Practical Challenges and Opportunities in Cloud Security Theory: Semantics, Humans, Metrics

Greg Shannon@cert.org

March 27th, 2014

Theories for Cyber Security?

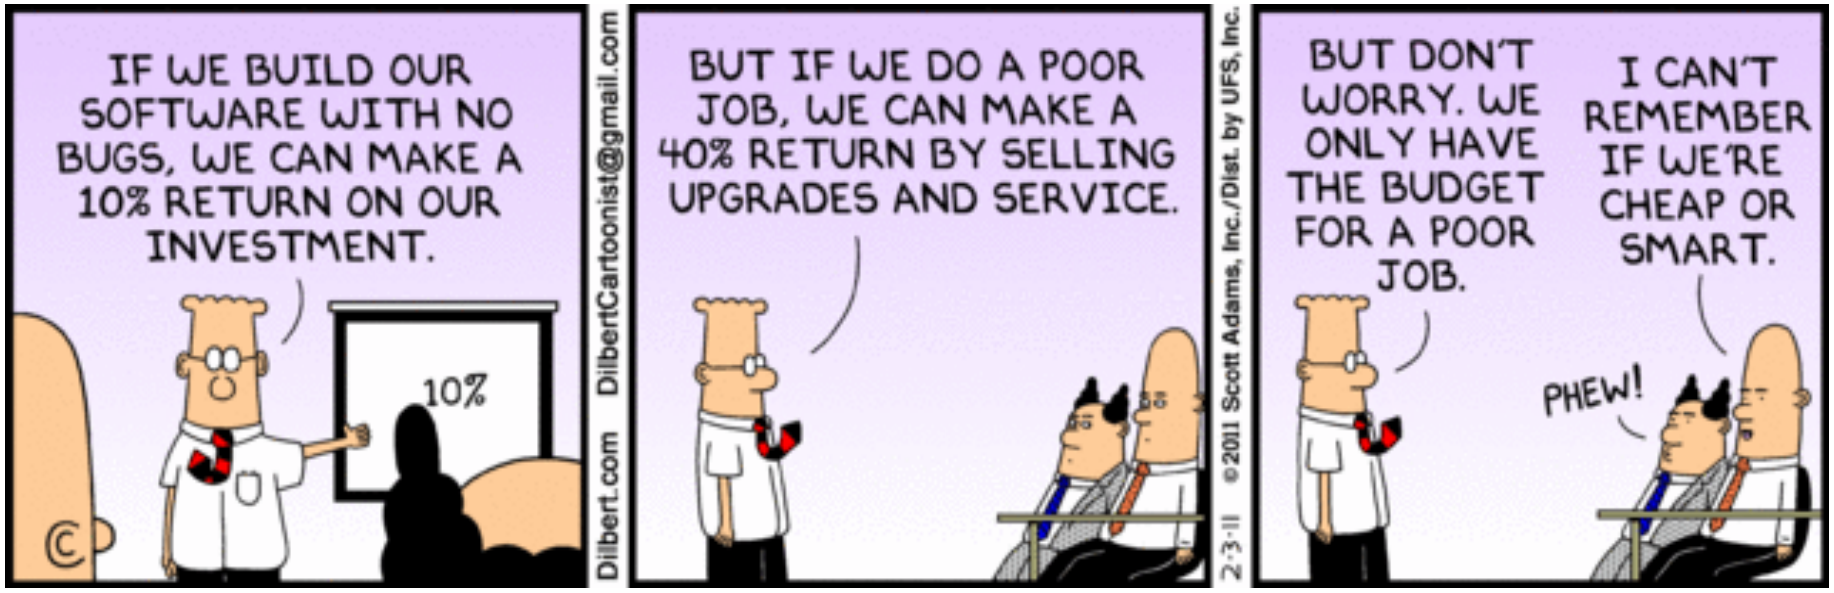

http://www.dilbert.com/strips/2011-02-03/

CERT® Division at Carnegie Mellon

- **Created by DARPA in response to the Morris Worm (1988)**
  - cert.org, ~300 technical staff, ~10% research
  - Part of the Software Engineering Institute, a Defense (DoD) FFRDC operated by CMU
- **Mission: anticipating and solving national cyber security challenges**
- **CERT focuses on**
  - **Customers** with *Pain in Cyber Security*
  - **Data** collected
  - **Trust** as a 3rd party for gov't, law enforcement, industry, academia
  - **Operational experience and capabilities** for cyber-security failures at scale

- **What are we learning from experience?**
  - Failures – many
  - Successes – few?
- **If what we're learning isn't helping, why?**
- **What are the opportunities for theory?**

Delusions (2012): Observations & Questions

- **Exploits continue to rely on known avoidable programming mistakes.**
  - If code **correctness** is improving, why?
- **Unimplemented or ineffective policies continue to be an enabling element of major incidents.**
  - Given all of the information assurance **policies**, why?
- **Many significant intrusions remain undetected for weeks, months, years.**
  - Given all of the **monitoring** and auditing technologies, why?
- **Even sophisticated victims are challenged to quickly and effectively investigate, mitigate and recover from attacks.**
  - If proficient **response** capabilities exist, why?

Some Observations on Cloud Security

- **Yet another abstraction layer?**
  - Layers are highly valuable
  - But not so much for security
  - There will be more layers...

- **People matter?**
  - Always have, always will
  - Too complicated, best to ignore them?
  - Empirical models of population entropy?

Some Examples

- **Assumptions: If we assert P given assumptions A**
  - What are the implications when an assumption changes?
  - Which assumption are required for P to hold?
  - Which set of assumptions must be broken for ¬P to hold?
  - Example: value(@X)+1 > value(@X)? AKA, integer wrap

- **Semantic Gaps: Optimizing Compilers**
  - What is a language's definition v. what a programmer believes it is?
    - http://www.kb.cert.org/vuls/id/162289
  - Confounded by multiple compilers and hardware platforms

- **Abstraction Gaps: Metadata computing**
  - Example: Malware outside of the instruction trace
    - Julian Bangert, Sergey Bratus, Rebecca Shapiro, and Sean W. Smith. 2013. The page-fault weird machine: lessons in instruction-less computation. In Proceedings of the 7th USENIX conference on Offensive Technologies (WOOT'13). USENIX Association, Berkeley, CA, USA, 13-13.

