The Future of Cybersecurity Risk Measurement

Jack Jones

The future of cybersecurity risk measurement is…

Automation and AI

Any questions?

Clarification: risk measurement vs. other measurements

Maturity models

Vulnerability rating models

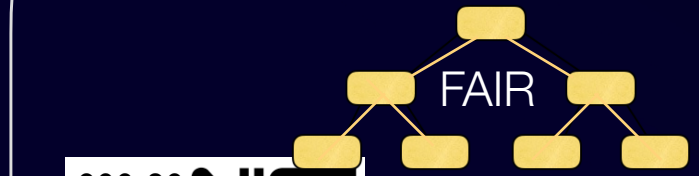

Risk measurement models

Likelihood and Impact

Risk management models

Credit scoring-like models

What we're going to cover today…

Part 1: Risk measurement past and present

Part 2: Fireside chat — Why is CRQ worthwhile for organizations?

Part 3: The future

Can you relate to one or more of these?

- “Religious battles” over risk ratings
- Too much to do — everything’s important
- How many mediums equals a high?
- Difficulty explaining expensive cybersecurity improvements
- What should the thresholds be for KRIs and KPIs?
- Executives that are too quick to accept risk

These exist in large part due to risk measurement problems

Risk Measurement Past and Present

Mid-to-late1980's - early attempts to quantify information security risk in economic terms failed, creating a perception that it couldn't be done.

Late1990's - Interest in cyber risk quantification (CRQ) resurfaced, but approaches were complicated and deemed impractical. This reinforced the perception that qualitative measurement was the only choice.

Early-to-mid 2000's - FAIR developed and published. Was seen as relatively easy to understand and use. Adoption was slow due to established perception that CRQ can't be done.

Early 2010's - FAIR adopted by The Open Group as an international standard.

Mid 2010's - FAIR Institute established. Now with over 13,000 members.

Late 2010's - the question regarding CRQ began to change from "Can it be done?" to "Should it be done?"

Today - the question regarding CRQ is beginning to change from "Should it be done?" to "How do we do it?"

Which of these is most important to fix first?

Issue #1

An audit discovered that privileges are not consistently being updated for user accounts with access to a customer service application containing PII.

Issue #2

A security assessment determined that the organization was unlikely to be able to identify when a cyber criminal breaches its network perimeter.

Both were identified as “high risk”

How do you like your risk measurement?

Fast, cheap or good.

Pick two.

How much time and expense was applied to those risk analyses we did just now?

When you analyze it quantitatively…

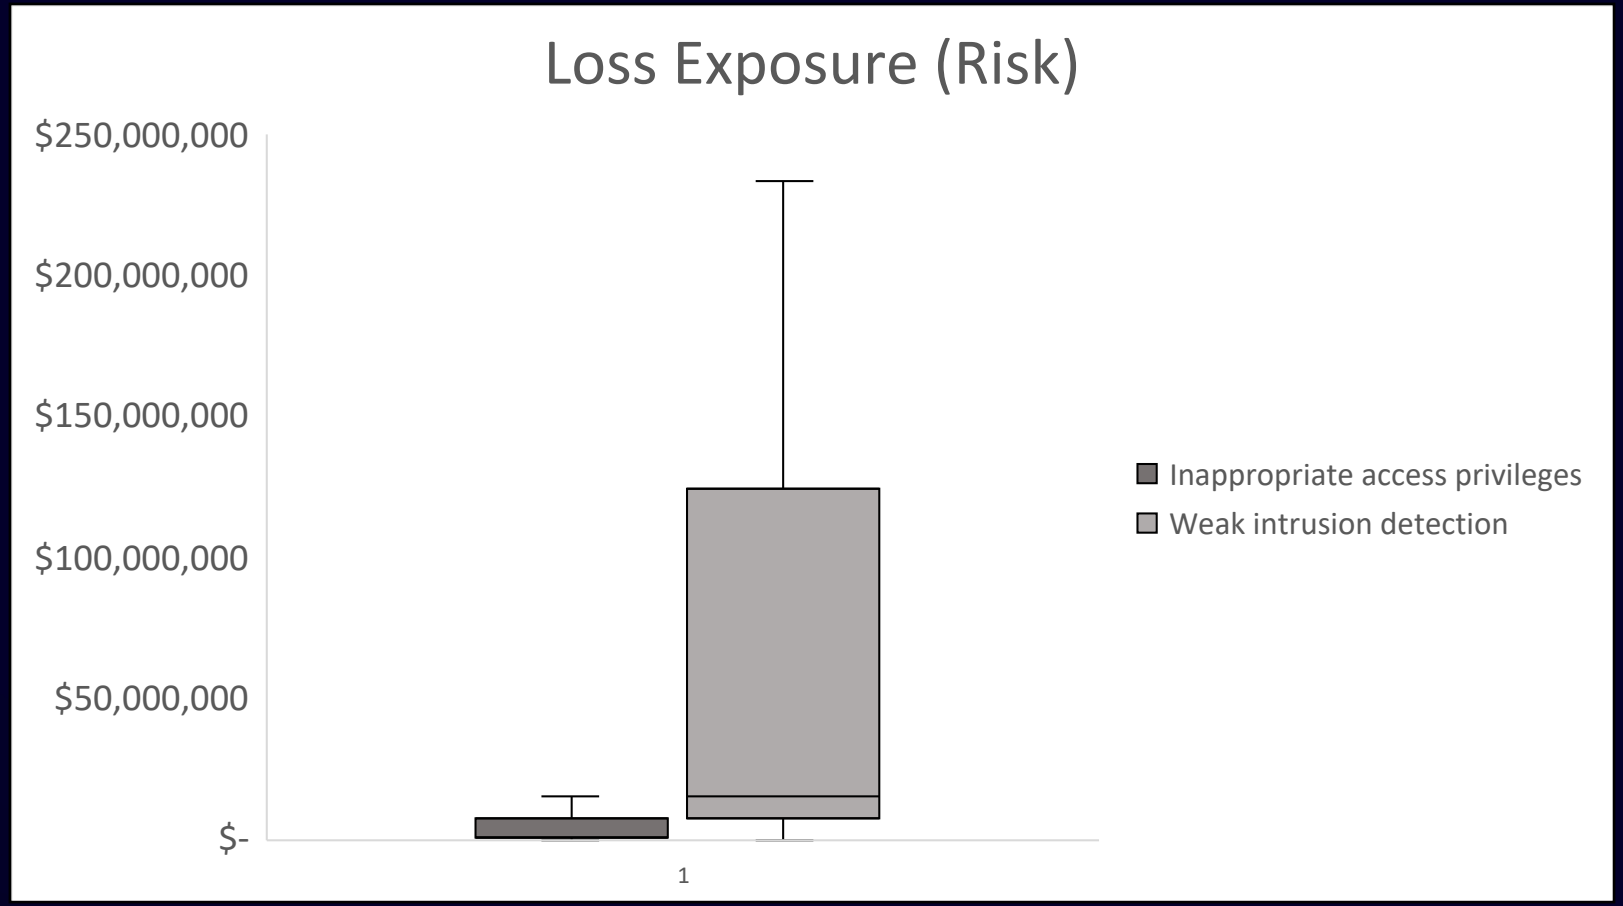

Key take-away…

There’s always a trade-off to be made in the speed, cost, and quality of risk measurement.

Evaluating risk measurement

Three fundamental elements of risk measurement

1. The scope of what's being measured

2. The model being used

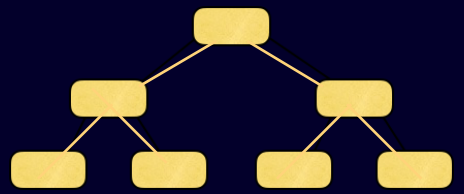

3. Data

Getting clarity thru effective scoping

A measurement example

How fast are they going?

Qualitatively

Challenges…

- Is your “Fast” the same as mine?
- Which car am I referring to?
  - One in particular? (Slowest? Fastest?)
  - An average for all of them?
- Which part of the track am I referring to?
  - Corners?
  - The straightaway?
  - Average over the entire track?
  - This lap, or an average for the entire race?

Measuring speed

Requires three elements:

1. The scope of what's being measured
   - Which car(s)?
   - Which part of the track?
   - Which lap(s)?

2. An analytic model
   - What data do we need? (time, distance)
   - How do we apply the data? (speed = distance/time)

3. Data

Measuring risk

Every risk measurement involves three elements:

1. The scope of what's being measured
   - What asset?
   - What threat?
   - Which vector?
   - Which controls are relevant?
   - What type of event (e.g., C, I, A)?
2. An analytic model
   - What data do we need?
   - How do we apply the data?
3. Data

A scoping example

Outage of key business systems due to cybercriminals succeeding in a ransomware attack via a phishing e-mail campaign.

- Asset
- Threat
- Effect
- Method
- Vector

Key take-away…

Without clear scoping, the odds of measuring risk accurately are much lower…

…regardless of whether you're doing qualitative or quantitative measurement

But scoping takes time

Models

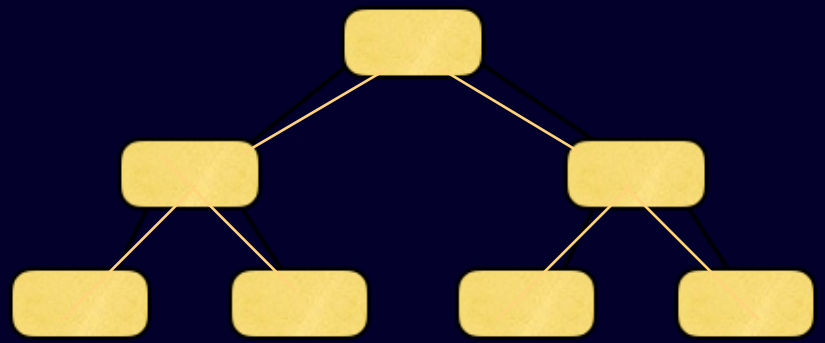

A model is a simplified representation of reality used to describe, simulate, or make forecasts.

