FloCon 2024

20th Annual Open Forum for Large-Scale Data Analytics

SITUATIONAL AWARENESS: BEYOND THE NETWORK

Foresight: Using Incident Reports to Improve Measurability of Risk Exposure for Predictability

JANUARY 9–11, 2024

Brett Tucker
Technical Manager, Cyber Risk Management
CERT Division, Software Engineering Institute

Risk Analysis: How Do We Know if We Are Correct?

Does an incident happen within our probability estimates?

For example, Does 1 in 10 yr. event happen 1 in 10yrs?

(But we have only been doing this for 3 years…)

Does comparing risks with other systems help? (Or is each system too unique?)

Can we use Incident Reports to build our Risk Model?

- Can we use incidents to estimate likelihood?
- What information would we need? And how much?
- Are we collecting the right information now? Are we collecting too much info?
- If we had the information, how do we use it in our Risk Model?

Can this be done cheaply and efficiently? Or do we need help?

Can we build a community to build this "*anonymous*" incident data base?

Challenge Problem:

Given the vast number of cyber incidents that may be reported in the both the private and public sectors, what is the correct information for collection about that incident that can help us improve the assessment of risk both qualitatively and quantitatively?

Proposed Solution:

This body of research focuses upon the analysis of incident reporting (i.e., nature of data collected). Specifically, we are seeking to identify those data elements that ideally contributes to a more predictive understanding of cyber attacks.

In this model, is there a way to aggregate those data sets to determine what information provides the greatest context to inform risk-based decisions to avoid or mitigate future attacks?

Potential Impact for Our Community

Ultimately, we are looking to improve the assessment of risk, both qualitatively and quantitatively, with greater standardization and fidelity given the potential wealth of incident report data.

By doing so, we will enable the risk-based decisions with improved confidence control selection and reduce risk exposure. The goal is to reduce risk exposure with a predictive model that enables better risk-based decision making.

Approaches Considered

Carnegie Mellon University Software Engineering Institute

Our Approach

- What if we could map controls to risk?
  - Some models already exist (e.g., MITRE ATT&CK and D3FEND)
- What if we could identify patterns of critical controls across risks?
  - How well they perform and how often
- How is this better than just “Do All The Things™”?
  - Costs must be optimized for return on risk investment.

Data is always a challenge:

- Automation for volume?
- Synthetic data for modeling
- Authentic data from where?

Risk Analysis Current Ideal

The goal of risk analysis is to decide what to spend the budget on this year. These cybersecurity projects are ranked into three groups:

- Definitely fix
- Argue over which project to fix this year (*With the remaining money*)
- Accept this risk (*do nothing*)

This process works fairly well except when the decisions are tight.

(All the while hoping that the risk they accepted isn't the one that will get them).

Risk Analysis: “What If?”

What if, a company could easily determine the risk – to the level required by an Insurance company? (Comparable across all your systems)

What if, you knew which parameters actually mattered in determining risk?

- Which controls prevent an attack
- Which controls are redundant
- How deep is your defense in depth? (Single point of failures?)

What if, you could tune your cybersecurity projects to maximize defense and reduce cost?

Problems with the Standard Approach

Our risks are rank ordered by guesses.

What controls are critical to the mission?

It's difficult to tell whether more than one practice is protecting something.

How can you tell if you're missing a control?

Risk does not tell you which control to use.

There are questions about completeness:

What? Where? Why? When? How?

Immediate Observation: Business Impacts

Carnegie Mellon University Software Engineering Institute

Inconsistency in collection identified (of nine tools and requirements explored)

- Structured and predefined impact information necessary
  - Business Impact Analysis (BIA) non-prescriptive
  - This may be for a good reason – faster, cheaper, and proprietary
  - BIA can be hard to do among all other actions
  - Some reports ask for function impacted