The Real Opportunity

- **Energy == Security?**
- **Example**
  - Fuzzing combined with formal methods for symbolic execution now suggest that finding fully exploitable vulnerabilities is like mining bitcoins.
  - E.g., 3 CPU years → 1 binary for a proven exploit
  - The most power + efficiency "wins"?
  - Continuous searching!
- **More Generally**
  - Computational investment in applying formal methods to identify threats early/often

- **Semantics: Sophisticated threats leverage the inherent semantic gaps between levels of abstraction as well as abstractions and actual artifacts. Semantic gaps ensure that an abstraction can't reason about, describe, or mitigate threats at lower layers enabled by the completeness of an artifact at the higher layer (e.g., a program on your computer). Furthermore, any(!) assumption made about a system, theory, artifact, etc. serves as a possible point of attack. Recent research has sought to "tighten" abstractions so as to ensure that abstractions don't ignore "undocumented functionality" in lower layers. However, further research is needed in minimizing assumptions (i.e. axiom narrowing) as well as in formally representing semantic gaps.**

- **Humans: While a cloud is an engineering artifact, humans are present throughout design, implementation, installation, operation, as users, as owners, as attackers, etc. Furthermore, humans aren't logical nor are they fully predictable. Humans behave as distributions with peculiar inference systems (e.g., decisions made in one second are fundamentally different from those made in ten seconds or ten days). A theory of how to account for the role of humans in homogeneous yet fractal artifacts, like a cloud, would enhance how we reason about and mitigate the security implications of human frailties and malfeasance.**

- **Metrics: A securon? A threaton? What are fundamental units of abstraction on which a science of security might build useful metrics? The practical impact of the lack of well-founded security metrics makes security investment decisions impossible. Can enforced homogeneity in a cloud artifact enable a different kind of metric and measurement?**

Explorations of Science in Cyber Security

Dr. Greg Shannon@cert.org
Chief Scientist
October, 2012
+1 (412) 268-8545
www.sei.cmu.edu/about/people/shannon.cfm

My Science of Cyber Security Rabbit Hole

- Alice wants to consider ideas from Greg for cyber security research (October, 2006)
- Bob works for Alice
- Greg gives Bob a digital media artifact with ideas (a compact disk with power point files)
- Bob says, “I should scan this for viruses,” and then uses the files without scanning them
- Greg realizes he just fell down a “rabbit hole”

CERT® Program at Carnegie Mellon

- Created by DARPA in response to the Morris Worm (1988)
  - cert.org, ~300 technical staff, ~10% research
  - Part of the Software Engineering Institute, an FFRDC operated by CMU
- Mission: anticipating and solving national cyber security challenges
- CERT brings to the table two decades of experience with security failures
  - **Customers** with *Pain in Cyber Security*
  - **Data** collected
  - **Trust** as a 3rd party for gov't, law enforcement, industry, academia
  - **Operational experience and capabilities** for previously *unexperienced* cyber-security failures at scale

A Science of Cyber Security?