"All models are wrong, but some are useful."

George Box

But there are different types and degrees of "wrongness"…

What's wrong with these models?

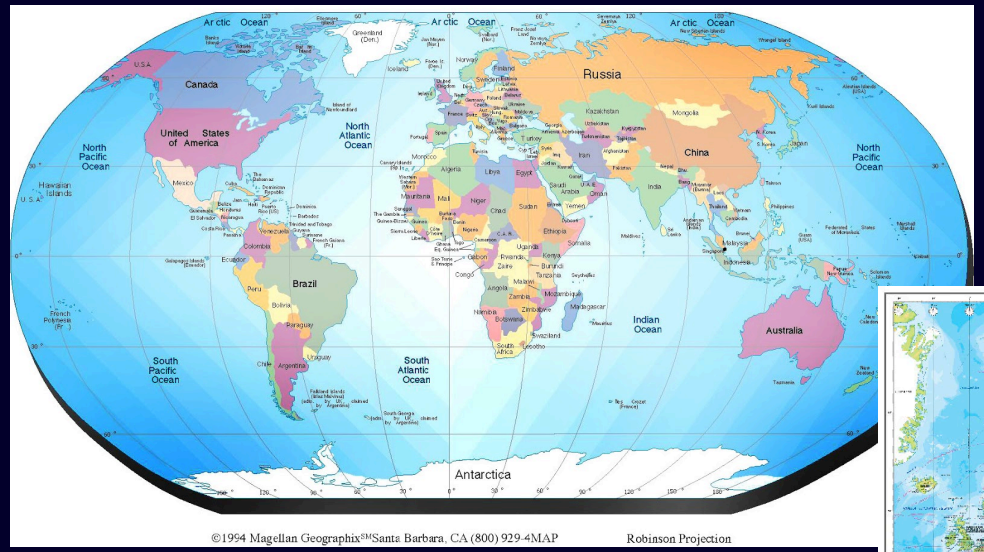

They're wrong, in that they aren't perfect replicas.

A different kind of wrong…

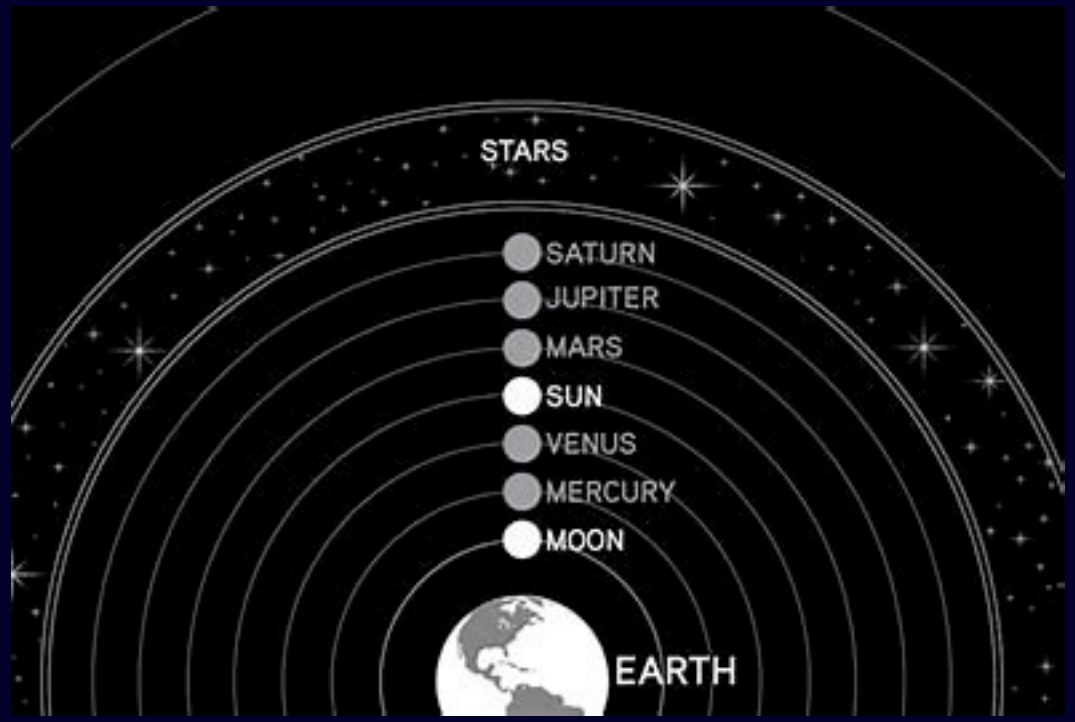

An example cybersecurity risk model (not FAIR)

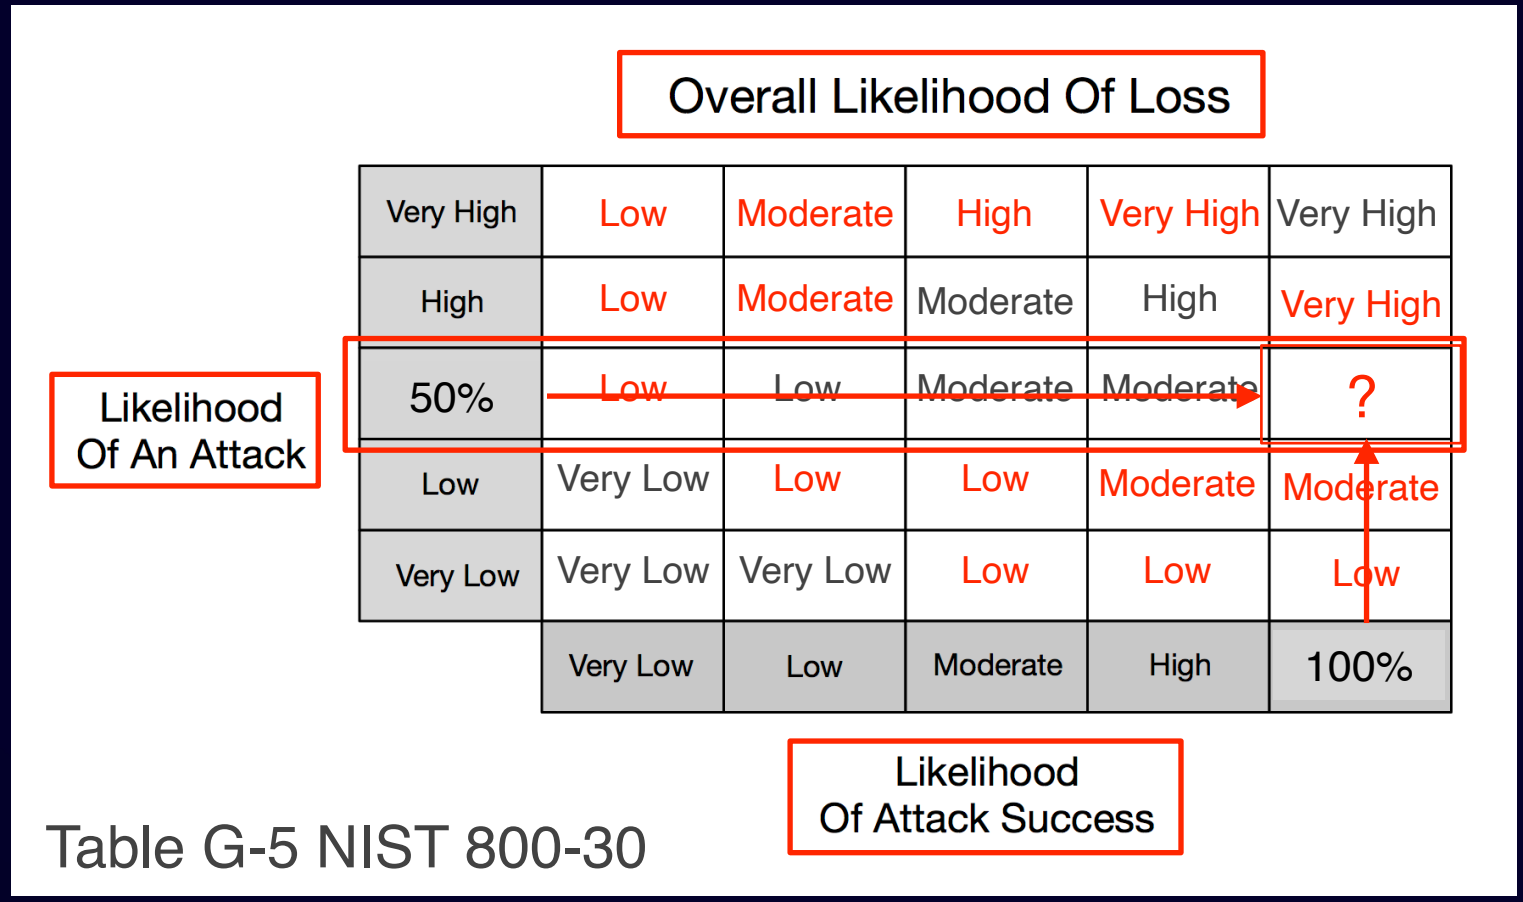

Likelihood Of An Attack	Very Low	Low	Moderate	High	Very High
Very High	Low	Moderate	High	Very High	Very High
High	Low	Moderate	Moderate	High	Very High
50%	Low	Low	Moderate	Moderate	?
Low	Very Low	Low	Low	Moderate	Moderate
Very Low	Very Low	Very Low	Low	Low	Low
	Very Low	Low	Moderate	High	100%

Table G-5 NIST 800-30

What is the most commonly used cyber risk measurement model?