Some Impacts are Qualitatively Defined

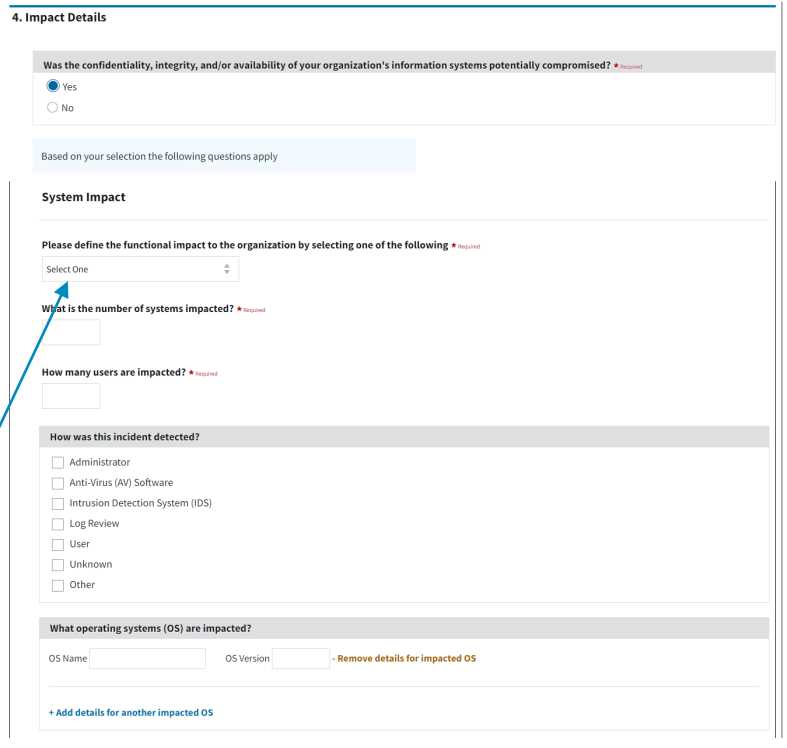

Incident Reporting System | CISA

Mandated by FEDRAMP Too

CISA Defined Impacts

2. **Information Impact** – Describes the type of information lost, compromised, or corrupted.

NO IMPACT – No known data impact.
SUSPECTED BUT NOT IDENTIFIED – A data loss or impact to availability is suspected, but no direct confirmation exists.
PRIVACY DATA BREACH – The confidentiality of personally identifiable information (PII) [6] or personal health information (PHI) was compromised.
PROPRIETARY INFORMATION BREACH – The confidentiality of unclassified proprietary information [7], such as protected critical infrastructure information (PCII), intellectual property, or trade secrets was compromised.
DESTRUCTION OF NON-CRITICAL SYSTEMS – Destructive techniques, such as master boot record (MBR) overwrite; have been used against a non-critical system.
CRITICAL SYSTEMS DATA BREACH - Data pertaining to a critical system has been exfiltrated.
CORE CREDENTIAL COMPROMISE – Core system credentials (such as domain or enterprise administrative credentials) or credentials for critical systems have been exfiltrated.
DESTRUCTION OF CRITICAL SYSTEM – Destructive techniques, such as MBR overwrite; have been used against a critical system.

Qualitative impact good for:

- Categorization
- Rapid documentation
- Low-cost analysis

Better definition and quantification could:

- Improve risk-based decisions through return on risk investment calculations.
- Assist in understanding threat actor intentions and motivations.

Note That Recovery is Also Qualitatively Considered by CISA

Immediate Observations: Additional to Collect Upon

Consideration for similar events

- May lend itself to frequency of occurrence

Defined response actions with indication of costs

- Secondary impacts provide improve priorities and risk decisions

Taxonomy for vector of attack

- Specific knowledge of attack vectors improve control selection

Improved Reporting = Greater Burden

Overcoming Additional Collection Costs

Additional reporting requirements means more time and money.

- Business cases necessary to demonstrate improved efficacy
  - General cases for all
  - Sector specific cases

- Phased roll-out of collection
  - Prioritize new asks and slowly introduce to the community
  - Specific direction may yield greater results
  - It is unclear if additional or new tooling would be necessary

- Executive advocacy a must
  - Can government procurement expectations be the driver?

Additional Impact of Research

CISA CPGs Use Case Example

The measurement of control efficacy has come up lately with **CISA Cyber Performance Goals (CPGs)** since they:

- Establishing baseline practices to reduce risk exposure
- Prioritizing security practices
- May lead to a greater understanding of aggregate risk to the nation

Challenges exist in terms of measuring cost, complexity, and impact provided

- Approaches for quantification may vary based on context

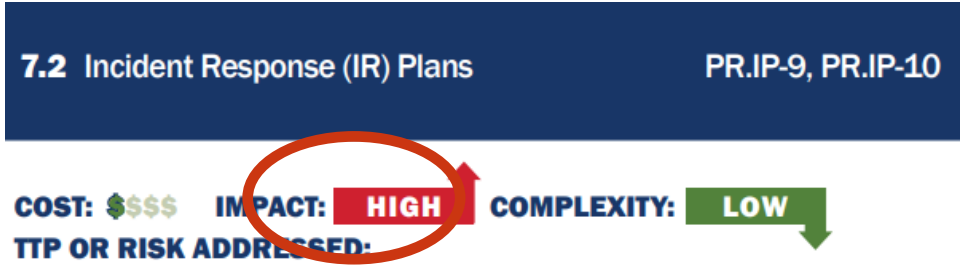

Why Do We Want to Do This?
Imagine the Possibilities

We reduce risk exposure with cyber controls.