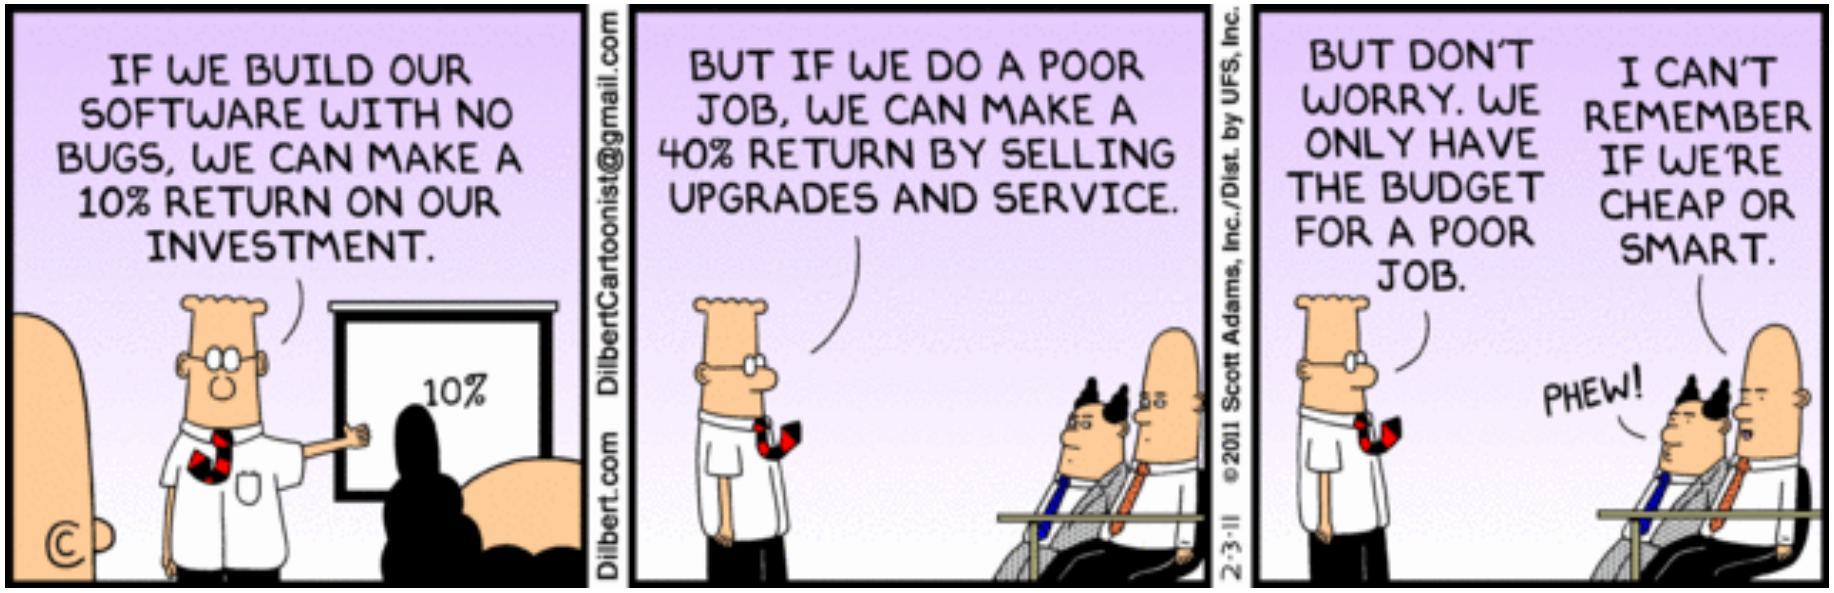

http://www.dilbert.com/strips/2011-02-03/

Security v. “InstaGram”?

From: Greg Shannon <shannon@cert.org>
Subject: **Houston, we have a (secure) coding problem?**
Date: April 10, 2012 8:04:59 PM EDT
To: recertch <recertch@cert.org>

http://thenextweb.com/2012/04/10/instagrams-ceo-had-no-formal-programming-training-hes-a-marketer-who-learned-to-code-by-night/

http://tnw.co/HVZ96z

So, the Instagram founder created a billion dollar company
(purchased by Facebook this week)
by teaching himself to code on CodeAcademy.
(Plus, he had some great market-savey ideas.)

As far as I can tell, CodeAcademy has little (if any)
guidance or lessons on secure coding.

Does anyone know to what extent that site discusses security?
Or similar such sites?

Delusions: Observations and Questions

- Exploits continue to rely on known avoidable programming mistakes.
  - If code **correctness** is improving, why?
- Unimplemented or ineffective policies continue to be an enabling element of major incidents.
  - Given all of the information assurance **policies**, why?
- Many significant intrusions remain undetected for weeks, months, years.
  - Given all of the **monitoring** and auditing technologies, why?
- Even sophisticated victims are challenged to quickly and effectively investigate, mitigate and recover from attacks.
  - If proficient **response** capabilities exist, why?

The Scientific Method (Simplified)