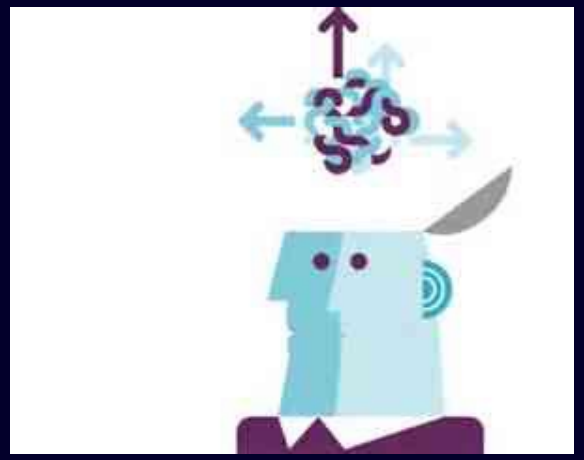

Mental models

What scope?

What formula?

What data?

The Factor Analysis of Information Risk (FAIR) Model

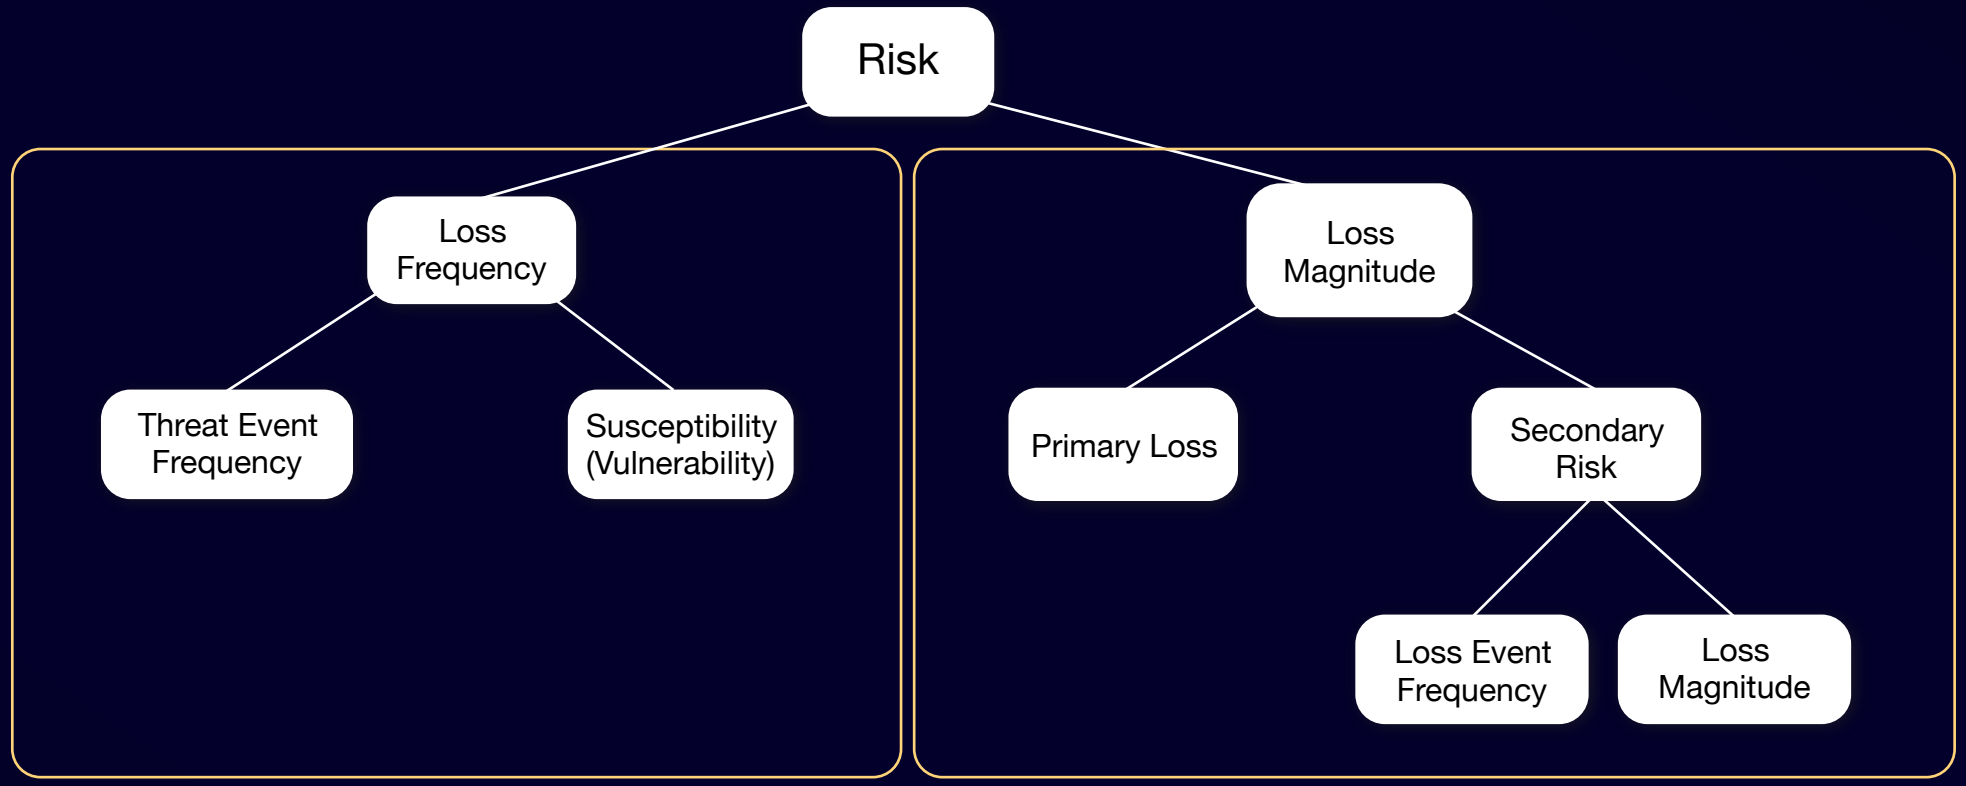

Loss Event Frequency

Loss Magnitude

FAIR advantages

- Is a clearly defined model that defines the relationships between factors
- Normalizes terminology
- Is an open, international standard
- Has a professional certification available thru the Open Group
- Supported by a large global community of users thru the FAIR Institute
- Included in numerous college courses
- Complements, rather than replaces existing frameworks
- Can be used quantitatively or qualitatively
- Supports both quick-and-dirty and deep analyses
- Can be used to measure any type of risk

Key take-aways…

- No models are perfect, but some are fundamentally broken
- The most commonly used, risk measurement model is the individual professional's mental model, where we don't know
  - The scope
  - Formula, or
  - Data
- FAIR is an open standard model that is pragmatic, flexible, and well-vetted

But what about data?

“We don’t have enough data.”

“You have more data than you think you do.”

“You need less data than you think you do.”

Douglas Hubbard

Author of “How to Measure Anything”

The problem of uncertainty…

How tall am I?

Uncertainty is inevitable. It's simply a matter of whether it's accounted for in measurement inputs and outputs.

Questions for any risk analysis…

- What data do we need? The scope and risk model tell us this
- How do we apply them? The model tells us this

If the analysis is scoped clearly and you're using a well-defined model, then data will be far less challenging to gather and apply.

Key take-aways…

- Data is never the limiting factor in risk analysis. You just may have more uncertainty than you're comfortable with.
- The key is to faithfully reflect your uncertainty by the width of your range or distribution.

Remember these?

- “Religious battles” over risk ratings
- Too much to do — everything’s important
- How many mediums equals a high?
- Difficulty explaining expensive cybersecurity improvements
- What should the thresholds be for KRIs and KPIs?
- Executives that are too quick to accept risk

How else does poor risk measurement affect an organization?

Key takeaways…

- Historically, and even predominately today, our profession has focused on fast and easy risk measurement, without a clear understanding of what "good" measurement looks like or requires.

- Good risk measurement requires:
  - A clear scope of what's being measured
  - A well-defined and vetted model
  - Data (input values that account for uncertainty)

Additional resources

- Measuring and Managing Information Risk: A FAIR Approach
  - Jack Jones & Jack Freund
- How to Measure Anything in Cybersecurity Risk
  - Douglas Hubbard & Richard Seiersen
- FAIR Institute (www.fairinstitute.org)
- The Open Group (www.opengroup.org)

Questions?

Fireside chat:
Why is CRQ worthwhile for organizations?

The Future of Cybersecurity Risk Measurement

Remember “Fast”, “Cheap”, or “Good”?

A “good” risk measurement method does not have inherent flaws in scoping, modeling, or the use of data.

You have to be able to defend your results

Cyber risk measurement of the future will…

1. Be quantitative
2. Account for controls physiology
3. Be based on open, non-proprietary models
4. Where possible, leverage automation and AI
5. Be performed by trained and certified professionals

#1 - The future is quantitative

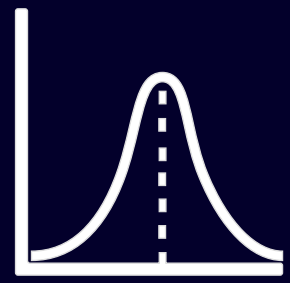

Good risk measurement does not require quantification.

