text{payoff of } a_i]) * w_p +$
 $U_D(E[\text{detection of } a_i]) * w_d$

$C_1 = \$1000$
 $Pr_1(s,1) = 0.9$
 $Pr_1(s,2) = 0.1$
 $D_1(s,1) = 0.01$
 $D_1(s,2) = 0.1$
 Payoff(t) = \$0
 Payoff(s) = \$0

$C_{DN} = \$0$
 $Pr_{DN}(s,1) = 1$
 $D_{DN}(s,1) = 0$
 Payoff(s) = \$0

$Attr(a_{DN}) =$
 $U_C(\$0) * w_c +$
 $U_P(\$0 * 1) * w_p +$
 $U_D(\$0 * 0.9 + \$0 * 0.1) * w_d$
 $= 0.3$
 $= 0.28$

Attr(a_1) = 0.28

Attr(a_{DN}) = 0.3

$\beta^1(s) = a_{DN}$

INFORMATION TRUST INSTITUTE

Consider an adversary attack decision in state s with $N = 2$

Attractiveness of attack step $a_i =$
 $U_C(E[\text{path cost of } a_i]) * w_c +$
 $U_P(E[\text{horizon payoff of } a_i]) * w_p +$
 $U_D(E[\text{path detection of } a_i]) * w_d$

$Attr^2(a_{DN},s) =$
 $U_C(\$0) * w_c +$
 $U_P(\$0) * w_p +$
 $U_D(0) * w_d$
 $= 0.3$

$Attr^1(a_2,t) = 0.85$	$Attr^1(a_1,s) = 0.28$	$Attr^1(a_{DN},t) = 0.3$
$Attr^1(a_{DN},t) = 0.3$	$Attr^1(a_{DN},s) = 0.3$	$Attr^1(a_1,s) = 0.77$



$\beta^2(s) = a_1$

$Attr^2(a_1,s) =$
 $U_C(\$500 * 0.9 + \$0 * 0.1 + \$1000) * w_c +$
 $U_P(\$1000 * 0.8 + \$0 * 0.2) * w_p +$
 $U_D(0.038 * 0.9 + 0.1 * 0.1) * w_d$
 $= 0.77$
 $= 0.85$

INFORMATION TRUST INSTITUTE



Optimality of the Original ADVISE Decision Rule

- **Bellman's Principle of Optimality**
 “an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision”
- The original ADVISE decision rule implements a **provably optimal policy** when the attractiveness function is
 - wholly linear (cost and payoff only) **OR**
 - wholly multiplicative (detection only).
- The original ADVISE decision rule does **not** always produce an optimal decision when the decision rule combines
 - additive rewards (cost and/or payoff) **AND**
 - multiplicative rewards (detection).





Practical Implications of Algorithm Optimality



- Adversaries modeled using this algorithm exhibit “worst case” behavior, that is, they always select a next attack step that is best for them considering
 - Adversary attack preferences
 - Adversary planning horizon
 - Available attack steps
 - Attractiveness function definition



Case Study

- Investigates the effects of architectural changes on the security of an electric power distribution system
- In particular, analyze the security impact of adding radio communication between substations and poletop reclosers





Adversary Profiles: Decision Parameters


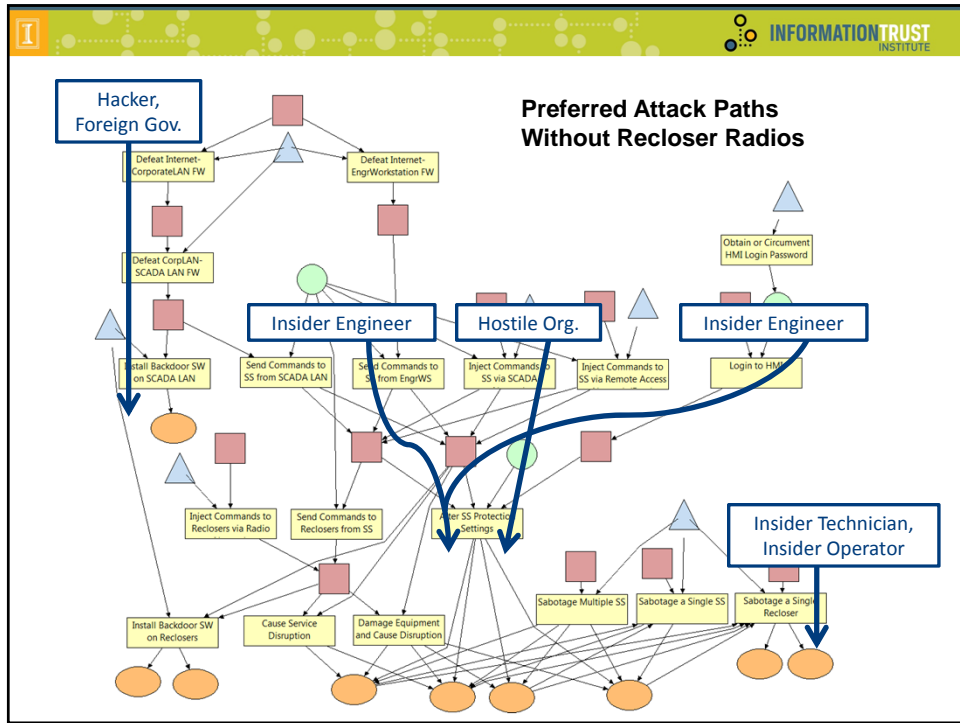
	Foreign Government	Hacker	Hostile Organization	Insider Engineer	Insider SCADA Operator	Insider Remote Technician
Cost Preference Weight	0	0.2	0.05	0.2	0.2	0.2
Detection Preference Weight	0.5	0.4	0.2	0.1	0.1	0.1
Payoff Preference Weight	0.5	0.4	0.75	0.7	0.7	0.7