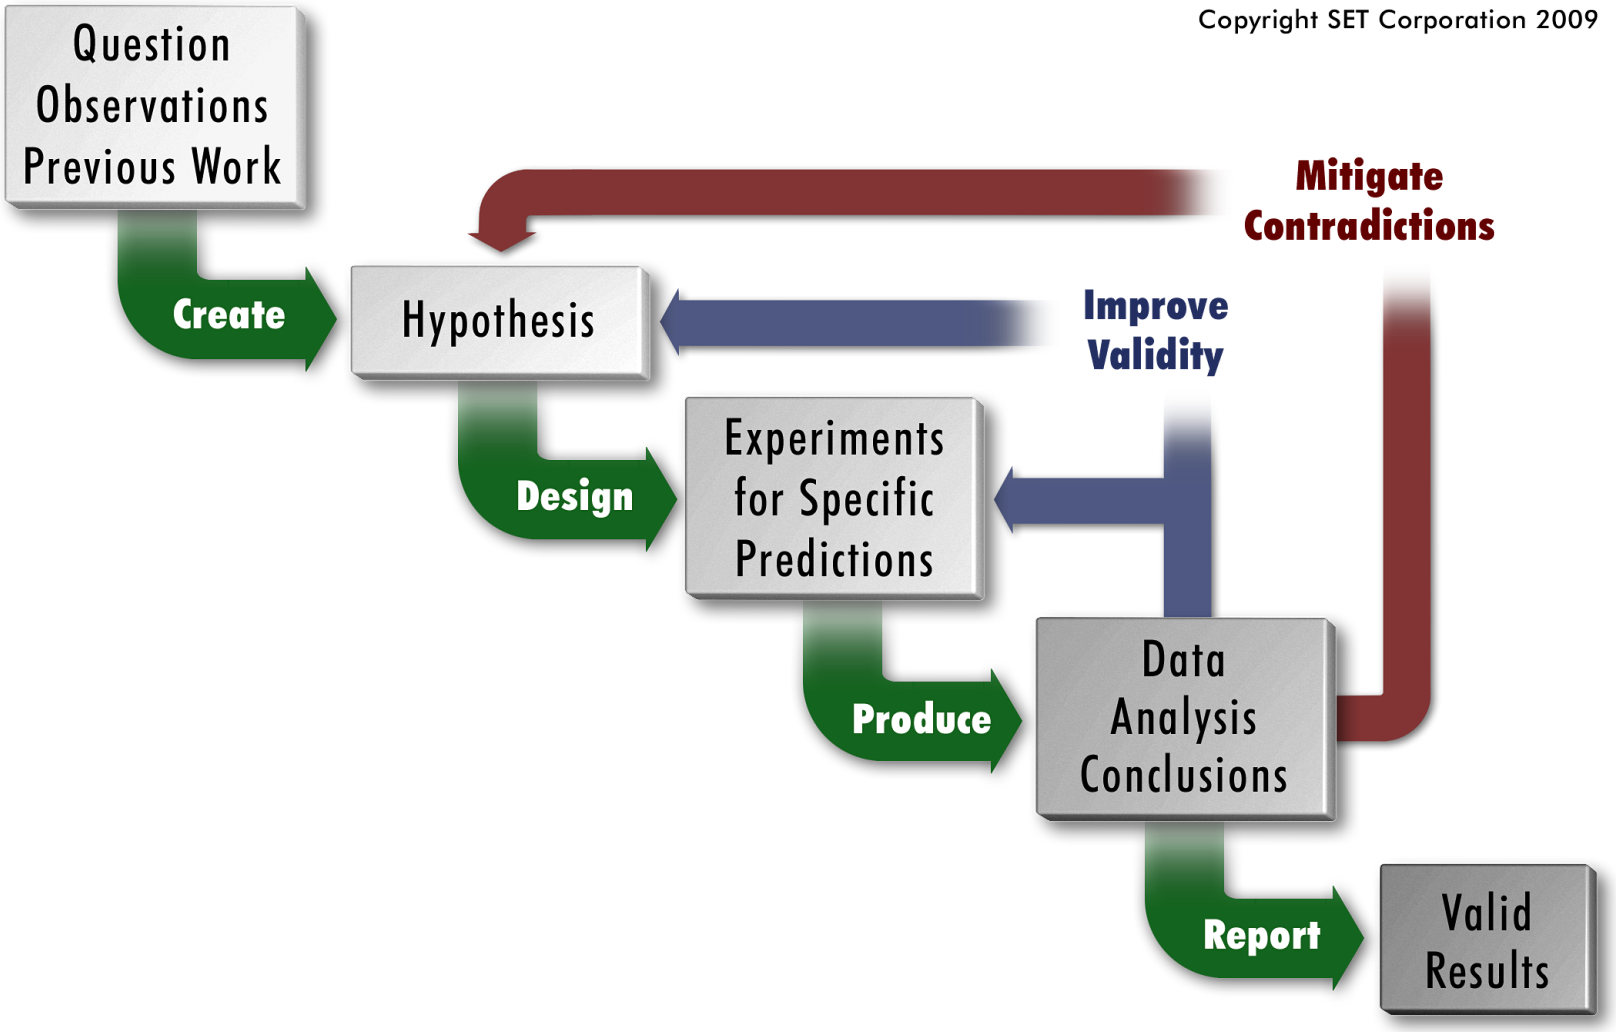

Long Paths to Scientific Understanding

- Health ~ Religiosity →→→ Health ~ Hygiene
  - Required an understanding of the underlying phenomenologies that degrade health as opposed to the causes of health per se.

- Bloodletting
  - Widely accepted treatment in 1800 for fever, swelling.
  - "Medical statistics" led to better treatments

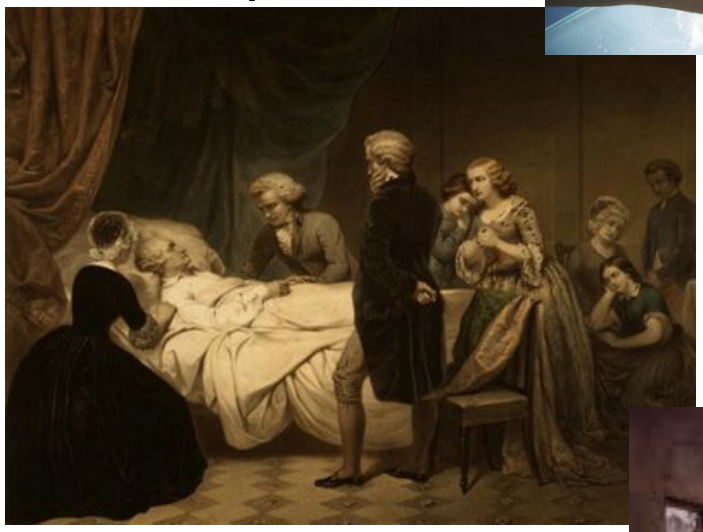

- Alchemy
  - Broad support; fervently practiced by Newton
  - Eventually overcome by modern chemistry