But the future of cybersecurity risk measurement is quantitative.

Immutable limitations of qualitative measurement

- Inherent imprecision
- Difficult to clearly define ordinal scales
- Inability to aggregate risk
- Inability to do cost-benefit analyses
- The inputs and outputs can’t be validated
- Creates an apples-to-oranges problem for executives

Ordinal scales

One vs. many?

Timeframe?

Severity	Probability
5. *Catastrophic* Likely to result in death.	5. *Frequent* Hazard highly likely to occur.
4. *Critical* Potential for severe injury.	4. *Probable* Hazard will be experienced.
3. *Moderate* Potential for moderate injury.	3. *Occasional* Some manifestations of the hazard are likely to occur.
2. *Minor* Potential for minor injury.	2. *Remote* Manifestations of the hazard are possible but not likely.
1. *Negligible* No significant risk of injury.	1. *Improbable* Manifestations of the hazard are very likely.

Table I. Qualitative probability and severity scaling.

Math on ordinal values

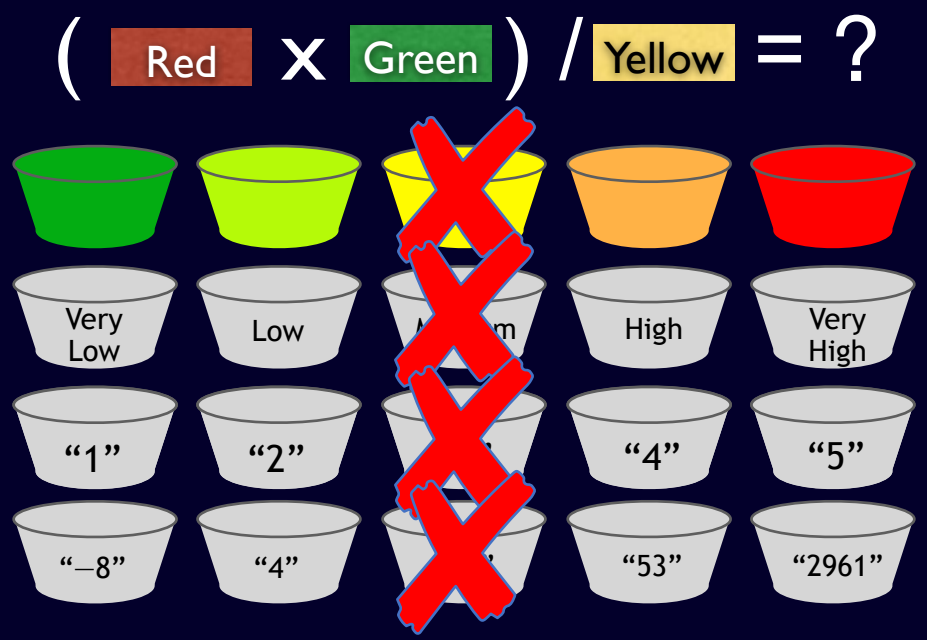

No unit of measurement

Data validation

Which of these can be validated?

- Our network logging is “Medium”, or “2”
- We log 75% of the traffic on our network

Key take-aways…

- The future of cyber risk measurement is quantitative because qualitative risk measurement has too many inherent limitations.
- Math on ordinal values generates fundamentally unreliable results.

#2 Controls Physiology???

FAIR Controls Analytics Model (FAIR-CAM)

What is it, and why do we need it?

Ask yourself these questions…

- What's the most valuable control in your cybersecurity program?
- What's the least valuable control?

Would your answers be the same as someone else's in your organization?

In the practice of medicine, which is more important?

Anatomy?
(The parts of the system)

OR

Physiology?
(How the system works)

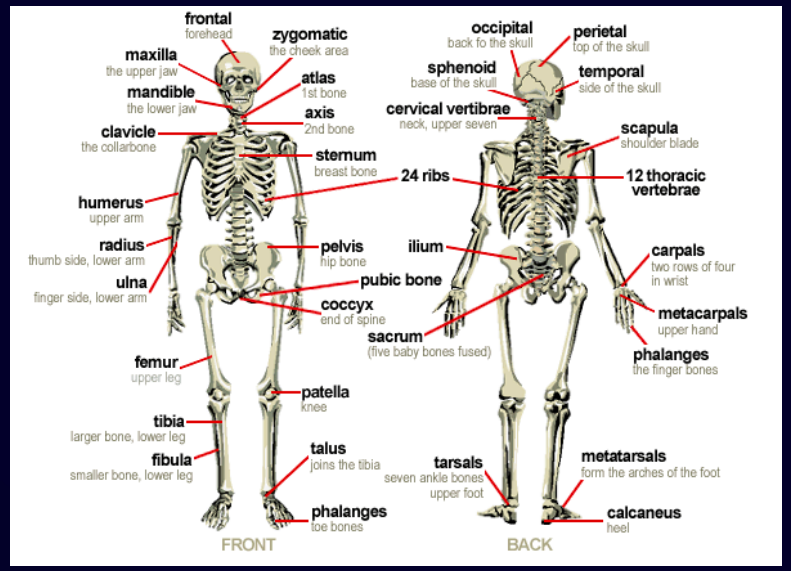

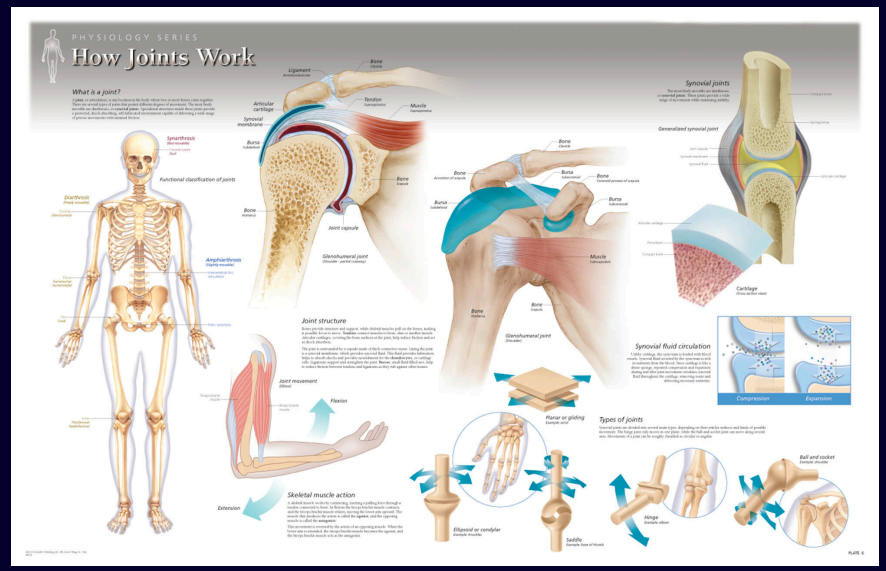

Neither. You need to know both.

Human Anatomy vs. Physiology

- **Anatomical component**: Spleen
  - Size: Approximately 1 x 3 x 5 inches
  - Weight: Approximately 7 oz
  - Location: Upper-left abdomen

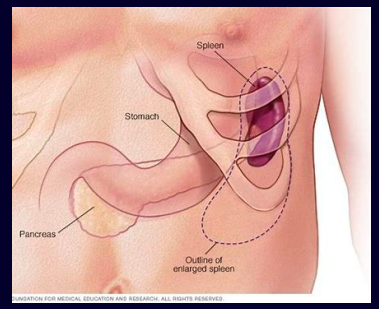

- **Purpose**: Supports the immune system

- **Physiology**
  - Function: Blood filtering via white pulp and red pulp
  - Depends upon: Arteries, veins, nerves, lungs, etc…
  - Is depended upon by: Liver, brain, etc…
  - When missing or damaged is partially compensated for by: Lymph nodes, etc…

In other words, it's part of a system.

Cybersecurity Anatomy vs. Physiology

- **Anatomical component**: Awareness training
  - Content: Passwords, phishing, clean desk, etc.
  - Periodicity: Annual
- **Purpose**: Informs personnel about their responsibilities
- **Physiology (how it functions within the system to reduce risk)**
  - Function: Improves decision-making, which reduces human error
  - Depends upon: Policies, risk appetite, risk measurement, etc…
  - Is depended upon by: Authentication, system security, access privileges, physical security, data protection, etc…
  - When deficient, may be partially compensated for by: DLP, password enforcement, Anti-malware, etc.