- The Foreign Government adversary is very well-funded but risk-averse.
- The Hacker is resourced-constrained.
- The Hostile Organization is moderately well-funded and more driven by payoff than the others.
- The Insider Engineer, Insider Technician, and Insider Operator are resource-constrained but willing to take risks.

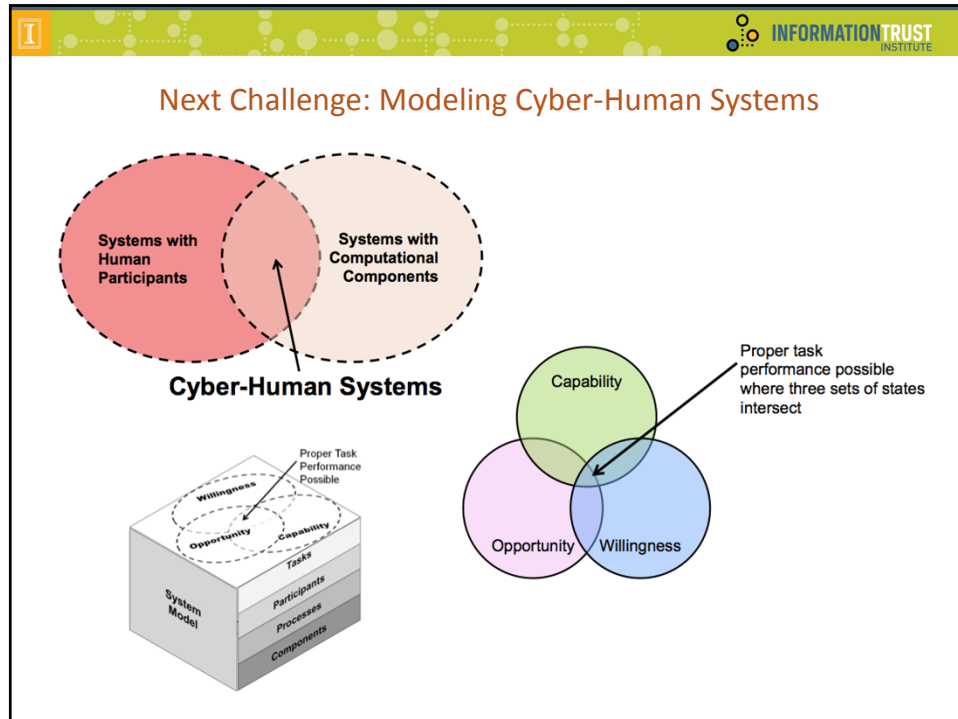





Security Metrics

- Average Number of Attempts**
 - Report for each attack step
 - Gives insight on preferred attack path of adversary
- Probability of Attack Goal Achieved at End Time**
 - Report for each attack goal
 - Gives insight on what goals the adversary is actively pursuing and reaching
- Average Time-To-Achieve-Goal**
 - For attack goals where the above probability metric is 1 (or close to 1)
 - Gives insight on the speed of the adversary's attack











Conclusions

- Since system security cannot be absolute, quantifiable security metrics are needed
- Metrics are useful even if not perfect; e.g., relative metrics can aid in critical design decisions
- The ADVISE formalism, and its implementation in Mobius-SE
 - Is rich enough to adversary, user, and system behavior
 - Natural for security analysts
 - Semantically precise
- Mobius-SE is in alpha-test, and has been distributed to 10 organizations (industry, govt., & academics) who are using it in real case studies
- Work is on going on modeling human user behavior



Thank you!

Bill Sanders
perform.csl.illinois.edu
whs@illinois.edu

38