CERT Research Landscape

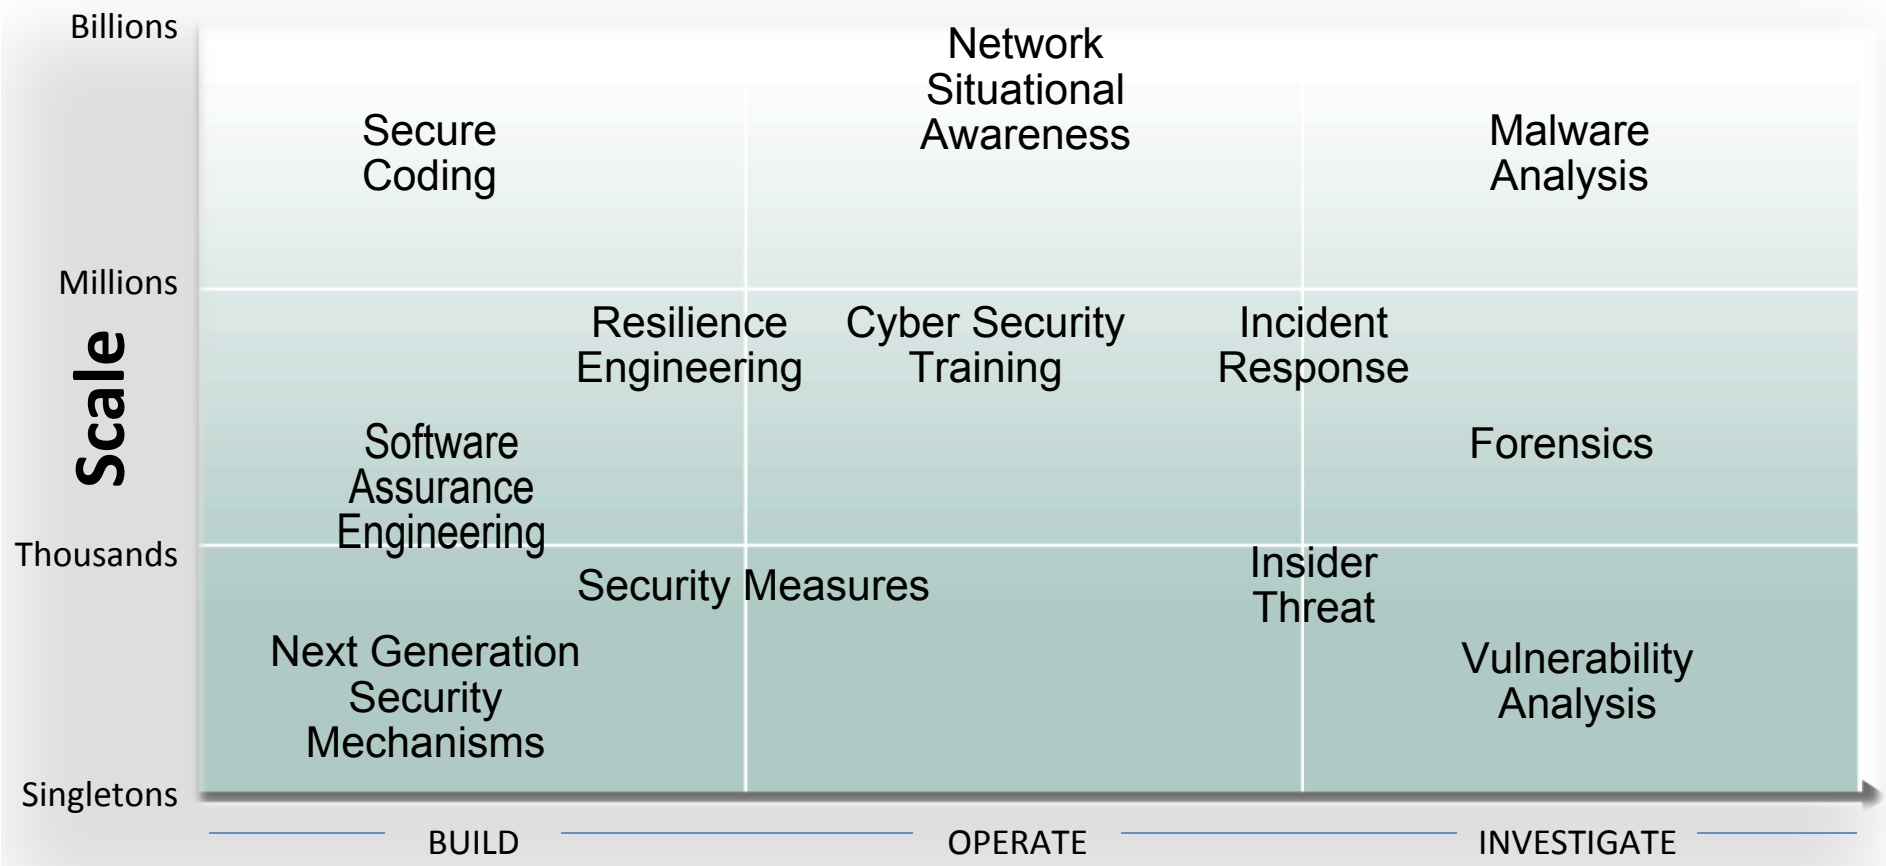

CERT Real Data for Research

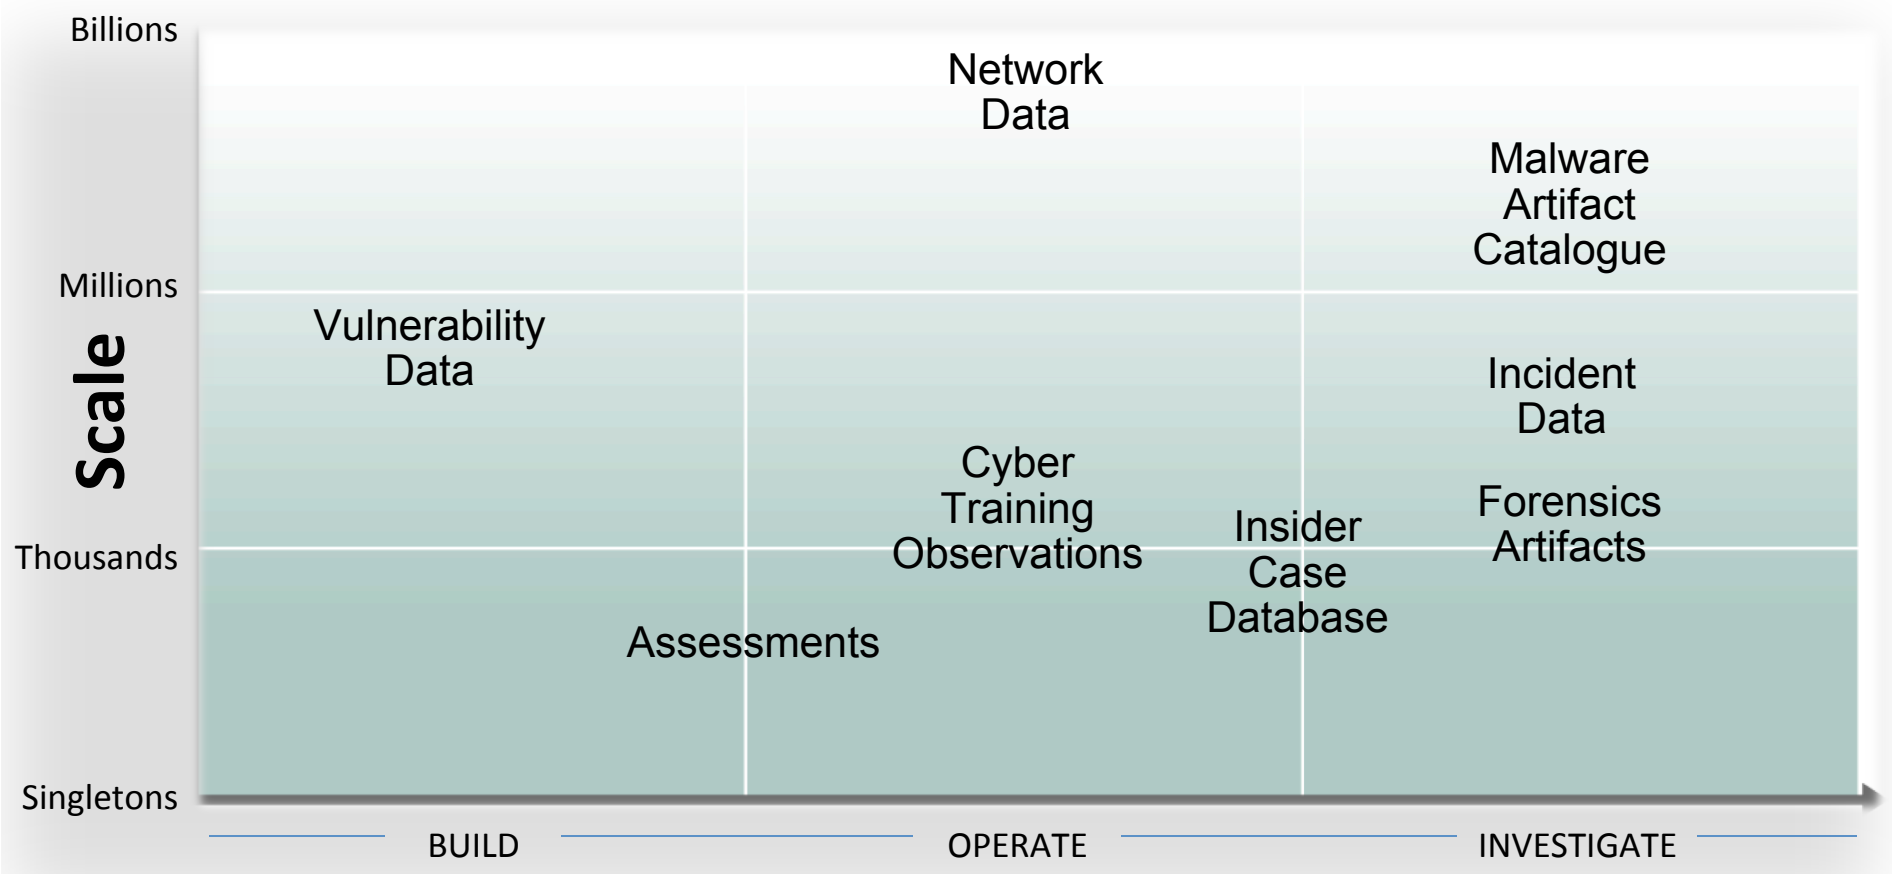

Software and Systems Lifecycle

CERT Synthetic Data for Research

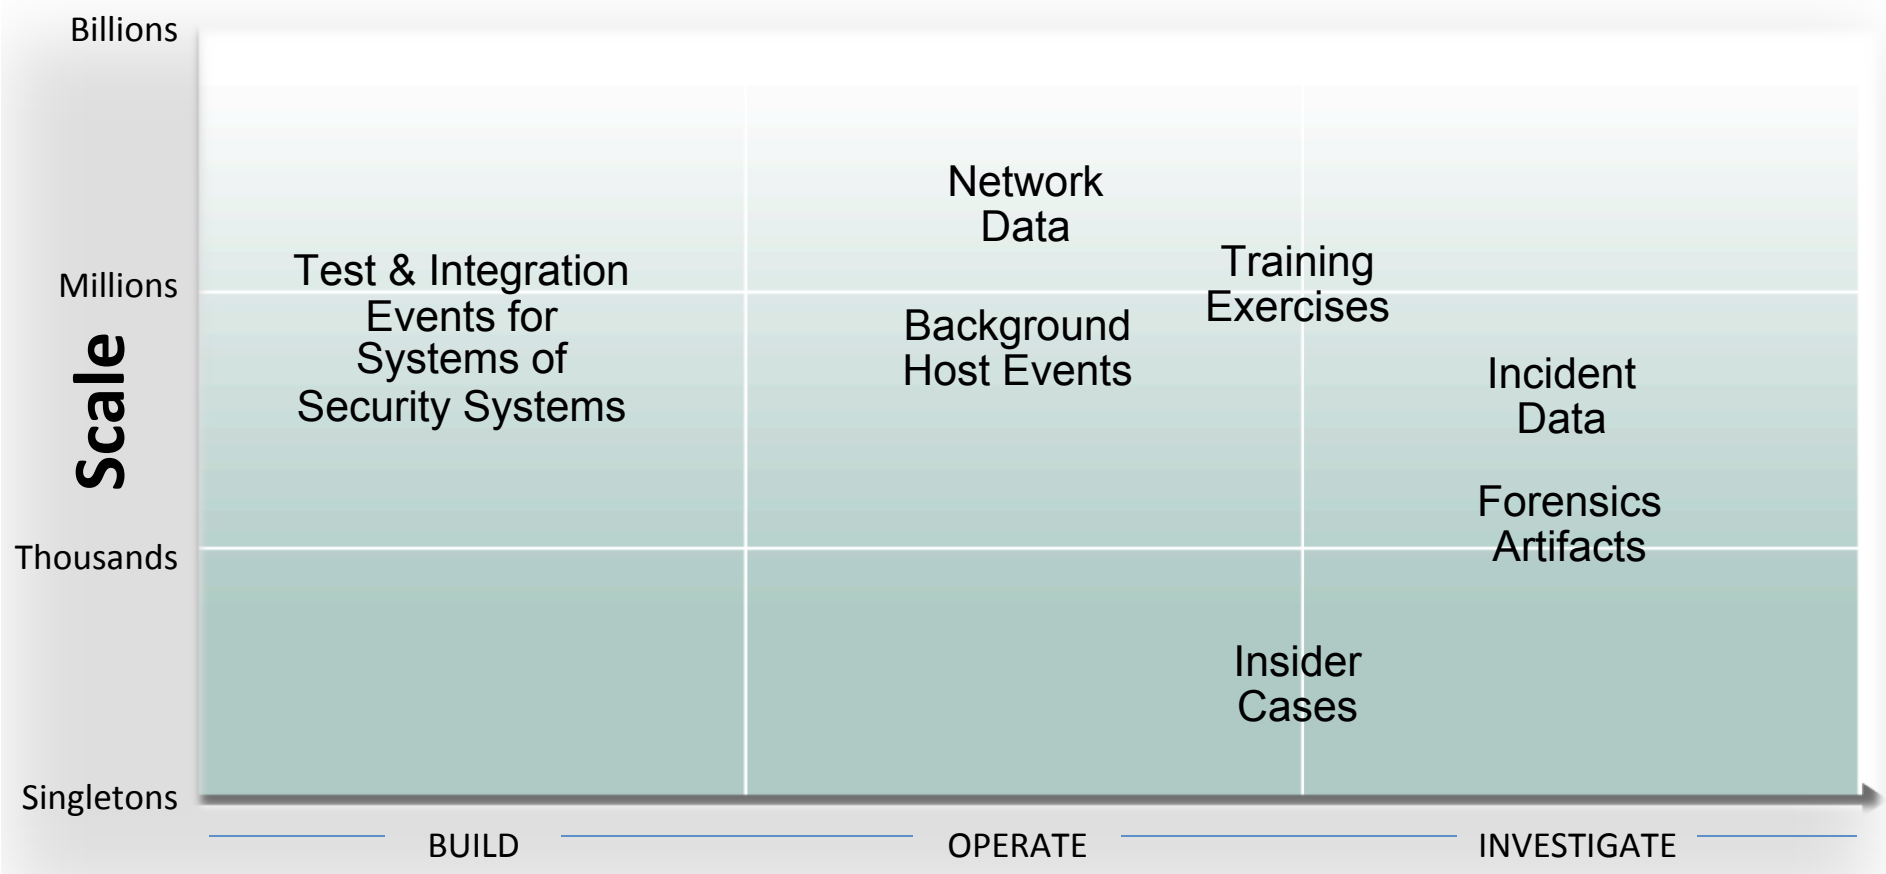

Realistic Malicious Events

- Inserting “needles” in data “haystacks”
- What should a needle look like?
- How to avoid easy “tells” (synthetic artifacts)?
- How to create and insert “needles” at scale for experiments, tests, exercises, evaluation, integration?
- A “DARPA hard” problem

Synthetic Cyber Security Data Challenges

- Synthetic “normal” data
- Modeling and simulation realistically at scale
- Malicious faults vs. “random” faults
- “Turing Tests”

U.S. Gov't and Cyber Security Science

- Dept. of Defense CTO, Hon. Zachary Lemnios
  - http://www.acq.osd.mil/chieftechnologist/areas/cyber.html
  - Campaigns in cyber measurement, modeling, simulation
  - Nov'10 JASON Report on: Science of Cyber-Security
  - DARPA-BAA-08-43, National Cyber Range