Direct vs. indirect effects on risk

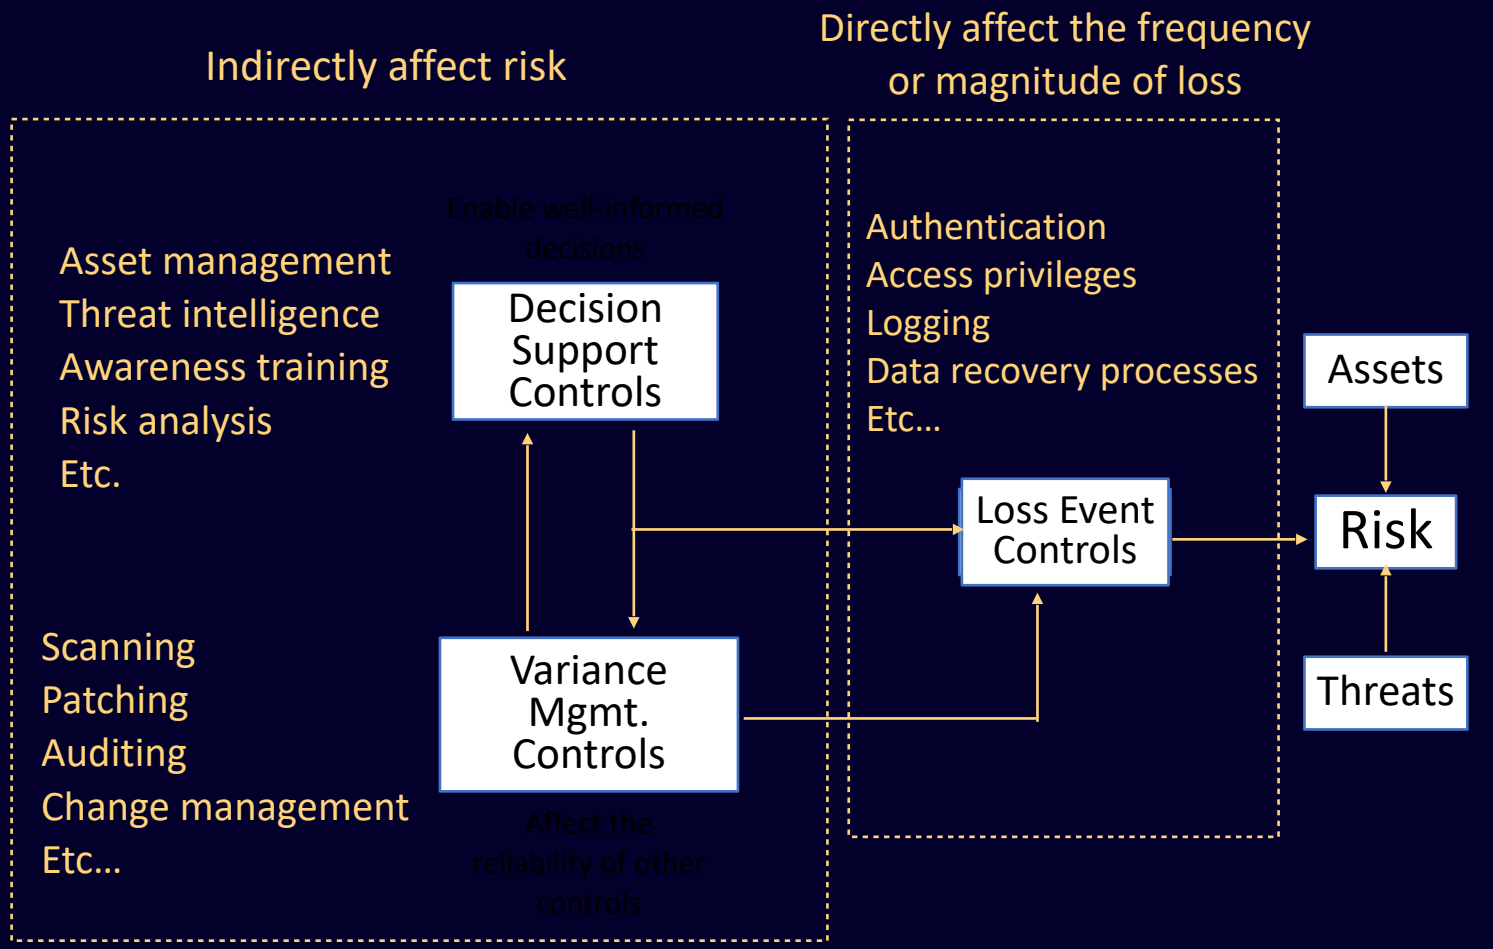

An example of direct vs. indirect effects…

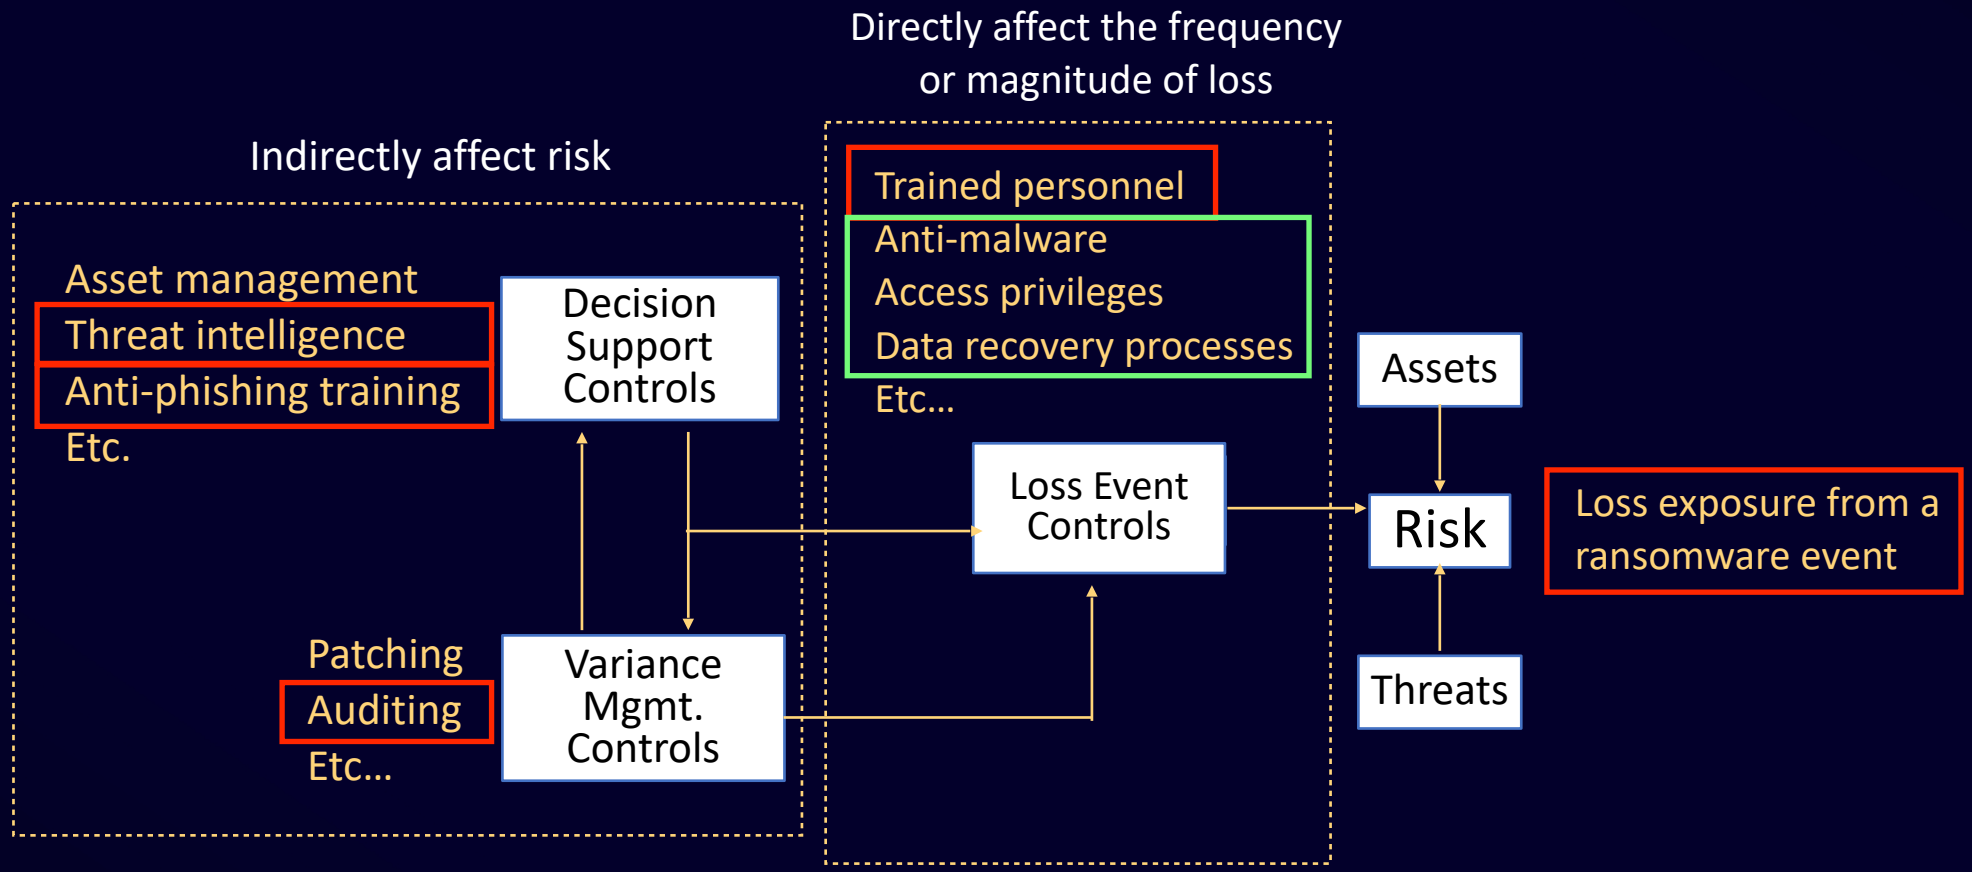

Loss Event Control Functions

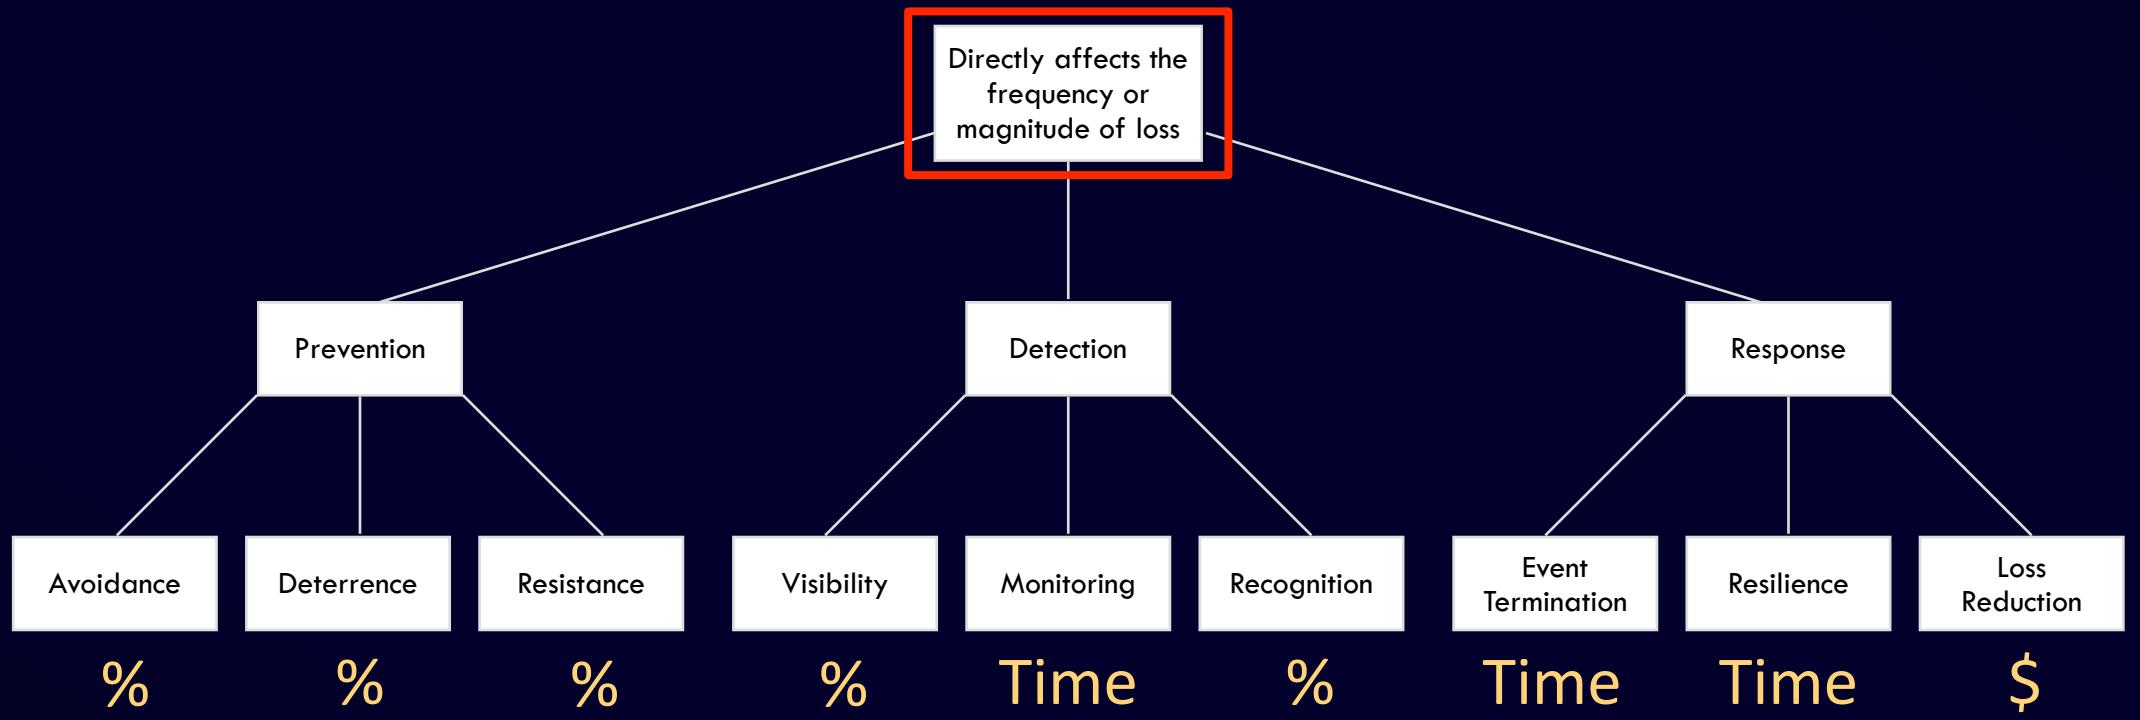

Variance Management Control Functions

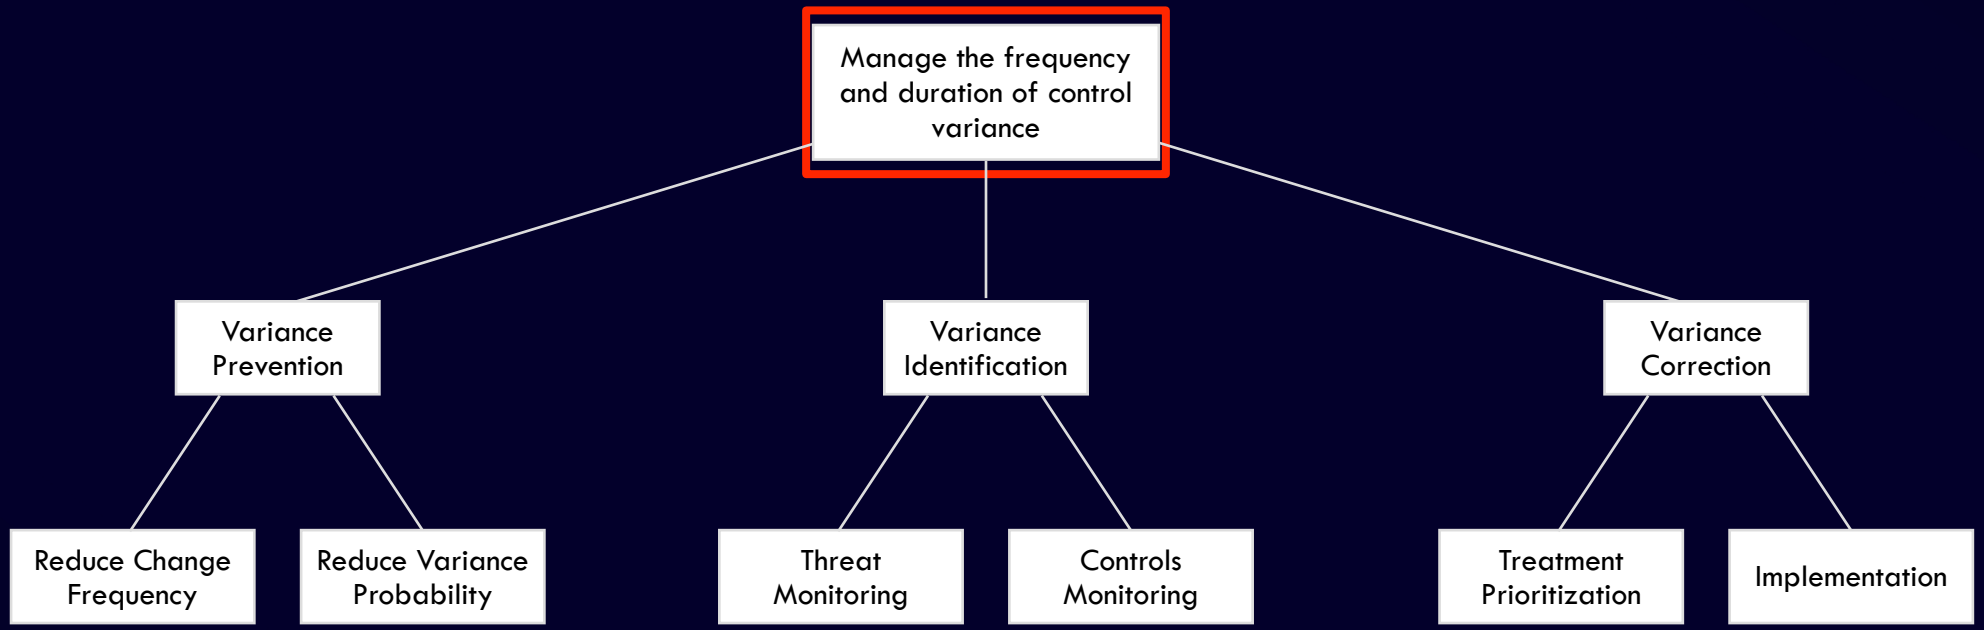

Decision Support Control Functions

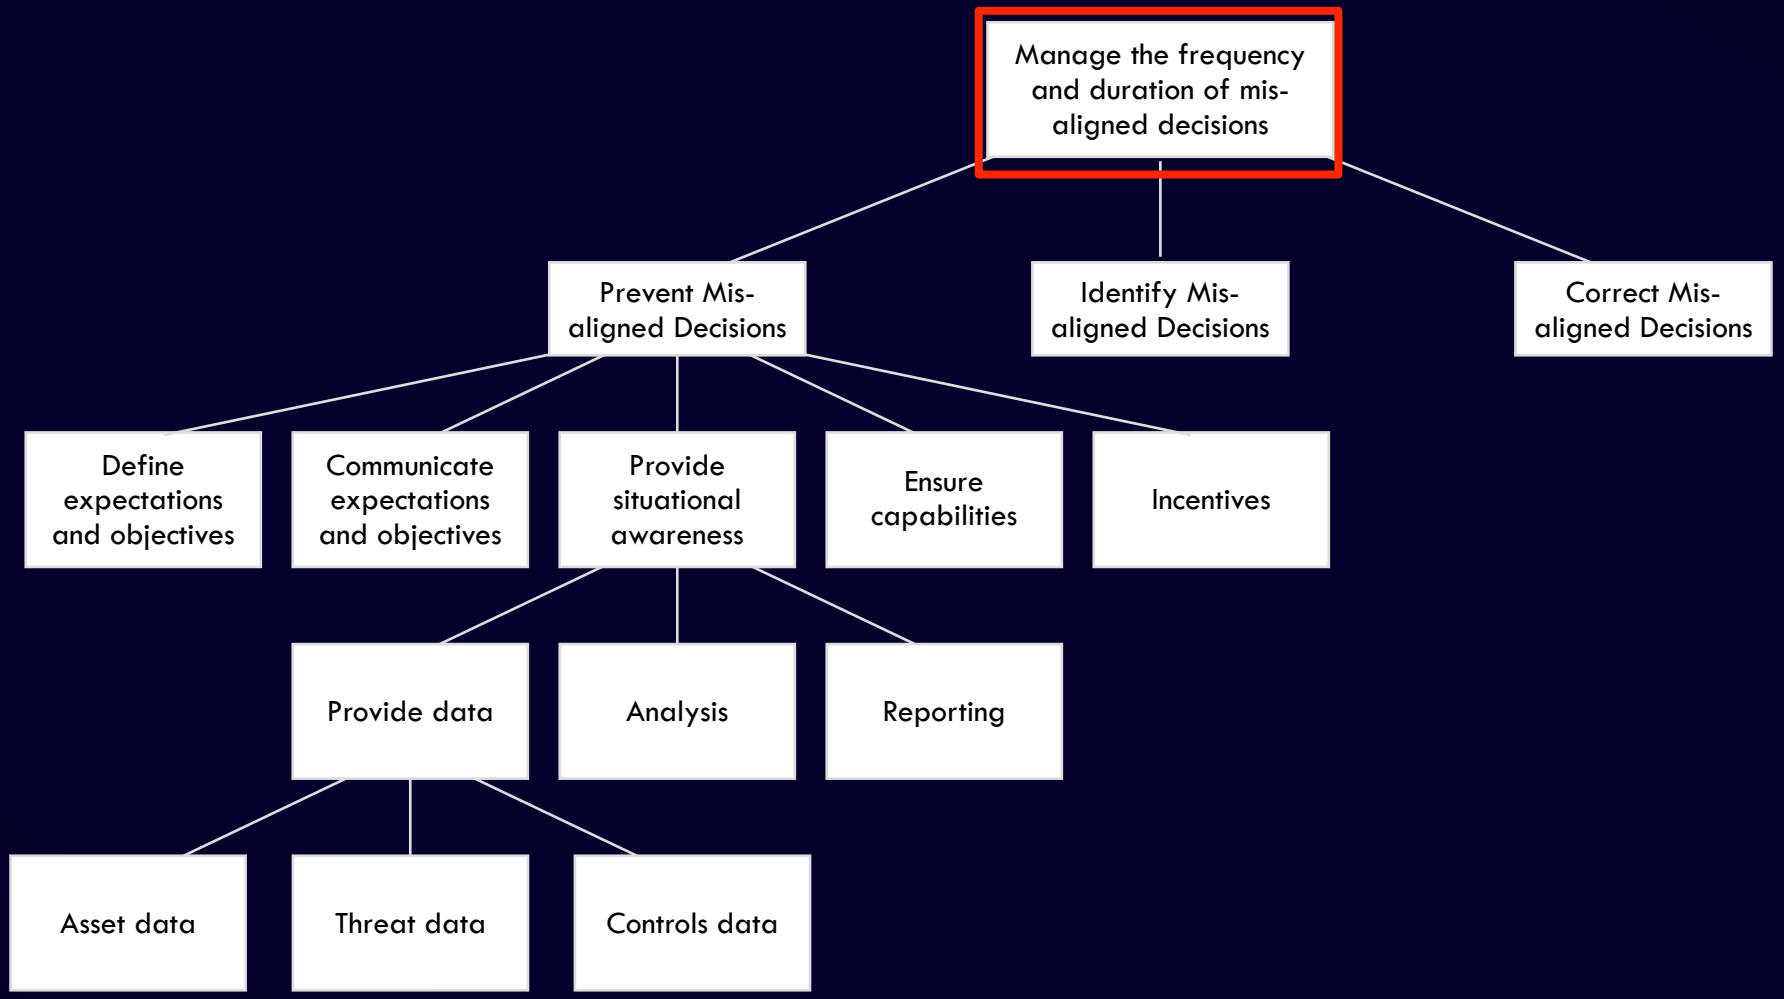

Common control frameworks focus on anatomy

CIS Controls 8.0

NUMBER	TITLE/DESCRIPTION	ASSET TYPE	SECURITY FUNCTION	IG1	IG2	IG3
3.4	**Enforce Data Retention** Retain data according to the enterprise's data management process. Data retention must include both minimum and maximum timelines.	Data	Protect	●	●	●
3.5	**Securely Dispose of Data** Securely dispose of data as outlined in the enterprise's data management process. Ensure the disposal process and method are commensurate with the data sensitivity.	Data	Protect	●	●	●
3.6	**Encrypt Data on End-User Devices** Encrypt data on end-user devices containing sensitive data. Example implementations can include: Windows BitLocker®, Apple FileVault®, Linux® dm-crypt.	Devices	Protect	●	●	●
3.7	**Establish and Maintain a Data Classification Scheme** Establish and maintain an overall data classification scheme for the enterprise. Enterprises may use labels, such as "Sensitive," "Confidential," and "Public," and classify their data according to those labels. Review and update the classification scheme annually, or when significant enterprise changes occur that could impact this Safeguard.	Data	Identify		●	●