The goal is to quantify the efficacy of cyber controls.

- Prioritization of controls can inform capital investment related to:
  - Procurement
  - Implementation
  - Management

  } Expensive!
- Improve risk-based decision making
- Attempts at qualification leave gaps for interpretation

Challenging to Account for Context – Can we Standardize the Measurement?

Measure Risk to Optimize Cybersecurity Investment

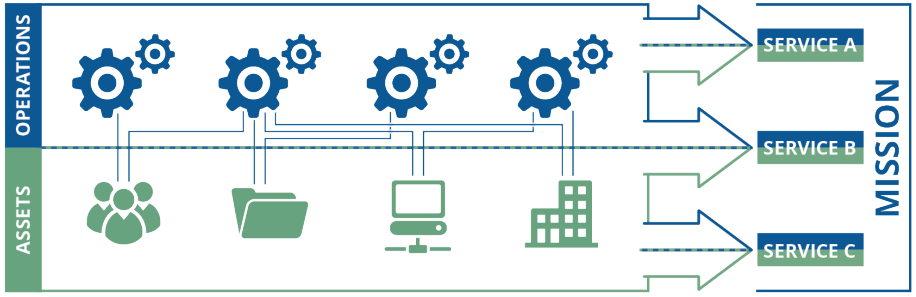

People: those who operate and monitor the service
Information: data associated with the service
Technology: tools and equipment that automate and support the service
Facilities: where the service is performed
Third Party Providers: external suppliers that we rely upon to deliver service

Risk Quantification and Management

- Risk-informed cyber control selection and evaluation
- Econometrics of cybersecurity and return on cybersecurity investment
- Cybersecurity to resilience transformation

Resilience Diagnostics

- Adversary emulation and penetration testing
- Rigorous measurement of capabilities and benchmarking
- Cyber incident study and control analysis

Call for Action

Research continues

- Significant volume of incident reports to be reconciled
- Are we missing something?
  - Some reporting methods better than others
  - Analysts have varied aptitude for completing reports
  - Organizations differ on resource investment for reporting
- We can collect a lot of data and we can analyze it to get information
  - Is there a return on risk investment?
  - Can we demonstrate that our collections improve risk decisions?

Backup Slides

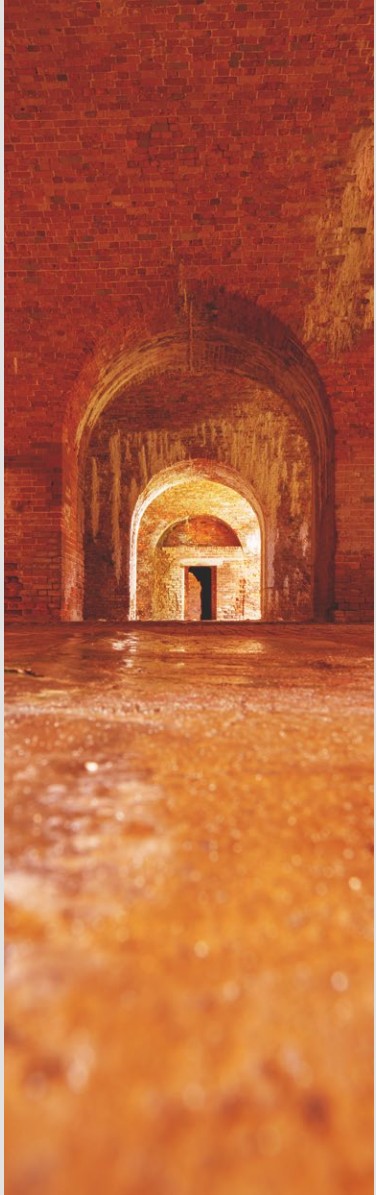

Current State vs What Could Be

How we do it now

Likelihood \ Impact	Negligible	Minor	Moderate	Significant	Severe
Very Likely	Low	Moderate	High	High	High
Likely	Low	Moderate	Moderate	High	High
Possible	Low	Low	Moderate	Moderate	High
Unlikely	Low	Low	Moderate	Moderate	Moderate
Very Unlikely	Low	Low	Low	Moderate	Moderate