    - Goal #7: Use "the scientific method for rigorous cyber testing"
- White House support for "Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity"
  - National Science and Technology Council / NITRD / OSTP
  - Discussion on validity and the scientific method for cyber security
  - http://www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf
- National Academy study on Future Research Goals and Directions for Foundational Science in Cybersecurity
  - http://sites.nationalacademies.org/CSTB/CurrentProjects/CSTB_066764

Customer Driven Scientific Validity

- A result is scientifically valid when it is the product of a methodical process; when it is well documented, quantifiable, statistically sound, and reproducible; and when it produces principles that explain a testable class of phenomena.
  Results are analyzed for confounds; unmitigated confounds are identified and characterized.

Customer Driven Operational Validity

- A result (report, technology, capability, practice, policy, or process is operationally valid when it delivers in practice the measurable properties it was intended to deliver.
  Operational validity applies only to the properties actually observed, demonstrated, or measured in practice.
  For example, a capability realistically demonstrated on 1,000 systems is operationally valid for 1,000 systems, but not yet for 10,000 systems.

- Missing: importance of stating limitations

2012 LASER Workshop

- Learning from Authoritative Security Experiment Results
  - LASER 2012, July 18-19 in DC
  - PC Co-Chairs Matt Bishop (UC Davis) and Greg Shannon (CERT/CMU)
  - www.laser-workshop.org
  - Experimental failures, methods, confounds, mitigations
  - Co-sponsored by NSF

- Highlights
  - 40+ attendees including George Jones from CERT
  - Prof. Stuart Firestein from Columbia on *Ignorance: How It Drives Science*
  - Pro. Roy Maxion from CMU on *The Science of Security: Getting There from Here*
  - Government Panel w/ S.King, D.Dion-Schwarz, B.Martin, K.Landwehr on *The Role of Risk and Failure in Research*