3.8	**Document Data Flows** Document data flows. Data flow documentation includes service provider data flows and should be based on the enterprise's data management process. Review and update documentation annually, or when significant enterprise changes occur that could impact this Safeguard.	Data	Identify		●	●

"Protect", "Identify", etc. aren't specific enough to support empirical measurement and analysis. (They're also often misapplied.)

Dependencies between controls aren't accounted for.

Controls that fulfill multiple functions aren't accounted for.

Mapping

CIS Control Framework		Serves Which Functional Domain(s)			Loss Event Control Functions								
					Prevention			Detection			Response		
CIS Sub-Category	Title	LEC	VMC	DSC	Avoidance	Deterrence	Resistance	Visibility	Monitoring	Recognition	Event Termintion	Resilience	Loss Reduction
3.13	Deploy a Data Loss Prevention Solution	X		X		X	X	X	X	X			
3.14	Log Sensitive Data Access	X		X		X		X					

Many controls fulfill multiple risk reduction functions

Mapping challenges

PR.PT-4: Communications and control networks are protected	**CIS CSC** 8, 12, 15 **COBIT 5** DSS05.02, APO13.01 **ISA 62443-3-3:2013** SR 3.1, SR 3.5, SR 3.8, SR 4.1, SR 4.3, SR 5.1, SR 5.2, SR 5.3, SR 7.1, SR 7.6 **ISO/IEC 27001:2013** A.13.1.1, A.13.2.1, A.14.1.3 **NIST SP 800-53 Rev. 4** AC-4, AC-17, AC-18, CP-8, SC-7, SC-19, SC-20, SC-21, SC-22, SC-23, SC-24, SC-25, SC-29, SC-32, SC-36, SC-37, SC-38, SC-39, SC-40, SC-41, SC-43

Many controls are too broadly defined to enable empirical measurement.

Questions we need answers to…

1. How does the control affect risk?
2. How effective is it designed/intended to be?
3. Does it depend upon other controls in order to be effective?
4. How reliable is it?
5. What loss event scenario(s) is it relevant to?
6. Are other controls in place that also are relevant to the scenario(s)?

Loss Event Preventative Control Analysis

# of Assets	100
Contact Frequency	20

the population size of assets that are reachable by threat agents

the number of times per year that threat agents come into contact with the asset population

Controls	Toggle On/Off	Intended Performance	Reliability: Variance Frequency (per yr)	Reliability: Variance Duration (days)	Override?	Reliability: Variance Frequency (per yr)	Reliability: Variance Duration (days)	Aggregate Performance (per instance)
System configuration	Y	99%	2.0	7	N	0	0	
Anti-malware	Y	80%	2.0	14	N	0	0	
Admin restrictions	Y	90%						
URL filtering	Y	80%						
Allow/Disallow solution	Y	98%						

LEF	**0.07** annualized
	6.59% 12 month pro

Partial FAIR-CAM analysis example
Preventative Loss Event Controls

All of these values can be empirically measured.