New improved process

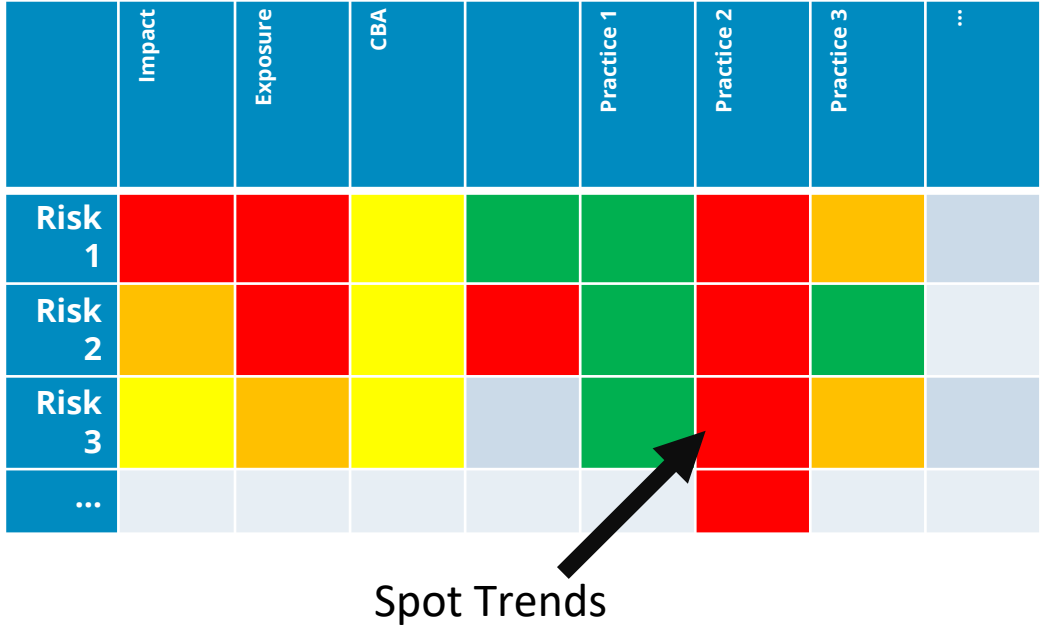

What about Black Swan Events?

How do you estimate rare events?

What about first-time events?

Why do black swan events happen more often than they should?

What if you have an active adversary deliberately trying to create one?

Factors for Measurement

Cyber controls come in many forms:

- Technical
- Administrative
- Physical

All have the same goal of preserving:

- Confidentiality
- Integrity
- Availability

Organizations must establish a defense-in-depth strategy that overlaps controls.

- Optimize investment with risk-based decision making
- Determine return-on-risk investment with proper measurement

Additional Factors for Consideration

Controls may be selected to achieve specific goals:

- Preventive Controls – defend your system from incidents occurring
- Detective Controls – seeking out errors or irregularities
- Corrective Controls – fixing identified errors
- Compensating Controls – addressing weaknesses of existing controls

With So Many Considerations at Hand, We Must Consider the Context of the Control Implementation as Well as the Additive Effect of Controls When Used Together.

Various Means of Measurement

Cyber control efficacy may be measured by:

- Penetration Testing
- Security Operation Center Analysis
  - Enabled by SIEM systems and logging
- Compliance and Regulatory Audits
- Incident Response Measures
  - Mean Time Detection, Recovery Point Objective, Recovery Time Objectives

Measures of Efficacy May Consider Resiliency of Organization or Threat Prevention

Measuring Efficacy of Control

Control efficacy measurements may depend on:

- **Strategy** – prevent threat or mitigate impact
- **Type** – administrative, technical, physical
- **Periodicity of Measurement** – alert or monthly reporting
- **Governance** – what does the leader need to make the right decision

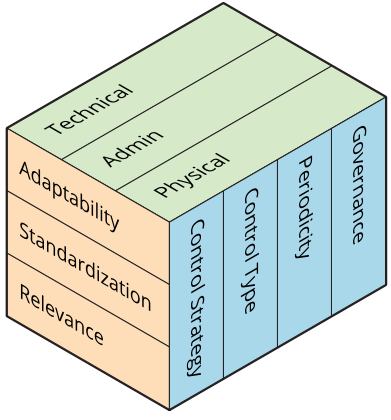

Measures of control efficacy should possess:

- **Adaptability** – fit the context to achieve objectives
- **Standardization** – used throughout the enterprise with consistent tolerances
- **Relevance** – drive decision making to render value

Other Attributes May Apply Depending Upon Organizational Context

Iterative Development May Accommodate a Dynamic Environment