2010 CERT Research Report

www.cert.org/research/2010research-report.pdf

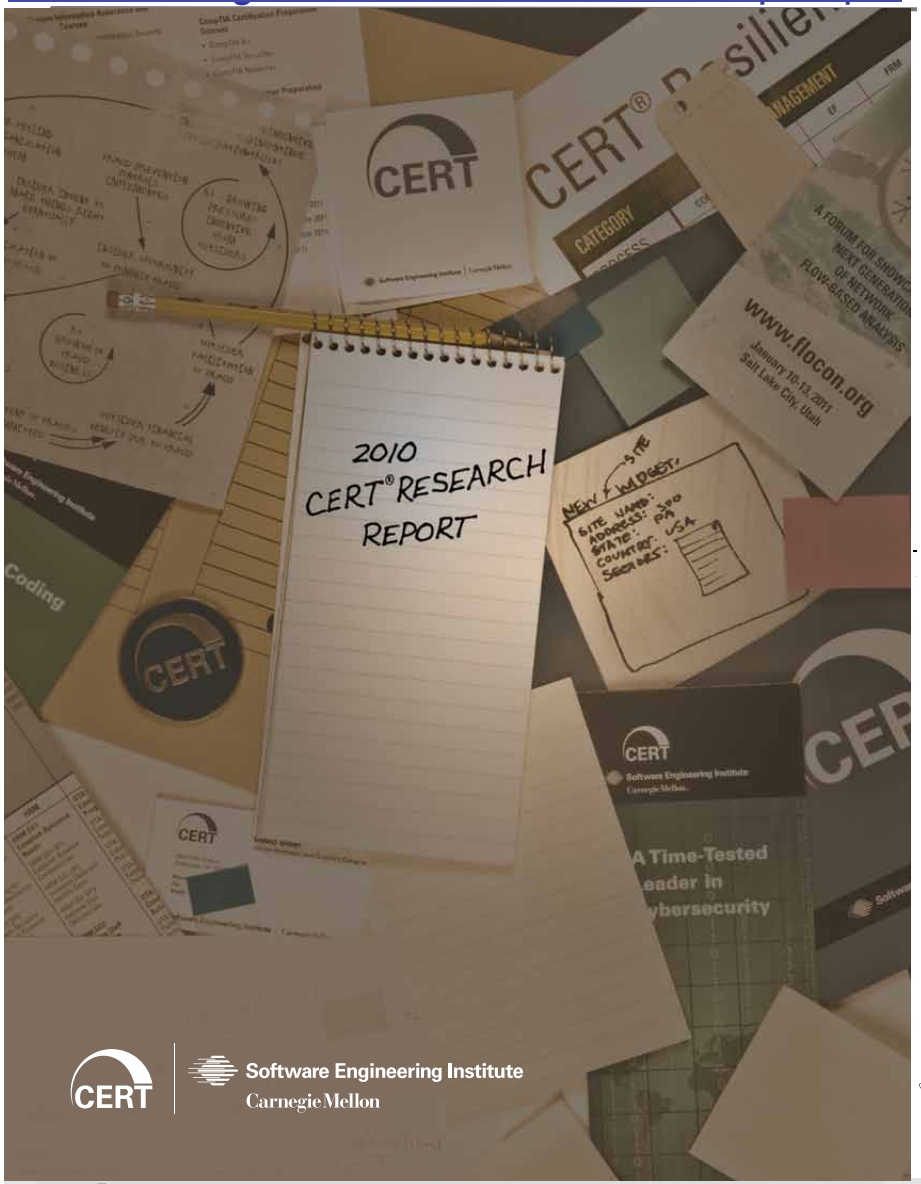

New Functions in Malware Binaries (2011)

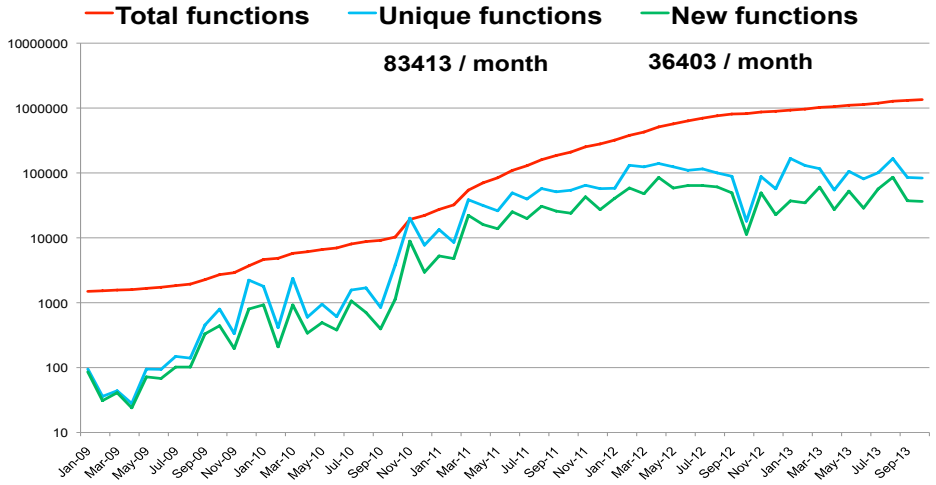

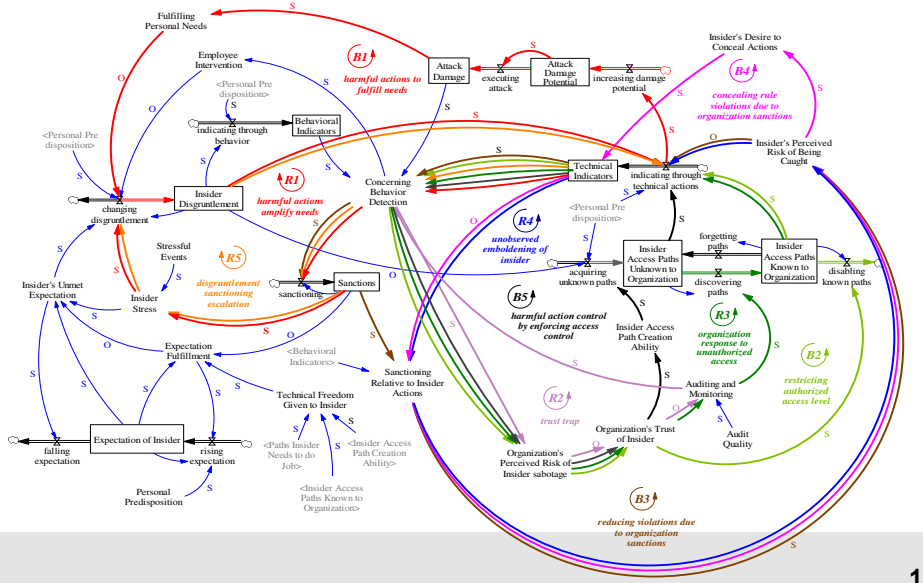

Thank You