Loss Event Preventative Control Analysis

# of Assets	100
Contact Frequency	20

the population size of assets that are reachable by threat agents

the number of times per year that threat agents come into contact with the asset population

Controls	Toggle On/Off	Intended Performance	Reliability: Variance Frequency (per yr)	Reliability: Variance Duration (days)	Override?	Reliability: Variance Frequency (per yr)	Reliability: Variance Duration (days)	Aggregate Performance (per instance)
System configuration	Y	99%	2.0	7	Y	2	30	**99.9888%**
Anti-malware	Y	80%	2.0	14	N	0	0	
Admin restrictions	Y	90%	0.5	30	N	0	0	
URL filtering	Y	80%	1.0	1	N	0	0	
Allow/Disallow solution	Y	98%	1.0	30	N	0	0	

LEF	**0.22** annualized
	22.33% 12 month probability of occurrence

FAIR-CAM's differentiating characteristics

- 24 functions (vs. NIST CSF's 5)
- Accounts for direct and indirect effects on risk
- Accounts for relationships and dependencies between controls
- Accounts for controls that affect risk in more than one way (context sensitivity)
- Defines units of measurement for control functions
- Enables empirical measurement of control efficacy and risk reduction value
- Complements existing control frameworks

Key take-aways…

- The controls landscape is complicated and highly nuanced. If an analysis doesn't account for controls physiology (especially automated analyses), the results are unlikely to be accurate.
- FAIR-CAM describes how controls affect risk and defines units of measurement.
  - Enables empirical measurement and validation of control efficacy and value
- FAIR-CAM is one of the keys to enabling automation as well as the use of AI.
- FAIR-CAM is an open standard (covered by a creative commons license).

#3 - Open, non-proprietary models

It’s a matter of trust

Why is proprietary encryption considered dangerous?

What if NIST 800-30 was a proprietary model?

And by the way, just because something is patented, that doesn’t mean it works.

Key take-away…

- There are many ways to get risk measurement wrong, and very few ways to get it right.
- Consequently, if we can’t examine it, we shouldn’t trust it.

#4 - Automation and AI

Automated risk measurement…

Anti-malware data
CIS controls scores
CVSS scores

Threat intelligence services
Organization's own logs

NIST CSF scores
System config data
Program maturity scores

Verizon DBIR
Advisen
Insurance providers

Which of these is more important?

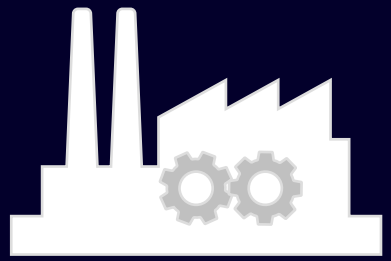

PR.AC-1: Identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users and processes

PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy

Is it twice as important? Three times…?

Control relevance is context sensitive!

How does a control affect risk?

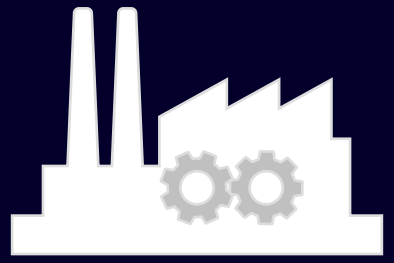

Does logging affect likelihood, or magnitude?

PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy

Control dependencies…

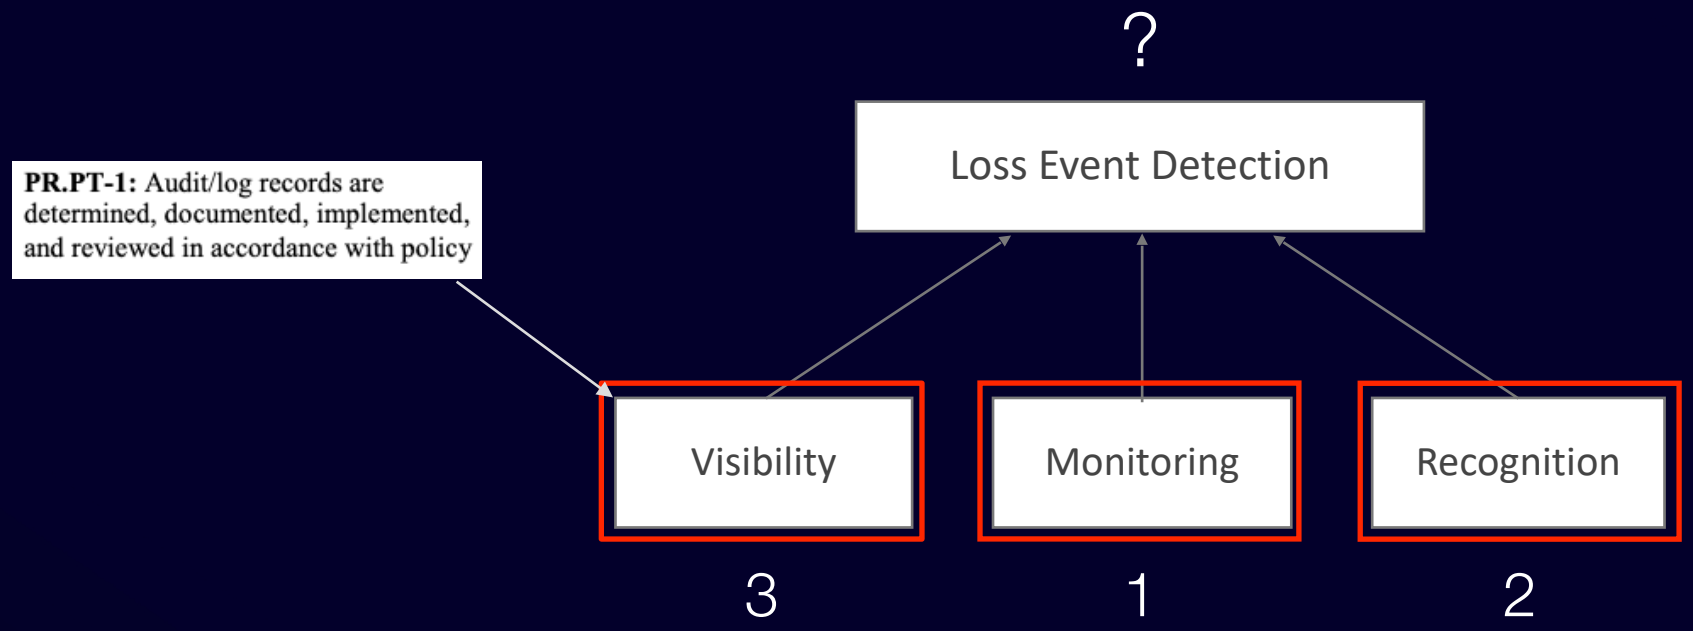

VISIBILITY: There has to be data that contains evidence of a breach (e.g., logs)

MONITORING: Someone or something has to review the data (e.g., manual reviews, SIEM, etc.)

RECOGNITION: Exploit signatures, malware signatures, baselines of normal activity, etc.

Key take-aways…

- Control relevance is highly context sensitive.
- Controls often have key dependencies with other controls.
- If automation fails to account for these (and other nuances), analysis results will be inaccurate.

Artificial Intelligence

What's the difference between automation and AI?

The analytic model is learned rather than designed, which means it's only as good as its training.

Training AI — it's all about the data

A risk measurement AI is supposed to estimate the likelihood and magnitude of loss from the scenario data it's presented with.

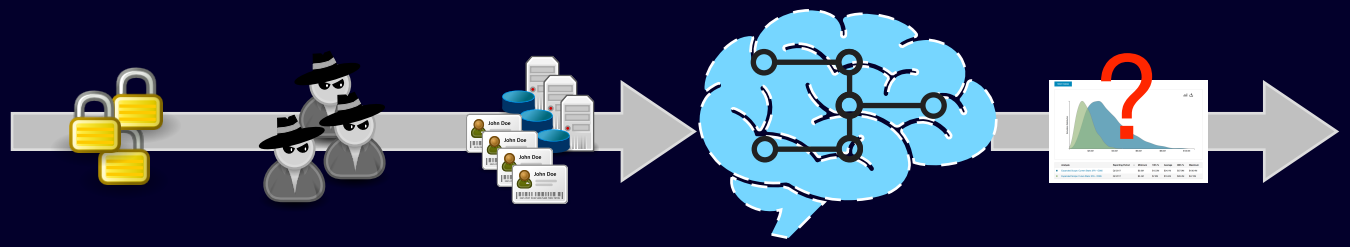

Its simulated neurons and synapses get reinforced when the estimate is correct.

Simplified training example

- Inputs:
  - Asset: Sensitive customer information
  - Threat agent: Disgruntled insider
  - Event type: Data leakage
  - Method: Data exfiltration
  - Vector: E-mail
  - Controls in place: Authentication, access privileges, anti-malware, logging, etc.
- AI output:
  - Likelihood: 10%-15%
  - Impact: $5M-$20M

What is the AI's estimate going to be checked against?

Either empirical loss data has to exist for that scenario, or someone has to have done the analysis beforehand.

The problem of training bias

- Bias exists when an AI makes inappropriate decisions due to poor training data
- It is one of the most prevalent and difficult to manage dimensions of AI

https://www.technologyreview.com/2019/02/04/137602/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/

The problem of opaqueness

- The ability to audit and analyze how an AI arrived at a particular result is non-trivial.
- Unless the AI is auditable, it should be considered just as prone to failure as proprietary models.

Key take-aways…

- The quality of AI is dependent upon the quality of training
- Bias and opaqueness are serious challenges that need to be addressed before AI can be trusted

#5 - Risk analysis professionals

Who ya gonna trust?

- Who's allowed to measure risk where you work?
- Have they been trained and certified in risk analysis?
- What methods do they use?
- Are they critical thinkers?

Key take-away…

Risk analysis and measurement should be considered a distinct discipline, just as forensics, penetration testing, DevSecOps, and others are.

Wrapping up

Overall key takeaways…

- “Good” risk measurement is defensible
- In order to evolve and mature, we have to recognize and correct what isn’t working
- The future of risk measurement will:
  - Be quantitative
  - Account for controls physiology
  - Rely on open, non-proprietary models
  - Where possible, leverage automation and AI
  - Be performed by trained and certified professionals
- In order for automation and AI to generate results we can trust, they can’t rely on the methods of the past and present

Questions?